Impact of AI in diagnostics and pathology

A Joint Workshop Report from Health Innovation Oxford & Thames Valley and the Centre for Personalised Medicine, University of Oxford.

February 2026

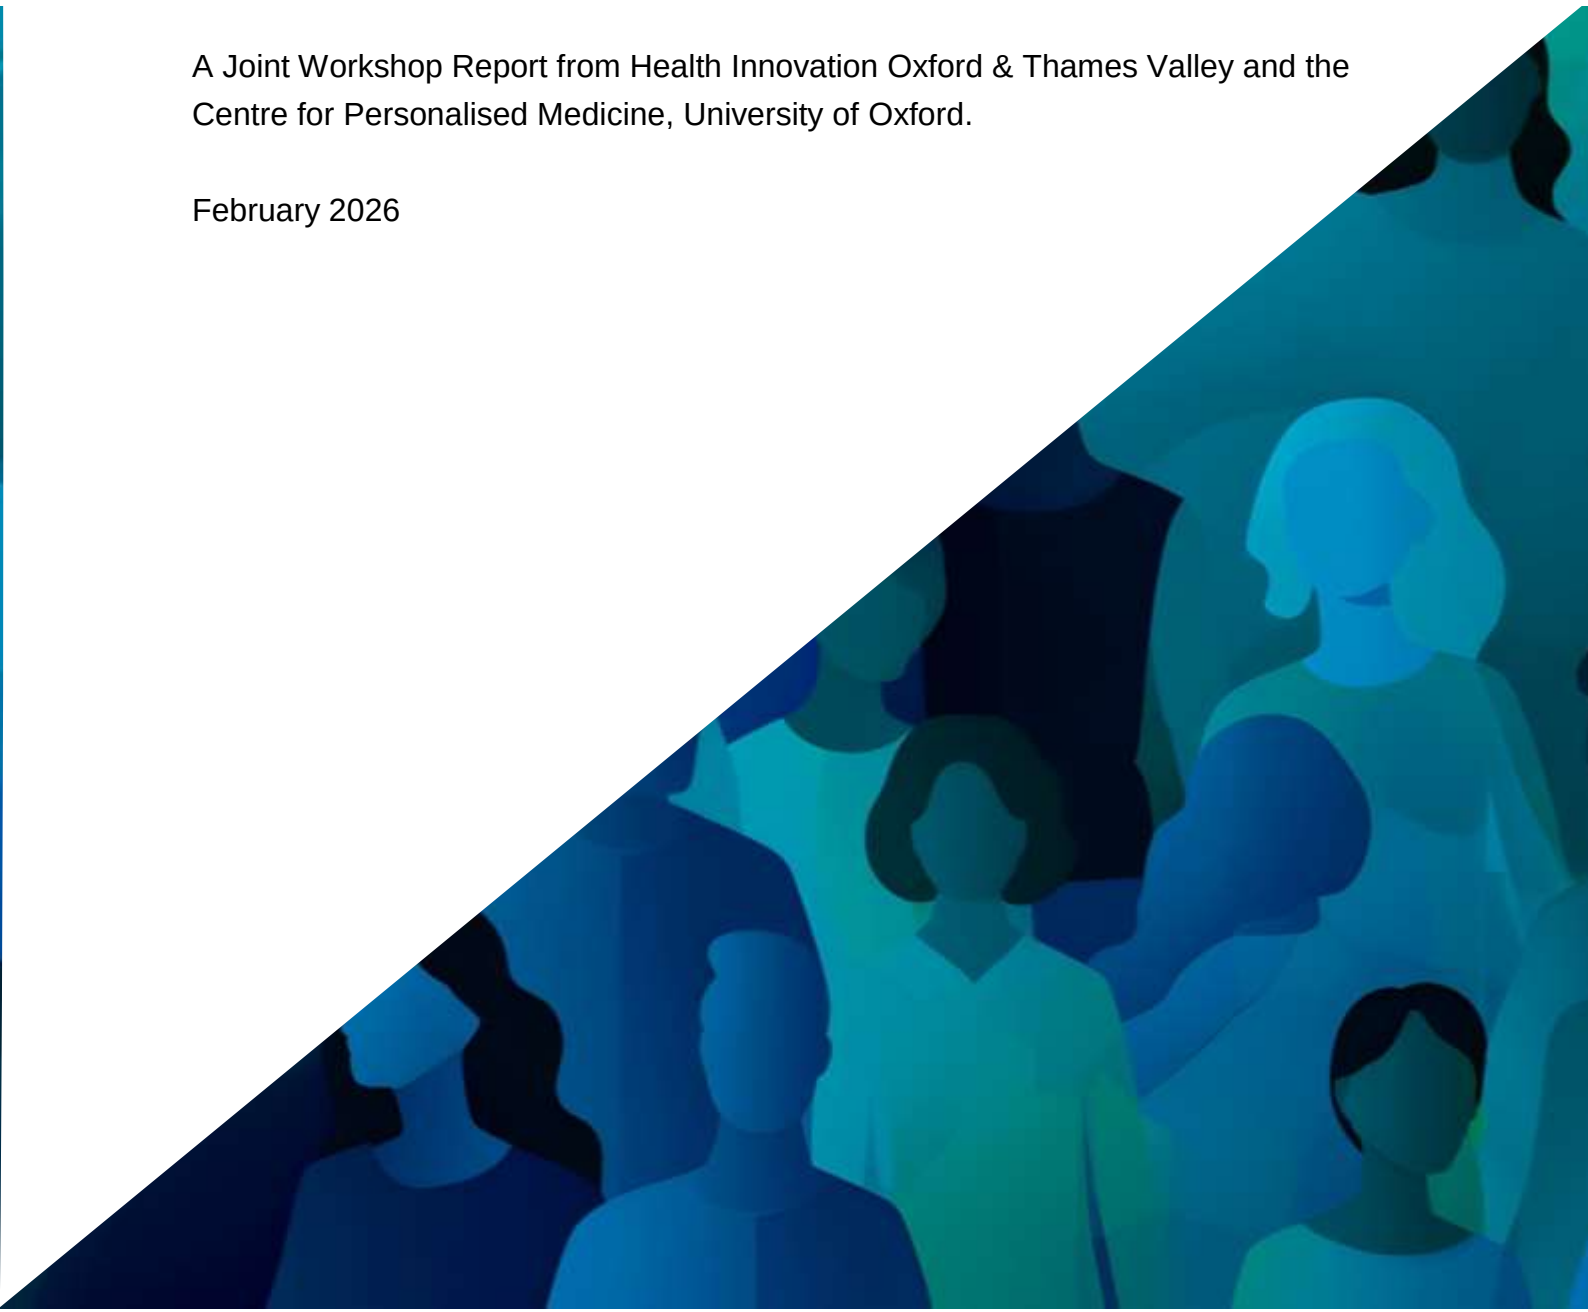

Executive Summary

This report summarises a roundtable discussion convened to address the question: *"How can we harness the diagnostic power of artificial intelligence (AI) in pathology while addressing the challenges of workforce disruption and the evolving role of human expertise?"*. Participants broadly agreed that AI offers potential administrative efficiency gains, improved diagnostic accuracy and new capabilities such as large-scale risk prediction. However, they also had concerns regarding unrealistic expectations, inadequate existing digital infrastructure, regulatory uncertainty, variable trust among clinicians and patients, training, workload duplication, and impacts on workforce morale. Furthermore, it was strongly voiced by participants that AI should not be seen, or relied upon, as the solution to the current workforce crisis. Overall, participants concluded that AI should complement rather than replace human expertise, and that its integration into healthcare must be evidence-based, ethically grounded and supported by robust governance and appropriate resourcing.

