



Beyond productivity: AI and NHS workforce implications

June 2026

Dr Shaaz Mahboob, NHS England (London)
Ambra Caruso, HIN South London
Shilpi Shobowale, HIN South London
Richard Stockley, Imperial College Health Partners
Mandy Trinh, UCLPartners
Ilias Zapantis, UCLPartners

Acknowledgments

This work was commissioned by NHS England (London), which identified the workforce implications of non-clinical AI and ambient voice technologies as an important area for exploration. The three London Health Innovation Networks (UCLPartners, HIN South London and Imperial College Health Partners) were commissioned to undertake this work as a collaborative project.



Contents

Acknowledgments	2
Foreword	5
Executive summary	7
Key insights	9
Key recommendations	10
Introduction	11
Scope and aim	12
Methodology	13
Workforce implications of AI adoption	15
1. AI adoption is outpacing organisational readiness and is driven by operational pressures.....	16
Adoption is driven by workforce pressures rather than strategic transformation.....	17
2. AI is redistributing work, not reducing it.....	18
AI is redistributing effort towards oversight rather than reducing workload.....	19
Roles, responsibilities and accountability are constantly evolving.....	20
Implications on workforce structure and planning.....	20
3. Workforce planning, requirements and formal structures have not yet adapted to the rapid pace of AI adoption.....	21
Capability requirements are changing faster than workforce development.....	22
Professional judgement is becoming a critical workforce skill.....	22
Formal workforce structures are lagging behind AI-enabled work.....	23
Leadership, culture and trust influence workforce readiness.....	24
4. Workforce engagement is essential to ensure benefits are maximised and harms are avoided.....	25
Governance has not kept pace with adoption.....	25
Trust, clarity and permission influence workforce adoption.....	25
Workforce benefits are not consistently captured.....	27
Evaluation frameworks do not reflect workforce value.....	27

5. Fragmented and ad-hoc AI adoption limits workforce benefits	28
Fragmented local adoption creates uneven workforce experience	29
Lack of ownership and coordination limits workforce benefits	29
Scaling requires workforce and organisational alignment	30
Next steps and recommendations	32
Place the workforce at the centre of AI adoption	33
Embed AI into workforce planning and role design	33
Shift from one-off training to continuous capability building	33
Clarify roles, responsibilities and accountability	34
Align governance, ownership and “permission to use” with real-world practice	34
Redefine how productivity gains, value and benefits are measured to reflect workforce realities	35
Future proof the workforce and progression pathways	35
Enable system-level shared learning to improve equitable adoption of AI	36
Conclusion	37

Disclosure

This report was synthesised and produced with the assistance of generative AI under human direction, with the final review and assurance of the authors, who accept ownership and responsibility for the contents of the submission.

Foreword



Dame Caroline Clarke
NHS England (London)

The NHS is entering a defining period of transformation. Across London, and throughout the wider health and care system, artificial intelligence and digital technologies are reshaping how care is delivered, how services are organised and how our workforce is supported.

The pace of change is significant, and while the opportunities are considerable, so too is the responsibility to ensure these technologies are introduced thoughtfully, safely and in ways that strengthen both patient care and workforce experience.

This report on the impact of non-clinical AI, including ambient voice technology (AVT), on the NHS workforce in London comes at an important moment. AI is no longer a future consideration for the NHS. As this work demonstrates, it is already being adopted across organisations and teams, often in response to very real operational pressures. Staff are increasingly using AI-enabled tools to manage administrative burden, support documentation and improve workflow efficiency, helping to release more time for patient care. In many cases, this adoption is happening organically and at pace, driven by frontline teams seeking practical solutions to day-to-day challenges.

Importantly, this report moves beyond the broader narrative around AI and productivity to examine what these changes mean in practice for the workforce. Drawing on both published evidence and direct insights from NHS staff and leaders across London, it provides a grounded and credible assessment of how AI is already beginning to reshape work, roles and organisational culture.

The findings highlight that AI is not simply replacing tasks, but redistributing work, with increasing emphasis on oversight, validation, professional judgement and new forms of decision making. They also highlight that many of these changes are emerging ahead of formal workforce planning, governance arrangements and capability frameworks.

This makes the report particularly valuable. It brings together practical experience from real-world implementation with wider system learning, helping to identify not only the opportunities AI presents, but also the risks and tensions that require careful attention. The report reinforces that successful adoption will depend not only on technology itself, but on leadership, workforce engagement, trust and clear governance. It also highlights the

importance of ensuring staff feel confident, supported and empowered to use these tools safely and effectively.

London is uniquely placed to lead this next phase of innovation and transformation. Our region brings together world-class clinical expertise, academic excellence, diverse communities and a strong ecosystem of data, digital and life sciences partners. Through collaboration across NHS organisations, research institutions and system partners, we have an opportunity to shape a model of AI adoption that is responsible, equitable and grounded in the realities of frontline care and operational delivery.

AI has the potential to help build a more sustainable, productive and patient-centred NHS, one that better supports both the people who deliver care and the communities we serve.

Dame Caroline Clarke, NHS England (London)

Executive summary

The NHS has entered a period of rapid technological and organisational transformation. Artificial intelligence (AI), including non-clinical AI applications and ambient voice technology (AVT), is rapidly being adopted, often informally, across health and care settings to support administrative work, operational processes and clinical documentation.

While much of the wider discussion on AI has focused on clinical applications and productivity gains, less attention has been given to the implications for the workforce, particularly in relation to how work is changing in practice, how roles and responsibilities are evolving, and what organisational conditions are required to support safe and sustainable adoption. The pace and scale of adoption mean that workforce implications are no longer theoretical or future-facing but are already emerging across NHS organisations and teams.

This report presents the findings of a collaborative programme of work commissioned by NHS England (London) and delivered by the three London Health Innovation Networks (UCLPartners, HIN South London and Imperial College Health Partners). The work examines the implications of non-clinical AI and AVT for the NHS workforce in London. It draws on a structured review of published and grey literature alongside insights from multidisciplinary roundtable discussions involving NHS leaders, clinicians, operational teams and workforce representatives from across London. Together, these sources provide a grounded assessment of how AI is currently being used, how it is experienced by staff and the emerging implications for workforce planning, organisational culture, governance and service delivery.

The report finds that AI adoption is outpacing organisational readiness, often driven by frontline operational pressures rather than coordinated strategic transformation. Staff are already integrating AI into day-to-day administrative, operational and documentation tasks, frequently without consistent governance or alignment with workforce strategy. This has created a growing disconnect between how work is performed in practice and how roles, responsibilities and workforce models are formally defined and supported.

While this report focuses on AI and AVT usage in non-clinical functions, it is important to recognise that it is not possible to separate AI and AVT usage into clinical and non-clinical applications. Many administrative and operational processes directly influence clinical care and workflows, with the capacity for both benefit and harm. As a result, AI deployment in non-clinical domains carries both opportunities and risks for patient care, not solely for efficiency. This interdependency requires a whole-system perspective. Changes to administrative processes may improve clinical capacity, continuity of care and workforce

experience, while failures in governance or oversight may have downstream implications for safety and quality. Workforce, operational and clinical impacts must therefore be considered together, rather than in silos.

A central finding of the report is that AI is not simply reducing work, but redistributing it. While AI tools may streamline specific administrative tasks, improve workforce experience and enable better work-life balance, they also create new forms of work associated with oversight, validation, judgement and exception handling. Staff increasingly act as reviewers and decision-makers alongside AI-generated outputs, requiring different skills and forms of professional judgement. This challenges assumptions that AI will automatically deliver straightforward productivity gains and highlights the need to reconsider workload, performance expectations and workforce planning approaches. Where these additional responsibilities are not recognised, there is a risk that work becomes intensified rather than reduced.

At the same time, the report identifies a broader set of workforce and organisational risks that require greater attention. Entry-level and transactional roles are increasingly exposed to automation, raising concerns about workforce progression, skills development and the long-term sustainability of workforce supply. Adoption remains fragmented and uneven, with variation in access to tools, capability and governance creating inequities in workforce experience. Governance frameworks have not kept pace with real-world use, resulting in uncertainty around accountability, acceptable use and data handling. In some cases, staff are already using AI informally without clear organisational guidance, creating potential safety, legal and reputational risks. Without a more coordinated and workforce-centred approach, these risks may limit the benefits of AI and exacerbate existing workforce pressures.

Beyond these workforce risks, governance and accountability arrangements are also struggling to keep pace with real-world implementation. The report highlights growing uncertainty around acceptable use, ownership and responsibility, particularly where AI use is occurring informally or outside formal organisational structures. Existing approaches to evaluating impact are similarly underdeveloped and often fail to capture workforce-related outcomes such as workload sustainability, staff wellbeing and retention.

Taken together, these findings indicate that AI is already reshaping the workforce through changes in task profiles, skill requirements and role structures, but largely outside formal workforce planning processes. Without deliberate action, there is a risk that AI reinforces existing workforce pressures, widens variation in experience and delivers only partial benefits.

To address this, a coordinated and workforce-centred approach is required. This should include embedding AI into workforce planning and role design, shifting from one-off training to continuous capability development, strengthening governance and clarifying accountability, and reframing productivity to include workforce sustainability and wellbeing. It will also be critical to protect entry pathways and progression routes, enable system-wide coordination and shared learning, and build staff trust through transparent leadership and meaningful engagement.

The NHS is now at a critical inflection point. AI is already transforming how work is carried out, but workforce structures, governance and capability have not kept pace. With strong leadership, workforce engagement and a whole-system approach that recognises the interdependence between operational and clinical care, AI has the potential to support a more sustainable, equitable and effective NHS workforce.

Key insights



1. AI adoption is outpacing organisational readiness, limiting workforce benefits

AI adoption is accelerating across the NHS, often driven by operational pressures and frontline need rather than organisational readiness. In many cases, AI is being used informally and without fully established safeguards. As adoption accelerates, governance arrangements, workforce models, organisational structures and evaluation frameworks are struggling to keep pace. This is creating a widening gap between how work is evolving in practice and how it is formally supported, governed and planned.



2. AI is redistributing work, not reducing it

There appears to be a broad assumption that AI will lead to substantial productivity gains. However, the evidence that it will deliver transformative productivity gain in healthcare systems, in the short and medium term, is limited. Rather than delivering straightforward productivity gains, AI is redistributing effort towards judgement, validation and oversight. This challenges assumptions about efficiency and highlights the need to re-evaluate workload, capacity and performance expectations.



3. Workforce planning, requirement and formal structures have not yet adapted to the rapid pace of AI adoption

AI is a powerful technology, with potentially profound impacts in the healthcare sphere, and which is already available directly to healthcare workers for use of personal devices. The implications for roles, skills, career pathways and workforce progression are already emerging, particularly for entry level and administrative roles. However, workforce planning, training and role design have not yet adapted to these changes in a systematic way.



4. Workforce engagement is essential to ensure benefits are maximised and harms are avoided

The deployment of past digital technologies in the NHS have suffered from inadequate workforce engagement and a lack of opportunity for healthcare staff to participate in co-creating applications. Adoption is shaped not only by technology, but by whether staff feel confident, permitted and supported to use it. Clear governance, defined ownership, workforce engagement, and alignment with real workflow transformation are essential to move from fragmented local use to coordinated system level impact.

Key recommendations



1. Place the workforce at the centre of AI adoption

Embed workforce considerations into implementation through co-creation, and integrate AI into workforce planning, role design and capability development.



2. Align governance, ownership and “permission to use” with real-world practice

Ensure that governance frameworks reflect how AI is used in practice, providing clarity on accountability, acceptable use and decision-making responsibilities.



3. Redefine how productivity gains, value and benefits are measured to reflect workforce realities

Change how we define and measure the benefits of AI so it better reflects the reality of staff experience. This means looking beyond traditional efficiency measures and also considering factors such as manageable workloads, staff wellbeing, and other less tangible improvements, not just time or cost savings.



4. Future proof the workforce and progression pathways

Enable workforce transformation by focusing on maintaining meaningful entry level roles, adapting training and education pathways, and ensuring that AI adoption supports, rather than disrupts, career development.



5. Enable system-level shared learning to improve equitable adoption of AI

Strengthen mechanisms for system wide and cross-organisational learning, collaboration and evidence sharing to support more consistent and equitable adoption across the workforce.

Introduction

The NHS is facing a set of unprecedented and deeply interconnected systemic challenges that are reshaping both the delivery of care and the context in which its workforce operates.

Demand continues to rise, driven by an ageing population and the growing prevalence of long-term conditions, with around a quarter of the population accounting for **approximately 65% of NHS spending**. At the same time, access to care is deteriorating, with long waiting times and low levels of patient satisfaction, signalling increasing difficulty in meeting timely demand. Financial strain, rising costs and inefficiencies further challenge fiscal sustainability, while an outdated, hospital centric model continues to divert resources away from community and preventative care, contributing to poorer outcomes and higher costs. Progress is also constrained by legacy IT infrastructure, fragmented systems and limited interoperability, which increase administrative burden and restrict the potential for productivity gains.

These pressures are compounded by workforce shortages, burnout and productivity levels that remain **20-25% below pre-pandemic levels**. In parallel, there is a growing misalignment between current workforce skills and the requirements of future models of care. Without significant workforce redesign and improvements in productivity, meaningful system transformation and benefits realisation will not be achievable. Taken together, these factors create an urgent and unavoidable need for systemic change, with innovation adoption needing to be explicitly centred around workforce.

Artificial intelligence (AI) has been identified in “The 10 Year Health Plan for England” as one of the key technological advancements with the potential to address these challenges. As one of the “five big bets”, AI is expected to accelerate healthcare transformation, support productivity gains and improve workforce experience, enabling care teams to focus more on patient care.

While much of the policy and public discourse has focused on clinical applications, the use of AI in non-clinical settings across the NHS is expanding rapidly. AI solutions that support administrative work, operational planning and service management are increasingly being deployed as part of broader efforts to manage demand, improve efficiency and relieve pressure on overstretched staff. There is growing recognition that such applications may play a critical role in shaping how work is organised, how roles evolve and how services function. However, there remains limited clarity on what this means in practice for the NHS workforce, particularly within corporate and operational functions such as administration, workforce management and clinical documentation support, as well as in the use of ambient voice technologies (AVT) within clinical environments.

This raises important questions about how roles and responsibilities may change, what new skills will be required, how staff experience and trust can be supported, and what governance structures are needed to ensure safe, effective and empowering use. Without deliberate focus on these issues, there is a risk that AI adoption may fail to realise its intended benefits or could exacerbate existing system pressures.

To explore these challenges and inform practical next steps, NHS England (London) is working in partnership with UCLPartners, HIN South London, and Imperial College Health Partners to examine the impact of non-clinical AI, including the use of AVT, on the NHS workforce in London. This includes consideration of roles and responsibilities, skills, capability building and the cultural changes required to ensure that AI adoption meets workforce needs.