Introduction

AI is entering healthcare at pace, often without focused or interdisciplinary attention to the new issues it introduces. To address the potential applications of AI in diagnostics and pathology, Health Innovation Oxford & Thames Valley (HIOTV) in collaboration with the Centre for Personalised Medicine (CPM), Oxford University. convened a roundtable at St Anne's College, Oxford, on Thursday 20 November 2025. The discussion brought together clinicians, ethicists, researchers, technologists and members of the public under the Chatham House Rule, and was chaired by the CPM director, Professor Anneke Lucassen. The breadth of experience among attendees created a unique opportunity to explore a wide range of perspectives. This report does not attempt a comprehensive review of the field and related literature, but instead synthesises key themes that were identified through this round table discussion. We have included some quotes made during the discussion, but these are not attributed to any particular individuals.

What is AI good for, and what is needed?

Participants repeatedly returned to the question of where AI can usefully contribute, and how this aligns with real-life needs in healthcare systems, specifically the NHS. Two broad domains of application were discussed: support for clinical decision-making, and support for administrative and organisational tasks. While there was broad agreement that AI *can* assist with high-volume, repetitive tasks (for example, counting tasks in pathology or processing large numbers of routine blood tests), there was greater uncertainty about whether it can reliably deliver time savings in complex clinical decision-making; where it should fit in current workflows; and how these tools might be implemented in a coherent and minimally disruptive manner.

Supporting clinical decision-making

"There is a disconnect between assumptions of use in clinical care and headlines"

Expectations for efficiency, productivity and time saving from AI-supported clinical decision-making were perceived as high, but not yet consistently demonstrated in routine practice. Whether AI saves time depends heavily on how it is integrated into workflows and how clinicians perceive and interact with its outputs, including risks of "alert fatigue" whereby many false positives are generated, or duplication of work due to a need to 'double check' AI decisions, and the possibility of automation bias when clinicians over-rely on AI recommendations despite visible errors. Examples discussed included AI systems for image interpretation in radiology or pathology, where anticipated efficiencies (such as clinicians not reviewing AI-negative cases) may be eroded if clinicians feel obliged to re-check outputs due to prior experience of missed findings.

"There's the assumption that AI is a time saver but if you can already tell something e.g. see a fracture clinically and/or on X-ray, why do you have to go through AI if you can clearly already see it?"

AI supports a pathologist to help detect cancer cells.
The AI reports that 30 of the 100 samples show no evidence of cancer.
In theory these samples could "be put on the shelf", i.e. not reviewed. ***Time saved.***
However, the pathologist has previously found cancer within the apparently negative 30 samples.
Consequently, they feel they must review all samples themselves, especially as they are named in the results documentation. ***No, or little, time saved.***

Participants also highlighted areas where AI might enable genuinely new diagnostic capabilities that are difficult to achieve otherwise, such as reviewing very large and complex imaging or laboratory datasets (for example, in lung cancer screening or large-scale analysis of blood tests, genomics and proteomics). In addition, AI can support population-level risk prediction and standardisation of currently variable tests, though such developments would likely generate new forms and volumes of work by necessitating new preventive pathways and follow-up.

Supporting administrative and organisational tasks

There was strong agreement that some of the greatest immediate needs, and perhaps the most feasible gains, lie in administrative and organisational domains such as diary management, clinical documentation, workflow coordination and workforce management. Participants felt that AI is well suited to repetitive, data-heavy tasks of this kind, but that such applications tend to attract less attention and investment than high-profile clinical decision tools, despite their importance for day-to-day system functioning, system resilience and staff

experience. Concerns were raised that “back-office” functions may be under-prioritised if investment continues to focus on more high-profile clinical applications (with potential to generate headlines and political interest).

The human and the machine

“If full AI adoption was coming tomorrow, how do I protect the workforce, how do I protect the human interaction element of healthcare?”