This report brings together evidence from a structured programme of work, including a desktop review of published and grey literature on the adoption of non-clinical AI, with a particular focus on workforce implications, organisational processes and system level functioning. It also draws on insights from two roundtable discussions: one focused on the use of AI within corporate, human resources and finance functions, and a second exploring workforce experiences of using AVT. These discussions were designed to capture the perspectives of NHS staff who are actively engaging with AI enabled tools, providing grounded insights into current practice and emerging implications for the workforce.

Scope and aim

The scope of this review is deliberately focused on non-clinical AI applications, including those supporting administrative, operational, and scheduling functions, as well as the use of AVT to support clinical operations.

Diagnostic AI and clinical decision support tools are explicitly out of scope. The primary geographic focus is England, with particular relevance to the London context.

This work seeks to deepen understanding of the implications of AI for the workforce and to inform future workforce planning, policy and decision making. It aims to support more informed and pragmatic decision making by identifying key themes, tensions and considerations for leaders and policymakers.

This report is not intended as a technical assessment of AI tools or an implementation guide, nor does it seek to promote specific technologies or approaches. Instead, it adopts a workforce and system perspective, examining how AI is being introduced into existing roles, processes and workflows, how responsibilities and accountabilities are evolving as a result, and where pressures or unintended consequences may emerge over time.

Practical recommendations are intended to support policy development, workforce planning and organisation learning in the short to medium term, while also identifying opportunities for future collaboration, deeper exploration and continued development as AI adoption across the NHS evolves.

Methodology

This report adopts a structured, multi-method approach to examine the implications of non-clinical AI and ambient voice technologies (AVT) for the NHS workforce.

It draws on both published evidence and stakeholder insights, combining a targeted review of literature with findings from two multidisciplinary roundtable discussions. The distinction between clinical and non-clinical applications of AI follows an adaptation of the classification used in a Health Foundation report on non-clinical applications of artificial intelligence in the NHS.

The literature review was developed collaboratively by NHS England (London), UCLPartners, HIN South London and Imperial College Health Partners. An initial core set of sources was supplemented through targeted searches of UK healthcare literature, relevant examples from the wider UK public sector and selected international healthcare sources.

In parallel, two roundtables were convened, one focused on corporate and operational applications of non-clinical AI and the other on the use of AVT. Participants represented a broad range of NHS perspectives, including executive leadership, digital and transformation teams, clinicians, operational managers, workforce leaders and corporate functions such as finance, procurement and governance.

A structured evidence matrix and Generative AI large language models were used to synthesise findings across both the literature review and roundtable discussions, enabling direct comparison between published evidence and real-world implementation experience. This approach supported the identification of recurring themes, areas of convergence and important gaps between policy, research and practice.

The findings should be interpreted as indicative rather than definitive. They provide a contemporary assessment of how non-clinical AI and AVT are being adopted across the NHS and highlight emerging implications for workforce planning, capability, governance and organisational development.

Workforce implications of AI adoption



Workforce implications of AI adoption

Five overarching themes emerge from across the literature review and the roundtable discussions, providing a structured view of how non-clinical AI and AVT are shaping the NHS workforce.



1.

AI adoption is outpacing organisational readiness, limiting workforce benefits



2.

AI is redistributing work, not reducing it



3.

Workforce planning, requirement and formal structures have not yet adapted to the rapid pace of AI adoption



4.

Workforce engagement is essential to ensure benefits are maximised and harms are avoided



5.

Fragmented and ad-hoc AI adoption limits workforce benefits

In many areas, the roundtables reinforced existing findings from the literature, particularly in relation to uneven adoption, governance challenges and the importance of workforce capability. However, they also provided important nuance and, in some cases, challenged prevailing assumptions in policy and research. This was particularly evident in relation to productivity, where expected efficiency gains were not always realised as anticipated, the nature of skills required to work effectively with AI, and the day-to-day experience of staff as they adapt to new tools and workflows. Furthermore, the roundtables provided a more practical and immediate account of how AI is being used and experienced in real-world settings, highlighting issues around informal use of AI, permission, confidence, access, workflow and the additional judgement staff needed to review AI outputs.

Taken together, these themes reflect both the current state of adoption and the emerging implications for workforce roles, responsibilities and ways of working. They highlight a system in transition, where the impact of AI is already being felt in practice, but where organisational structures, workforce models and supporting frameworks have yet to fully adapt.

1. AI adoption is outpacing organisational readiness and is driven by operational pressures

The adoption of non-clinical artificial intelligence across the NHS is being driven primarily by immediate operational pressures rather than by a coordinated, system-wide strategy for transformation. Evidence from both the literature and stakeholder discussions indicates that organisations and individuals are turning to AI in response to workforce constraints, rising demand and the need to manage increasing administrative burden. For the workforce, this positions AI as a practical solution to workload pressures, particularly in administrative tasks, where the burden has become a key driver of adoption. As a result, uptake is often led by individuals and teams seeking to address immediate challenges, rather than being implemented as part of a planned approach to workforce or service transformation.

In this context, AI is largely experienced as an efficiency tool rather than a transformation lever. Current use cases are predominantly focused on discrete, task-based applications such as documentation, correspondence, information retrieval, scheduling and operational planning. These applications target repetitive and time-consuming activities, offering immediate and tangible benefits at a local level. However, from a workforce perspective, this means that the impact of AI is concentrated on specific tasks rather than on broader changes to roles, workflows or service models.

Adoption is driven by workforce pressures rather than strategic transformation

A consistent finding across all sources is that adoption remains uneven, ad-hoc and, in some cases, informal. Implementation is frequently initiated at the level of individual teams or departments, often without a unifying organisational or system-wide framework. In parallel, there is growing evidence, particularly from the roundtables, of widespread informal use of generative AI tools by staff to support day-to-day work. This suggests that adoption may be more advanced in practice than is formally recognised, with use in some cases outpacing organisational visibility, governance and oversight. For the workforce, this creates a gap between how AI is used in practice and how it is formally supported, with staff often navigating new tools without clear guidance or consistent support.

The roundtable discussions further highlighted a strong operational “pull” from the workforce, with AI increasingly seen not only as a means of improving efficiency but as a necessary response to sustained workload pressures. Participants from the corporate roundtable noted that:

“ people who are already under the most pressure are the ones jumping on AI first, because they are looking for practical ways to get through the work. At the other end, however, there are still staff who do not want to engage with it, or who lack the confidence to try. ”

This illustrates an immediate and practical need, as staff seek practical solutions to manage demand and maintain performance. Consequently, adoption is driven as much by frontline demand as by organisational strategy, reinforcing a predominantly bottom-up pattern.

This pattern contributes to significant variation in access, maturity and use across organisations and teams. Differences in digital infrastructure, leadership prioritisation, governance arrangements and workforce capability mean that some areas are actively experimenting with and embedding AI into workflows, while others remain at an early stage of awareness or exploration. Evidence from the roundtables suggested that this variation exists both across organisations and within them, with teams adopting different approaches based on local leadership, risk appetite and access to tools. For the workforce, this results in uneven access to AI-enabled support, with some staff benefitting from efficiencies and workload relief, while others continue to operate without these advantages. As participants from the corporate roundtable observed:

“ Access to licences is already creating an uneven playing field. Some people have the tools and are starting to build them into daily work, while others want to use them but are waiting for access, capacity or support to catch up. ”

Important differences also emerge between types of non-clinical AI application. In corporate and administrative contexts, adoption tends to be more decentralised, informal and user-led, often involving general-purpose tools used at the level of individual tasks with limited integration into core systems. In contrast, the adoption of AVT is more likely to be organisationally sanctioned, implemented through structured programmes and embedded within clinical workflows. This creates differing workforce experiences, where some staff operate within supported and governed environments, while others rely on informal use and self-directed learning.

The absence of a coordinated approach presents challenges for oversight and alignment. Where adoption is dispersed, it becomes more difficult to ensure consistent governance, evaluate impact systematically or align AI use with broader organisational priorities. For staff, this can lead to duplication of effort, inconsistent expectations and limited opportunities to learn from others or benefit from shared approaches.

Overall, the current pattern of adoption reflects a system responding pragmatically to immediate pressures, with AI being used to support operational delivery rather than to a deliberate workforce design. While this has enabled early progress and demonstrated clear potential, it also results in uneven workforce experience and fragmented support. Without a more coordinated approach, there is a risk that the benefits of AI remain localised, with variation in access, capability and impact across the workforce, limiting its potential.

2. AI is redistributing work, not reducing it

The current impact of non-clinical AI across the NHS is most evident at the level of individual tasks and workflows, rather than at the level of whole roles or service models. Across both the literature and stakeholder discussions, there is consistent evidence that AI is beginning to reshape day to day administrative work, as AI tools are being used to draft correspondence, summarise information, support clinical documentation and streamline routine processes, enabling staff to complete existing activities more efficiently.

In this context, AI primarily functions as a form of work augmentation, enhancing human productivity rather than transforming it. While these changes are meaningful at the level of daily activity, there is limited evidence of fundamental role redesign or pathway transformation. AI is largely being layered onto existing processes rather than used to reconfigure how services are delivered. As a result, the structure of roles, team configurations and care pathways remain broadly unchanged, even where AI is being actively used. This reinforces a pattern of incremental improvement rather than systemic transformation. Unless whole pathway or system services are redesigned with AI capabilities as a key enabler for transformation, any benefits will remain limited and short-termed.

AI is redistributing effort towards oversight rather than reducing workload

However, insights from the roundtables suggested that the lived experience of this change may be more pronounced than is reflected in the literature. Participants described tangible shifts in how time is used, how tasks are prioritised and how work is distributed within teams. In some cases, AI is already altering the balance between administrative and higher value activities, even if formal role definitions have not yet caught up. This indicates that while structural transformation is limited, functional changes to roles are already underway in practice, creating a lag between how work is performed and how roles are formally described or understood.

A further nuance emerging from both the literature and roundtables is that AI is not necessarily reducing workload but redistributing it. While AI is often positioned as a means of streamlining processes and improving efficiency, the evidence suggests that it is more accurately shifting effort towards new types of activity. As participants from the AVT roundtable noted:

“ The question is whose burden is being reduced. Notes and letters may be generated more quickly, but the clinician still has to read them carefully, understand what the AI has written and check for errors. In some cases, that creates a different kind of cognitive burden rather than removing the work. ”

The desktop review supports this, showing that work is increasingly moving towards tasks such as checking, validation, judgement and exception handling, particularly within administrative, corporate and enabling roles.

Crucially, much of this new work remains largely invisible within formal workforce structures. Activities such as reviewing outputs, correcting errors and managing ambiguity are not consistently recognised in job descriptions, workload models or performance expectations. As a result, staff may take on additional responsibility without corresponding recognition or adjustment in how work is defined and measured.

This was reflected strongly in the roundtables, where participants frequently described AI as acting as a “first author”, with staff taking on the role of reviewer or overseer. One participant from the AVT roundtable clearly noted that:

“ The clinician is no longer always the first author of the note. The AI produces the first version, and the clinician becomes the reviewer, editor and quality assurer. That is a different skill, because you have to spot what is missing, what is wrong and what still needs clinical judgement. ”

In practice, this means that time saved in generating outputs can be offset by the time required to verify their accuracy, particularly where errors are subtle or context specific. As a result, the overall workload may not reduce as expected but instead becomes reconfigured. Furthermore, staff are having to learn and use evaluative skills to ensure accuracy and amend any errors.

This redistribution of effort has important implications. Where it is not recognised in workload planning and performance expectations, there is a risk that AI may intensify work rather than alleviate it, potentially undermining confidence in productivity claims. It also reinforces the importance of developing skills in judgement, evaluation and critical review, linking closely to the capability challenges identified in other themes.

Roles, responsibilities and accountability are constantly evolving

This evolving reality contributes to a lack of clarity, as the implications for roles, responsibilities and accountability remain emerging but not yet well defined. As AI begins to take on elements of tasks previously carried out by staff, questions arise around ownership, oversight and professional responsibility. This is particularly relevant in areas such as documentation and communication, where outputs generated or supported by AI continue to require human validation. The literature underlines these issues at a conceptual level, but the roundtables suggested that staff are already navigating these ambiguities in practice, often without clear guidance or consistent expectations. Responsibilities are frequently absorbed into existing roles without formal redefinition, leading to variation in how accountability is interpreted and applied.

Differences are also evident between types of AI application. In corporate and administrative contexts, the impact on work tends to remain task specific and individually experienced, with AI supporting discrete activities such as drafting, summarisation or data handling. These uses, while valuable, are less likely to alter team structures or workflows in a visible way. In contrast, AVT are having a more immediate and embedded impact on clinical workflows, particularly in relation to documentation and consultation processes. Evidence from the roundtables indicated that AVT is beginning to reshape how clinical time is organised, how attention is allocated during patient interactions and how documentation is produced and reviewed. In some cases, this is changing the experience of both clinicians and patients, even if the formal structure of roles remains unchanged.

Implications on workforce structure and planning

Beyond these task level changes; both the literature and the roundtables point to emerging implications for workforce structure and planning. The desktop review highlights early but underdeveloped evidence of redistribution of effort across roles, with uneven impacts depending on the nature of work. The roundtables added further depth, with participants describing entry level and more transactional roles as being most exposed to automation. In many cases, organisations are managing this change gradually through natural turnover rather than active redesign.

While this approach may reduce the need for immediate structural change, it raises longer-term concerns. Participants highlighted the potential for shrinking entry level roles, which could weaken traditional entry points into the workforce and disrupt established training and progression pathways. There were also concerns about the risk of deskilling, particularly where routine tasks that previously supported skill development is increasingly automated. These dynamics point to a tension between short term efficiency gains and longer-term workforce sustainability.

The discussions also emphasised the importance of staff morale and retention as key outcomes of AI adoption. Participants noted that the impact of technology on workload, job satisfaction and the ability to sustain performance over time should be considered alongside more traditional measures of productivity. In this context, the question of whether AI helps to support and retain the workforce becomes a critical consideration in its adoption.

Taken together, these findings suggest that AI is no longer a future workforce issue, but one that is already beginning to reshape workforce structures and dynamics. While the literature often frames workforce implications as emerging, the roundtables indicated that these impacts are already materialising in practice. Without more proactive workforce planning, there is a risk that these changes exacerbate existing shortages, create new inequalities across roles and undermine longer term system resilience.

Overall, the evidence suggests that AI is beginning to reshape how work is carried out, but not yet how work is formally organised. While efficiencies are being realised and new ways of working are emerging, the shift from incremental task level change to meaningful role and service redesign has not yet been fully realised. This emphasises the need for a more deliberate focus on how roles, responsibilities and workforce structures should evolve alongside AI adoption. Without this, there is a risk that changes to work remain informal, uneven and insufficiently supported, limiting the extent to which the full benefits of AI can be realised across the system.

3. Workforce planning, requirements and formal structures have not yet adapted to the rapid pace of AI adoption

The adoption of non-clinical AI and AVT is already shaping workforce dynamics across the NHS, with implications for skills, capability and organisational culture. Across both the literature and stakeholder discussions, there is consistent evidence that the impact of AI extends beyond how tasks are performed to influence how staff engage with their work, the types of skills required, and the expectations placed on different roles. These changes are contributing to variation in capability, confidence and experience across the workforce. As a result, AI is contributing to a shift in workforce requirements and behaviours that is not yet fully reflected in formal workforce structures, training approaches or organisational frameworks.