Discussions about where AI should be deployed were inseparable from questions about workforce, skills and the nature of professional work. Participants were concerned about narratives proposing that AI could solve the NHS workforce crisis or render some roles redundant, stressing instead the need to frame AI as assisting, not replacing, clinicians. They argued that if AI is presented as the primary solution to workforce pressures, this may reduce attention and responsiveness to broader policy and system reforms that are also needed. In addition, it was feared that this narrative could also prevent people from entering the healthcare workforce altogether. Participants were concerned that AI’s efficiency ‘promise’ might be linked to erosion of clinical expertise and skills, particularly in the context of a national health system that is already stretched and under-resourced.

Maintaining skills and developing new competencies

A recurrent theme was the risk of deskilling if clinicians develop an over-reliance on AI, particularly in areas where repeated exposure to normal and abnormal findings underpins clinical acumen. Participants emphasised the importance of understanding how to maintain and develop clinical skills when key elements of pattern recognition or classification are delegated to AI tools. The potential erosion of “relational knowledge” – informal sharing of expertise across colleagues and specialties – was also noted as a concern if workflows become more fragmented or intermediated by AI technology. Equally, there was acknowledgement that introducing new technologies into any clinical or administrative setting will necessitate developing new competencies, new skillsets, and training.

Keeping the job worth doing

With low morale and burnout already affecting many NHS staff, participants stressed that how AI changes the character of clinical work will influence workforce satisfaction and retention. There was concern that poorly designed AI implementations, particularly those that add alert burden or administrative overhead, could diminish job satisfaction, even if they bring some technical benefits. Participants also highlighted the ongoing importance of relational aspects of care and interprofessional work. Equally, the limits of purely empirical or

data-driven tools in addressing patients' need for explanation, reassurance and shared decision-making.

"AI uses purely empirical tools whereas within medicine and healthcare we want to find some sort of ground truths which AI cannot provide"

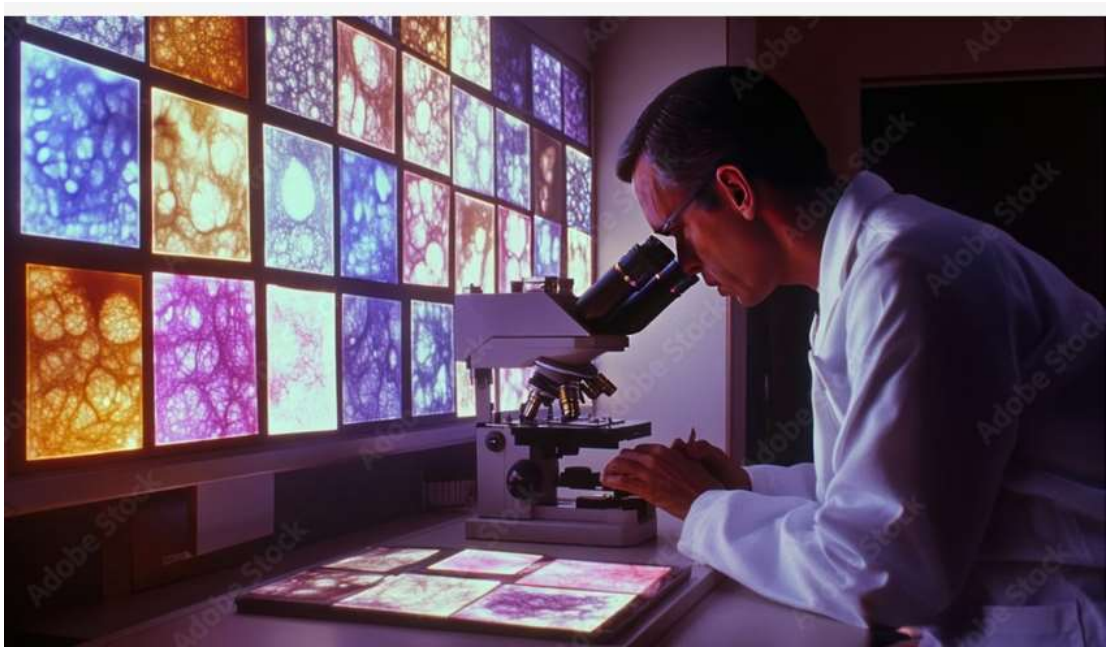

Transparency

Transparency emerged as critical for both professional and public trust in AI-enabled healthcare, but equally it was recognised that full transparency about exactly what AI does (and does not do) may be difficult to achieve. However, shedding some light on the black box of AI was seen as important - clear information about what an AI system is doing (and not doing), why it is being used in a specific clinical context, and importantly, in which populations and scenarios it performs well (as well as in which it cannot be relied upon). Additionally, issues pertaining to data privacy and data ownership were raised throughout by participants, including how data may be used in current and future training models and the security of the platforms used.