Capability requirements are changing faster than workforce development

A recurring theme is the need for new and evolving skill sets. The literature emphasises capabilities such as digital literacy, data awareness and the ability to critically assess AI generated outputs. Roundtable discussions reinforced this but also suggest that the skills challenge is more immediate and uneven than often described. Participants highlighted that while some staff are already confident and proactive in using AI tools, others lack even basic familiarity, leading to variation in confidence, uptake and effective use. This uneven distribution of capability mirrors broader patterns of digital maturity across the system.

In practice, this is creating a growing need for practical, role relevant upskilling, rather than abstract or generalised training. While the literature often frames workforce development in terms of building future capabilities, roundtable participants emphasised the importance of immediate, applied learning, focused on how AI can be safely and effectively used within specific roles and workflows. This extends beyond technical proficiency to include judgement-based capabilities, such as understanding the limitations of AI, validating outputs and knowing when to rely on or challenge AI generated content. As such, the critical workforce capability is not simply knowing how to use AI tools, but knowing when to trust them, when to question them, and how to apply professional judgement to their outputs.

Professional judgement is becoming a critical workforce skill

A more specific shift emerging from both the literature and the roundtables is the changing nature of skills required to work effectively with AI, moving away from technical expertise towards judgement and evaluation. The literature highlights that effective use of AI depends on practical capabilities, such as knowing when and how to use tools, interpreting outputs, and exercising professional judgement, particularly where outputs are ambiguous or incorrect. The roundtables added depth to this, illustrating how these shifts are already playing out in practice.

Participants described AI as relatively easy to begin using, but difficult to use well, with the ability to critically review and validate outputs emerging as a distinct skill. As a AVT roundtable participant noted:

“ The skill is changing from producing the document to evaluating it. The AI may write the first version, but the clinician has to become the reviewer, editor and quality assurer, spotting what is missing, what is wrong and what still needs professional judgement. ”

This was particularly evident in the use of AVT, where reviewing AI generated outputs to support clinical documentation during consultations was repeatedly described as requiring careful judgement. Participants highlighted the risk of automation bias, particularly among more junior or less confident staff, who may be more likely to accept outputs without sufficient challenge. At the same time, it was noted that junior and early career

clinicians often derive greater immediate benefit from these tools, while more experienced clinicians, with established ways of working, may find them more limiting. These dynamics point to variation not only in capability, but in how different parts of the workforce engage with and are affected by AI.

The discussions also surfaced broader implications for workforce development, including concerns that the automation of routine and entry level tasks may reduce opportunities for experiential learning and skill development. This raises questions about how traditional training pathways and progression routes may need to evolve in response to AI adoption, an issue that is not yet fully reflected in existing workforce frameworks or education curricula.

Formal workforce structures are lagging behind AI-enabled work

Alongside these shifts in skill requirements, both the literature and the roundtables highlight that training, support and role design are not keeping pace with the rate of change. The desktop review identifies a pattern where training is often front-loaded, one-off and focused on how to use specific tools, with limited provision for ongoing learning as roles evolve. Similarly, formal job descriptions and role profiles rarely reflect the additional responsibilities associated with AI-enabled work, particularly those relating to oversight, validation and risk management. In practice, the roundtables illustrated how these gaps are being addressed through informal mechanisms. Participants described a reliance on peer champions, prompt sharing and the development of local resource libraries to support adoption and build confidence. While these approaches can make AI use more accessible and less intimidating, they also contribute to variation in practice and are not consistently supported or recognised at an organisational level. The absence of clear expectations around AI use was also described as a source of uncertainty, with some staff unsure whether they would be judged for using, or not using, these tools.

For the workforce, this gap between formal structures and lived reality introduces risk. Where role expectations are unclear and learning is reliant on informal and inconsistent mechanisms, adoption can become fragmented and uneven. It also reinforces the disconnect identified in earlier themes between how work is changing in practice and how it is formally defined and supported. These findings point to the need for ongoing, embedded support and more explicit role design, ensuring that learning, expectations and responsibilities evolve in line with AI enabled ways of working. As a corporate roundtable participant put it:

“ Once the formal training happens, then people go back to their jobs and confidence drops away. What seems to help is a steady drip of real examples, prompts, champions and peer support, so people can see how AI fits into the work they are actually doing. ”

Leadership, culture and trust influence workforce readiness

Alongside skills and support structures, organisational culture and staff attitudes towards AI play a significant role in shaping adoption. The literature highlights issues of trust, perceived risk and resistance to change, and these are echoed in the roundtables. However, stakeholder discussions point to a more nuanced picture, where attitudes are not uniformly resistant but instead highly variable and often context dependent. In some areas, there is strong enthusiasm and openness to experimentation, particularly where AI is seen to alleviate workload pressures. In others, concerns around accuracy, safety, data governance and professional accountability contribute to hesitation or reluctance to engage.

Importantly, the roundtables suggested that culture is being shaped as much by local leadership and peer behaviour as by formal organisational messaging. Teams where leaders actively support experimentation and provide clarity on acceptable use tend to demonstrate higher levels of engagement and confidence. Conversely, where guidance is unclear or perceived as restrictive, staff may either disengage or continue to use AI informally without visibility or support. This shows a gap between formal policy and real-world practice.

Differences are also evident between corporate and AVT contexts. In corporate settings, the accessibility of general-purpose AI tools means that individual capability and initiative play a larger role, with staff often self-teaching and learning through experimentation. This can accelerate adoption but also increases the risk of inconsistent or inappropriate use where skills and understanding are limited. In contrast, the implementation of AVT is typically more structured, with greater emphasis on training, onboarding and supported use within defined workflows. This creates a more consistent baseline of capability but may also limit opportunities for informal learning and adaptation.

Across both contexts, there is a shared recognition that confidence and trust are as important as technical skill. Staff need to feel able to use AI tools safely, understand their limitations and be supported in exercising professional judgement. The literature identifies trust as a key factor in adoption, and the roundtables reinforce this, while also highlighting that trust is built through experience, peer learning and visible organisational support, rather than through policy alone.

Taken together, these findings suggest that AI is not only introducing new tools, but is actively reshaping the skills profile, expectations and cultural dynamics of the NHS workforce. However, these changes are emerging unevenly and largely outside of formal structures, with capability, confidence and approaches to use varying significantly across roles and organisations. This creates a risk that differences in access, experience and informal support may translate into widening variation in how AI is adopted and used in practice. Without clearer alignment between how work is evolving and how the workforce is formally structured, supported and developed, these changes may remain fragmented, limiting both the effectiveness of AI and its potential to deliver sustained improvements in workforce experience.

4. Workforce engagement is essential to ensure benefits are maximised and harms are avoided

The adoption of non-clinical AI and AVT is exposing gaps in governance, trust and how benefits are realised and experienced across the NHS workforce. There is consistent evidence that governance arrangements, accountability structures and approaches to evaluation have not kept pace with the speed and nature of AI adoption in practice. These gaps are not only systemic, but are directly experienced by staff, shaping how confidently AI is used, how risk is managed and how benefits are perceived. As a result, governance, trust and impacts are emerging not simply as organisational considerations, but as critical factors influencing workforce behaviour, experience and outcomes. This is significant, as governance functions not only as a mechanism for managing risk, but also as a means of providing staff with the confidence and clarity needed to use AI safely and consistently in practice.