Participants identified what they termed the "explainability paradox" – clinicians do not necessarily need to understand the detailed technical mechanisms of an algorithm in order to use it safely, yet they must be able to contextualise its outputs and explain them meaningfully to patients. This may create additional cognitive and time demands, as clinicians must integrate AI outputs into their reasoning while remaining accountable for decisions. Furthermore, AI explanations are fundamentally different from human reasoning. When a clinician explains a diagnosis to a patient, they typically draw on causal narratives ("I saw these features, which together suggested X condition, because..."). AI systems rarely

work this way; they identify statistical patterns in data that may not correspond to the causal mechanisms clinicians use to reason about disease. This mismatch between how AI "thinks" and how clinicians explain medical reasoning to patients can make it difficult for clinicians to bridge the explanatory gap – to translate algorithmic outputs into the kind of comprehensible, narrative explanation patients expect and need.

Public trust, data security and the patient journey

For patients, transparency about when and how AI is used, and about whether their data contributes to ongoing AI training or model updates, was seen as essential to maintain trust. Participants noted that advances in technology over the last 10-20 years mean that patient journeys are changing. People are used to searching the internet for their symptoms, using online AI tools to refine their questions, or using direct-to-consumer tests, often before they engage with a health professional. Rather than seeking help with a diagnosis, healthcare engagement often relates to perceived and estimated levels of disease and lifestyle 'risk'. What might a patient develop in the future and what can I do to prevent or ameliorate that disease from happening? This raises questions about how clinicians and patients may perceive and interpret AI-generated risk estimates as compared to a clinician's judgment, especially when they diverge. Concerns about cyber security and the implications of potential attacks on AI-enabled systems (both appropriation of data or blocking/interfering with systems) were also highlighted as fundamental issues that must be addressed clearly and openly.

Health and environmental equity

The equity implications of AI were a major focus. Because AI performance depends on training data, there is a real risk that systems will perform poorly in under-represented groups, thereby exacerbating existing health inequalities, especially where there is little commercial incentive to prioritise those at greatest risk of being left behind. At the same time, some participants noted that AI should not have intrinsic prejudices and could, in principle, help democratise decision-making and support more equitable access if systems are designed and governed appropriately. However, this relies on larger datasets which are representative of entire populations and supportive of appropriate stratification.

Infrastructure inequities were also highlighted; for example, digital pathology is currently fully operational in only a minority of laboratories across England, meaning that access to AI-supported pathology is already subject to a “postcode lottery”. The ability of individual NHS trusts to invest in AI-supported pathology is likely to vary, based on their financial situation and existing digital infrastructure, particularly as some AI-supported pathology requires onward subscription fees in addition to up-front costs. As with earlier technologies such as facial recognition, participants viewed overall utility and fairness as a trade-off requiring explicit consideration of both benefits and harms across different populations. Environmental inequities were raised; for example, direct risks posed by AI to the environment and to [legally-binding] NHS Net Zero commitments. Mechanisms include the impacts of AI on energy, water, and critical mineral extraction – primarily through the major infrastructure needed to build and maintain AI technologies, which is not equitably distributed by areas of highest AI usage.

Implementation, infrastructure and evidence

Concerns about implementation, infrastructure and the evidence base underpinned many of the other themes. Participants noted that implementation is rarely linear and emphasised the need to learn from past experiences with new technologies in the NHS, particularly when adding complex systems to already stretched services. Participants also discussed how to balance the need for localised approaches to implementation of these technologies alongside national standardization.