Governance has not kept pace with adoption

The literature emphasises the need for clear governance structures, particularly in relation to data protection, accountability and risk management. It also points to the challenges of demonstrating and evaluating impact, especially in the absence of standardised metrics or consistent evaluation approaches. These findings are reinforced by the roundtables, but with a stronger emphasis on the practical tensions that organisations are currently navigating between to enable innovation and maintain appropriate levels of oversight.

In practice, participants described a landscape where governance arrangements are still evolving and, in some cases, falling behind real-world use. This is particularly evident in the context of general-purpose AI tools, where staff may already be using to support day to day work without clear organisational guidance or formal approval. This creates uncertainty around acceptable use, data handling and accountability, and can lead to inconsistent approaches across teams. While the literature acknowledges these risks, the roundtables suggested that the scale and immediacy of this governance gap may be greater than formally recognised.

Trust, clarity and permission influence workforce adoption

Beyond formal structures, evidence highlights the critical role of trust, clarity and perceived permission in shaping how AI is used in practice. The desktop review identifies ambiguity in governance and accountability arrangements, with responsibility often pushed down to individual staff, even where AI systems are centrally procured or mandated. This can place a disproportionate burden on individuals, who may be expected to justify decisions influenced by systems over which they have limited control or visibility. The corporate roundtable discussions reinforced this, emphasising the importance of clear boundaries and guidance. Participants noted:

“ Give people the power to use the tools, but make sure the guardrails are clear and properly governed. The issue is not just whether people can use AI, it is whether they understand what they can do, what they should not do and where the boundaries are. ”

The roundtables added further depth, with participants frequently reporting uncertainty about whether AI use was permitted, even where tools were available. Inconsistent governance and limited access were described as factors driving some staff towards personal or external tools, increasing potential risks around data security and oversight. Participants observed that:

“ Staff will use these tools if they help them do the job. So if organisations do not provide a safe, approved route, people will find something else to use, and that creates a different kind of risk. ”

These findings highlight that when governance is unclear or fragmented, staff behaviour adapts in ways that may inadvertently undermine both safety and organisational control.

Trust emerges as a central factor in this dynamic. Mistrust of AI was reported as a barrier to adoption, linked to a perceived lack of robust evidence, scepticism towards new technologies, concerns around data access and storage, and anxiety about potential impacts on roles and workforce size. Participants also noted that trust is shaped by prior experience, with unsuccessful or poorly implemented technology rollouts contributing to caution or resistance. Where governance questions remain unresolved or inconsistently applied, this mistrust can be reinforced. These concerns extend beyond technical performance to include anxiety around job security, role change and professional identity. Where the implications of AI for workforce roles are unclear or not openly discussed, this can reinforce hesitation and reduce willingness to engage.

At the same time, the discussions suggest that clearer and more centralised governance, aligned to real workflows, has the potential to support trust, reduce uncertainty and enable more confident adoption. Participants highlighted the importance of governance approaches that reflect how AI is used in practice, rather than relying on rigid distinctions between clinical and non-clinical contexts. This reinforces the need for governance to be both enabling and practical, providing clarity on acceptable use while supporting innovation.

Another tension emerging from both evidence sources is the balance between enabling innovation and experimentation, while managing risk. Overly restrictive governance can limit experimentation and slow adoption, while insufficient oversight can expose organisations to safety, legal and reputational risks. Roundtable participants highlighted the challenge of striking this balance in a rapidly evolving technological landscape, particularly where national guidance is still developing and local interpretation varies. This contributes to further variation in how organisations approach AI adoption.

Workforce benefits are not consistently captured

The question of benefits realisation is closely linked to these governance challenges. While there is widespread expectation that AI will deliver productivity gains and efficiency improvements, evidence of measurable, system level impact remains limited. The literature notes that benefits are often localised or anecdotal, and that robust evaluation is lacking. Roundtable discussions supported this but also highlighted that benefits are often experienced in practice before they are formally measured or captured. Staff report time savings, improved workflow efficiency and reduced administrative burden, but these gains are not consistently quantified or translated into organisational metrics.

This lack of consistent evaluation also has direct implications for the workforce. Where benefits are not formally captured or recognised, particularly those relating to workload sustainability, morale and reduced cognitive burden, it becomes more difficult to validate staff experience, build organisational confidence or support wider adoption.

Evaluation frameworks do not reflect workforce value

A further issue emerging is that existing approaches to evaluating impact may not fully capture the types of benefits that AI is delivering in practice. The desktop review highlights weak and inconsistent evaluation of AI initiatives within non-clinical, corporate and operational functions, with benefits often uneven and difficult to evidence. Roundtable participants also emphasised the challenge of measuring return on investment where benefits are primarily qualitative, such as improved morale, reduced cognitive load, staff retention or the ability to sustain work and life balance. As one participant observed:

“ Some of the most important benefits are hard to put into a business case. If the value is that staff finish on time, feel less cognitively overloaded or are more likely to stay in their role, that is still a real benefit, even if it does not show up as a simple cash releasing saving. ”

This points to a misalignment between current measures of success and the realities of workforce experience, with existing frameworks often prioritising financial or quantitative metrics over those that reflect staff wellbeing, sustainability and the ability to maintain performance over time.

Discussions around AVT highlighted that benefits experienced by staff, such as finishing on time or reducing after hours administrative burden, were highly meaningful but rarely captured within formal evaluation frameworks. This points to a misalignment between current measures of success and the realities of workforce experience, with existing frameworks often prioritising financial or quantitative metrics over those that reflect staff wellbeing and sustainability.

This creates a disconnect between perceived and evidenced benefits. Without more appropriate and consistent approaches to evaluation, it becomes difficult to build a compelling case for scaling adoption, prioritising investment or understanding where AI

delivers the greatest impact. It also limits the ability to compare approaches across organisations or to identify what constitutes success in different contexts. Participants emphasised the importance of sharing learning from pilots and early implementations across organisations to avoid duplication of effort and ensure that emerging evidence is used effectively.

Differences are also evident between corporate AI and AVT contexts. In corporate settings, the use of general-purpose tools is often less visible and less formally governed, making it more challenging to monitor usage, manage risk or assess impact. Benefits may be realised at an individual or team level but remain largely undocumented. In contrast, the implementation of AVT is typically associated with more formal governance, clearer approval processes and greater scrutiny, particularly given its proximity to clinical workflows. This creates more structured opportunities for evaluation, but also introduces additional complexity in terms of assurance, safety and accountability.

Despite these differences, both contexts highlight a common issue: the absence of a consistent, system-wide approach to governance and benefits realisation. Fragmentation in how AI is implemented and assessed limits the ability to scale successful approaches, share learning or establish common standards. It also reinforces variation in risk tolerance and decision making across organisations.

Overall, the evidence suggests that governance and benefits realisation are critical enablers that have not yet kept pace with adoption. While the literature frames this as an emerging challenge, the roundtables indicated that organisations are already actively managing these tensions in practice. Without clearer governance, stronger alignment with real-world use, more appropriate measures of value, and greater attention to trust and workforce experience, there is a risk that AI adoption remains fragmented, with benefits that are difficult to evidence, scale or sustain across the system.

5. Fragmented and ad-hoc AI adoption limits workforce benefits

AI adoption across the NHS remains fragmented and largely ad hoc, limiting the benefits for the workforce. This reflects a pattern of local experimentation, early adoption and incremental progress that has not yet translated into coordinated, system-wide transformation. While there is clear evidence of innovation and growing use of non-clinical AI and AVT, this activity remains dispersed and insufficiently connected to broader strategic change. For the workforce, this results in inconsistent access to tools, variation in how work is organised and an uneven experience of AI-enabled change across organisations and roles.