System readiness

There was consensus that much NHS digital infrastructure remains outdated, fragmented and fragile, creating anxiety about adding further complexity through AI systems without first investing in “unseen” but essential digital foundations. In pathology, for example, lack of widespread digital pathology capability is already limiting the feasibility of large-scale AI deployment, regardless of algorithmic capability. Participants stressed the importance of local implementation and evaluation, including practical support for training and infrastructure, and of understanding the full costs beyond pilot projects, noting that high license fees can be difficult to sustain even where productivity gains are demonstrated.

Regulation was another area of concern, with participants noting that the breadth and speed of AI development make it challenging for regulatory frameworks to keep pace and expressing fears that under-regulation could lead to unsafe deployment in clinical settings.

Politics and economics

Policy documents such as the NHS 10 Year Plan (May 2025) have emphasised AI’s potential to transform healthcare, often positioning it as a core means to reduce costs, manage demand and address workforce challenges. This includes a pledge to make AI every nurse and doctor’s “trusted assistant – saving them time and supporting them in decision making”.

Participants discussed tensions between health policy and economic growth agendas, including the influence of private sector AI companies and investors who may seek rapid returns that outstrip the generation of robust evidence, careful implementation and patient experience, over commercial benefits. They cautioned that headline claims about dramatic workforce savings may be misleading if they ignore the proportion of clinicians’ work that involves communication, relational care and other tasks that are not easily automated, and argued that full economic models must include procurement, licensing, quality assurance, dataset updating, infrastructure upgrades and distributional effects such as equity impacts.

"Headline figures such as it will save 40% of your workforce, but actually it isn't 40% as image reading is only 20% of their workload anyway – the rest is talking to patients and hands-on care."

Mind the evidence gap

Participants drew an analogy between AI and other clinical interventions or tests, noting that while there is substantial research on whether AI systems *can* work under controlled conditions, there is much less evidence on whether they *do* work in routine practice and at scale. For example, there have been instances whereby the use of ambient voice technologies in General Practice have created unforeseen issues for end-users, with large volumes of text congesting electronic health records and creating barriers to accessing safe, concise documentation. Economic evaluations and cost-effectiveness analyses for AI tools were described as weak, and real-world implementation limited, resulting in a lack of robust outcome data on clinical effectiveness, safety and equity. The group advocated that AI should be subject to the same evaluative standards as other interventions, including systematic assessment of whether it works, for whom, in which settings, equity and environmental considerations, and whether it is worthwhile given opportunity costs. This includes the need for strong, robust regulation and legislation surrounding the use and application of AI tools – including liability and responsibility for unintended consequences and/or harms. Participants expressed concern regarding the outpacing of AI tool development as compared to regulation and legislation in this area.

Implications

Participants identified several key implications for future policy, practice and research. First, there is a need to "stop and think" about what is genuinely needed in health and care systems and where AI can realistically help, avoiding simplistic narratives that present AI as a "silver bullet" or a substitute for addressing structural workforce and system deficits. Second, the evidence base must be strengthened to include not only technical performance but also impacts on equity, relationships, skills, morale and full economic costs across the system.

Third, developers, clinicians and patients should work together to produce meaningful explainability resources and training materials that support clinicians in using and communicating about AI tools responsibly. Finally, the continuing importance of relational care should be explicitly recognised and embedded in training, implementation and evaluation. This will ensure that the integration of AI in pathology and diagnostics enhances, rather than undermines, the human elements that are central to good care.

Health
Innovation
Oxford & Thames Valley

CENTRE *for* PERSONALISED MEDICINE
St Anne's College, Woodstock Road,
Oxford, OX2 6HS, United Kingdom

Centre for Human Genetics, Roosevelt Drive,
Oxford, OX3 7BN, United Kingdom

HEALTH INNOVATION OXFORD & THAMES VALLEY
Magdalen Centre, Oxford Science Park
OX4 4GA