The literature highlights the challenges of scaling innovation within the NHS, pointing to structural, organisational and cultural barriers that can limit the spread of successful approaches. While often framed as a system-level issue, the roundtables indicated that

these challenges are experienced directly by staff. Participants described variation in how AI is introduced, supported and embedded, leading to differences in workload, expectations and opportunities to benefit from AI. As a result, workforce experience is shaped more by local context than by a consistent organisational strategy.

Fragmented local adoption creates uneven workforce experience

At present, much of the activity is locally driven and context specific, reflecting the operational pressures of individual organisations and teams. While this has enabled rapid experimentation, it has also resulted in duplication of effort and limited opportunities for shared learning. For staff, this can create inefficiencies and frustration, with similar challenges being addressed in parallel and little visibility of existing solutions. It also contributes to variation in capability development and confidence, reinforcing differences in how AI is used. AVT roundtable participants highlighted this clearly, noting that:

“ People are reinventing the wheel in different places, without always knowing who has already tried this, what they learned or how to skip the first ten steps. We need a way of sharing those experiences so organisations can learn from each other rather than repeating the same work. ”

Lack of ownership and coordination limits workforce benefits

A key factor underpinning this pattern is the lack of clear ownership for both AI tools and the processes they are intended to support. The desktop review highlights that ownership is often ambiguous once tools move beyond initial pilots, which can limit momentum and contribute to change fatigue. The roundtables echoed this, with participants describing siloed activity and limited coordination. For the workforce, unclear ownership translates into uncertainty about expectations, duplication of effort and inconsistent support, particularly where AI is perceived as an additional task rather than an integrated part of service delivery.

Practical constraints further reinforce this fragmentation. Participants noted that licence restrictions and uneven access to tools can limit who is able to use AI and how it is applied. This creates differences in workload and efficiency between teams and roles, with some staff benefiting from AI-enabled support while others continue to operate without it. Such variation risks creating inequities in both workload and experience across the workforce.

This lack of coordination also reflects a wider misalignment between local innovation and system-level strategy. While national and regional priorities increasingly recognise the role of AI, there remains a gap in translating these into coherent approaches that support consistent adoption. From a workforce perspective, this results in a lack of standardisation in how roles evolve, how skills are developed and how support is provided, reinforcing fragmented and inconsistent impact. AVT roundtable participants emphasised the need for a more aligned approach, noting that:

“ We need to be led by our own problem statements, rather than by whichever technology is being offered to us. Without a unified route for approval and adoption, different teams can end up buying or using different tools for different use cases, without a coherent solution across the organisation. ”

Differences between corporate AI and AVT contexts further illustrate this variation. In corporate settings, the decentralised and often informal nature of adoption means that workforce experience is highly dependent on individual initiative and local practice. In contrast, AVT implementation is typically more structured, providing clearer pathways for use and support. However, even in these contexts, variation persists across organisations and specialties, highlighting that more coordinated approaches do not automatically translate into consistent workforce experience.

Scaling requires workforce and organisational alignment

The roundtables also highlighted that scaling is not simply a technical process, but one that directly affects how work is organised and experienced. Without alignment across workflows, workforce capability, governance and organisational priorities, the impact of AI on the workforce remains uneven and difficult to sustain. Staff in different settings may face very different expectations, levels of support and opportunities to benefit from AI, even where similar tools are available.

At the same time, there is emerging evidence of increasing strategic attention and intent, particularly at regional and system levels. Participants noted a growing focus on coordination and shared learning, suggesting that the conditions for more consistent workforce experience are beginning to develop. However, this is not yet fully realised, and variation remains a defining feature of current adoption.

Overall, the evidence suggests that the NHS is in an intermediate phase, where fragmented and ad hoc adoption is limiting the extent to which AI can deliver consistent and equitable benefits for the workforce. While the literature frames scaling as a system challenge, the roundtables highlighted that its consequences are experienced directly by staff. Without clearer ownership, stronger coordination and alignment between local activity and system priorities, there is a risk that AI reinforces differences in workload, capability and experience, limiting its potential to support a more sustainable and effective workforce.

Next steps and recommendations

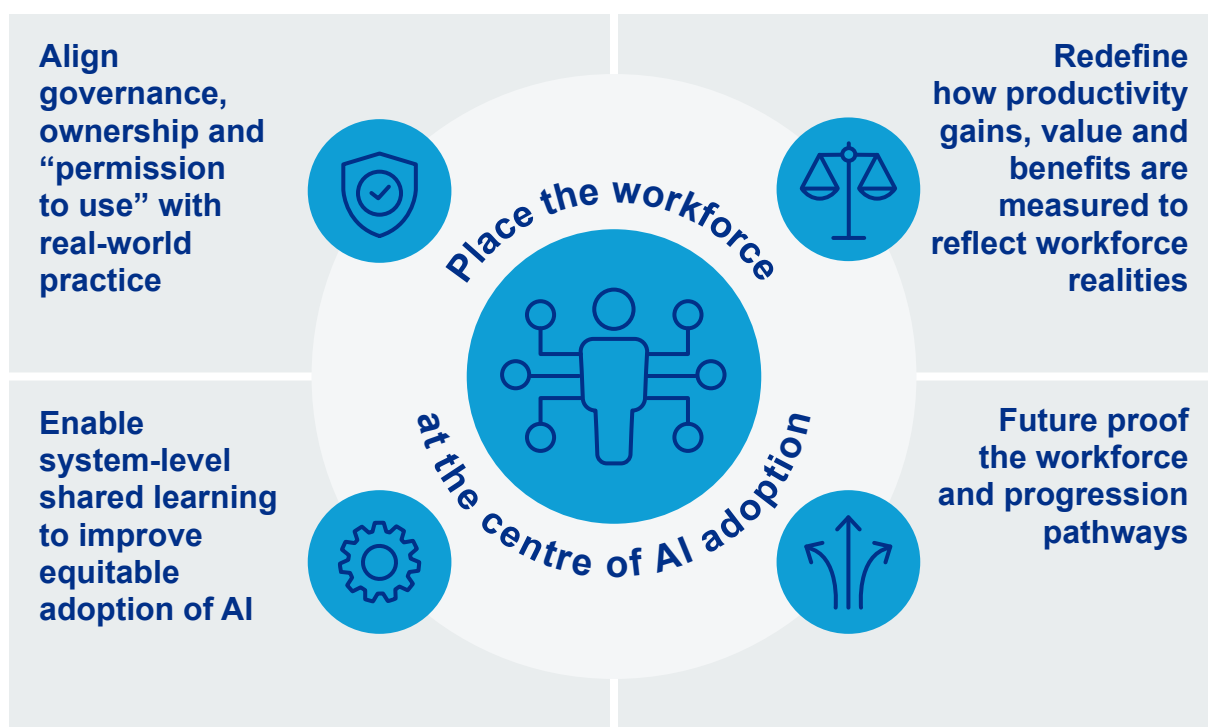


Next steps and recommendations

The findings of this report highlight that the NHS is at a critical stage in the adoption of non-clinical AI and AVT.

The findings of this report highlight that the NHS is at a critical stage in the adoption of non-clinical AI and AVT. While there is clear evidence of progress and growing use, adoption remains uneven, fragmented and insufficiently supported. Workforce structures, capability development and planning approaches have not kept pace with these changes. To move from reactive, fragmented adoption to sustainable and equitable impact, a more deliberate and workforce-centred approach is required. The recommendations that follow are intended to inform future policy development, strategic decision-making, workforce planning and organisational implementation across NHS England, integrated care systems and provider organisations.

Collectively, they aim to support a more coordinated approach to AI adoption that places workforce considerations at the centre of transformation and scale.





Place the workforce at the centre of AI adoption

Embed AI into workforce planning and role design

AI is already reshaping how work is carried out across the NHS, particularly through the redistribution of tasks and the exposure of certain roles, especially entry level and administrative functions, to automation. Despite this, workforce planning processes have not yet systematically incorporated the impact of AI.

System-level workforce planning leaders should incorporate AI-driven role redesign into the next iteration of workforce planning frameworks, with explicit attention to roles most exposed to automation, with a view to developing updated guidance that trusts can apply locally. This includes redesigning roles to reflect the increasing emphasis on oversight, validation and higher value activities, as well as reassessing the function and sustainability of entry level roles. While some organisations are currently managing change through natural turnover, this approach risks overlooking longer-term impacts on workforce supply, training pipelines and service resilience.

Embedding AI into workforce planning will require a more proactive and coordinated approach, ensuring that workforce models evolve alongside changes in how work is delivered.

Shift from one-off training to continuous capability building

The skills required to work effectively with AI are evolving, with increasing emphasis on judgement, interpretation and the safe use of AI-generated outputs. However, current approaches to training are often front-loaded, one-off and focused on specific tools, rather than supporting ongoing learning as roles and technologies evolve.

Those responsible for workforce development and education commissioning should review existing AI training offers against the competencies identified in this report, with the aim of establishing a minimum capability framework for non-clinical AI use that can be adapted at a local level. This includes focusing on practical, role-specific competencies such as critical evaluation, decision-making and understanding the limitations of AI systems. It also requires recognising and supporting informal learning mechanisms that are already emerging, including peer networks, shared resources and on-the-job experimentation.

A more sustained and structured approach to capability building will be essential to ensure that staff are able to use AI safely, confidently and consistently over time.

Clarify roles, responsibilities and accountability

As AI becomes more embedded in workflows, responsibilities for tasks such as oversight, validation and decision-making are becoming less clearly defined. In many cases, these responsibilities are being absorbed into existing roles without formal recognition or clear guidance.

This creates uncertainty for staff, particularly where they are expected to justify decisions influenced by AI systems. It also increases the risk of inconsistent practice and unclear accountability. There is therefore a need to define more clearly how AI should be used within different roles, and who is responsible for the outputs it generates or supports.

Updating job descriptions, role profiles and governance frameworks to reflect these changes will be critical. Greater clarity will not only support safe and consistent use but also help align expectations across organisations and teams.



Align governance, ownership and “permission to use” with real-world practice

The adoption of AI is shaped not only by technical capability, but by whether staff feel confident, supported and permitted to use it. The findings of this report highlight that mistrust, uncertainty and anxiety remain significant barriers, particularly where governance is unclear or where previous experiences with technology have been negative.

Building trust requires more than the provision of tools. It involves clear communication about acceptable use, transparency around how roles may change, and active engagement with staff concerns. In particular, conversations around workforce impact, including potential changes to roles or headcount, need to be handled sensitively and openly.

Organisations should also foster a culture that supports safe experimentation and learning, enabling staff to explore the use of AI without fear of judgement. Strengthening confidence and engagement will be essential to ensuring that adoption is both effective and sustainable.



Redefine how productivity gains, value and benefits are measured to reflect workforce realities

AI is often associated with expectations of reduced workload and increased productivity. However, the evidence suggests that it is more accurately redistributing effort, with increased emphasis on checking, validation and exception handling.

Where this shift is not recognised in workload planning and performance expectations, there is a risk that AI may intensify work rather than alleviate it. Staff may be required to take on additional oversight responsibilities without corresponding adjustments to time allocation or performance measures.

There is therefore a need to reassess assumptions about productivity gains and ensure that workload models reflect the reality of AI-enabled work. This includes recognising new types of activity, such as reviewing and validating outputs, and monitoring the impact on cognitive load and staff wellbeing. There is also a need to change how we define and measure the benefits of AI so it better reflects the reality of staff experience. This means looking beyond traditional efficiency measures and also considering factors such as manageable workloads, staff wellbeing, and other less tangible improvements, not just time or cost savings.

Aligning expectations with practice will be critical to maintaining workforce sustainability.



Future proof the workforce and progression pathways

The increasing automation of routine and entry level tasks is accelerating a fundamental shift in how the workforce develops and evolves. While these technologies may deliver short-term efficiencies and create capacity for higher-value work, it also risks reducing opportunities for skill development and progression, particularly for early-career staff.

This raises questions about how traditional training pathways and entry points into the workforce will need to evolve. Without deliberate action, there is a risk of weakening the pipeline of future staff and undermining long-term system resilience.

To enable workforce transformation, the healthcare system and organisations need to focus on maintaining meaningful entry level roles, adapting training and education pathways, and ensuring that AI adoption supports, rather than disrupts, career development. Monitoring the impact of AI on retention, morale and progression will also be important in shaping future workforce strategies.



Enable system-level shared learning to improve equitable adoption of AI

The current landscape of AI adoption is characterised by local innovation, but there are limited mechanisms for sharing learning on workforce impact across organisations. As a result, similar challenges are often being addressed in parallel, and opportunities to build a stronger evidence base are missed.

There is a need to improve the sharing of insights from pilots and early implementations, particularly in relation to workforce outcomes such as role change, skills development and staff experience. Developing more consistent approaches to evaluating workforce impact will also support better comparison and learning across organisations.

Strengthening system-level learning will be essential to avoid duplication, accelerate adoption and ensure that workforce implications are understood and addressed in a more coordinated way.

Conclusion

This report set out to examine the implications of non-clinical AI and AVT for the NHS workforce.

Drawing on both the literature and stakeholder discussions, a clear picture emerges: AI is already changing how work is carried out across the NHS, but the organisational structures, workforce models and supporting frameworks needed to manage this change have not yet kept pace.

Across all five themes, the findings suggest that AI is being adopted primarily in response to operational pressures and workforce challenges. While it is delivering benefits in specific contexts, particularly through reducing administrative burden, its impact is more complex than commonly assumed. Rather than removing work, AI is redistributing it, shifting effort towards judgement, validation and oversight. This has important implications for how roles are designed, how performance is measured and how staff are supported.

Crucially, the evidence suggests that AI is no longer a future workforce issue. Its effects on work, skills, career pathways and workforce dynamics are already being experienced across the NHS. However, workforce planning, capability development, governance arrangements and approaches to evaluation have not yet adapted in a systematic way. As a result, staff are often navigating changing expectations without consistent support, clear accountability or a shared understanding of how benefits should be realised and measured.

The findings also highlight the importance of trust, clarity and ownership in shaping adoption. Where governance is ambiguous and expectations are unclear, adoption becomes fragmented and inconsistent. At the same time, the report highlights the need to broaden how value is defined and measured, recognising workforce outcomes such as morale, workload sustainability, retention and staff experience alongside more traditional productivity metrics.

Taken together, the findings suggest that the NHS is moving beyond early experimentation with AI but has not yet reached a point of coordinated transformation. The opportunity now is not simply to scale technology adoption, but to ensure that workforce planning, organisational design, capability building and service transformation evolve alongside it.

Without deliberate action, there is a risk that AI adoption reinforces existing pressures, creates new inequalities and delivers only partial benefits. A more coordinated, workforce-centred approach offers the opportunity to improve workforce experience, strengthen workforce sustainability and support the ambitions of an increasingly AI-enabled NHS. The next phase of adoption will therefore depend not only on the technology itself, but on how effectively the NHS places workforce considerations at the centre of AI-enabled transformation.



ea@imperialcollegehealthpartners.com

imperialcollegehealthpartners.com



contact@uclpartners.com

uclpartners.com



hin.southlondon@nhs.net

healthinnovationnetwork.com



England

england.nhs.uk

June 2026