“Just because we can, doesn’t mean we should.”

A roundtable discussion on AI and diagnostics

November 2025

Contents

Summary

This report highlights key points from a roundtable discussion convened to discuss 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?"

It was widely agreed that AI offers the potential for efficiency gains, improved diagnostic accuracy and new capabilities, such as large-scale risk prediction. However, the discussion highlighted a number of challenges and concerns including unrealistic expectations, insufficient infrastructure, regulatory uncertainty, trust by clinicians and patients, alongside workforce disruption and impact on morale. The consensus was that AI should complement—not replace—human expertise, and its integration must be evidence-based, ethically sound, and supported by robust governance.

1. Introduction

The rapid pace at which AI is entering the healthcare system demands attention. This roundtable brought together experts including clinicians, ethicists, researchers, technologists and members of the public to explore the current landscape of AI in healthcare, particularly considering diagnostics, and pathology.

There is an abundance of literature already available on this topic, consequently this report does not try and cover everything there is to know or consider. Rather, it is the synthesises of the discussion into key themes.

The roundtable was held at St Anne's College in Oxford, on Thursday 20th November 2025. It was co-hosted by the Health Innovation Network Oxford & Thames Valley and the Centre for Personalised Medicine. Professor Anneke Luccason chaired the discussion.

2. What is it good for vs what do we want and need?

This broad question came up early in the discussion, and was circled back to frequently - what is AI good for, and how does this align – or not – with what is truly needed in healthcare? Two main areas for AI deployment in healthcare were discussed: clinical decision making and administration. Participants agreed that AI offers clear potential to create efficiency by carrying out repetitive, high-volume counting tasks but were less confident on its ability to save time in clinical decision making.

2.1 The sexy stuff - clinical decision making

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

Efficiency, timesaving and productivity are often promised for AI supported clinical decision making. However, this has yet to be fully evidenced in practice. Whether or not time is saved by AI will depend on a range of real-world factors, in particular how clinicians regard and interact with the AI:

- Pop-up Fatigue – if AI highlights many false positives, clinicians may simply switch-off attention to the results it flags – this sort of action can be observed in other automated workflows like drug interaction alerts.

- Trusting the tests results – an anecdote was shared:

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 and the results documentation. ***No, or little, time saved.***

- Automation bias – where the opposite behaviour may take place. Errors of automation bias tend to occur when decision-making is dependent on computers or other automated aids, and the human is in an observatory role but able to make decisions. People have been found to rate the automated decisions more positively than neutral, even in the face of observable error.

- Simply deploying AI without proper consideration of the pathway – for example, using AI in emergency care for fracture detection. Routine use for all cases might simply increase workload:

"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?"

Examples of where AI could be used to do things that can't be done easily otherwise were also discussed. These included in lung cancer screening for example where large amounts of complex, repetitive data need to be reviewed, which would be difficult to achieve without computational support. Similarly, making sense of large volumes of data for example the 1.5 million blood test/year which will increase with more routine use of genomics and proteomics tests for example.

AI could also enable new forms of diagnostic practice, such as population-level risk prediction for example precision prevention in breast cancer or the standardisation of tests that are currently too variable for widespread adoption. However, introducing these types of prediction pathways, leading to preventative care, would create a new volume of work with attendants.

2.2 The unsexy stuff - administrative tasks

It was widely agreed that AI may be most needed in areas where administrative tasks create a significant burden for the workforce – such as diary management, documentation and workforce management. AI has clear potential for repetitive, data-heavy tasks, however, participants felt this isn't headline grabbing and so a less attractive area of focus. The fear is that this will result in back-office functions being overlooked and so receive relatively less investment than more high-profile clinical applications.

3. 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?"

Inseparable from the discussions about where AI might be best deployed, were the conversations about the impact on workforce. There is a rhetoric that AI will make sections of the workforce redundant. Participants highlighted the need to change this narrative to AI assisting the workforce. Despite the well acknowledged NHS workforce crisis, participants felt strongly that AI cannot be the sole answer to address this "AI isn't good enough to do it by itself", especially as the introduction of AI presents new challenges and workload in itself.

There is a real risk that if AI is seen as the answer to the workforce crisis, that other solutions and policies may not be considered or developed.

3.1 Maintaining skills, alongside learning new things

The fear of deskilling the workforce was raised with arguments made that over-reliance on AI would reduce clinical skills. An emphasis on skills preservation is needed. To support this there needs to be a better understanding of how skills development, resulting in clinical expertise and acumen, can be maintained if tasks that generate people's capacity to determine normal from abnormal is done by technology.

The risk of eliminating relational knowledge, where knowledge is shared between colleagues and specialties, was also raised.

3.2 Keeping the job worth doing

"The workforce gets more of a buzz from being helpful and efficient, the workforce wants pride, enjoyment from work."

Low morale and burnout are already significant issues for NHS staff. How people feel about their job is key to workforce retention. Following the concern over "pop-up fatigue" participants felt that AI could reduce clinicians' job satisfaction. Although recognising that clinicians' experience may vary depending on speciality and personality – eg pathologists are quite 'techy' and this may influence how they view, and experience, new technology.

Emphasising and maintaining the relational aspects of care is, of course, essential for both staff and patient experience and outcomes. One participant noting "AI uses purely empirical tools" whereas within medicine and healthcare "we want to find some sort of ground truths" which AI cannot provide.

4. Transparency and trust

The importance of transparency is essential for many areas of AI deployment: staff and patients need to know what it is, along with how, and why, it is being used and for whom it works well. Deploying AI into healthcare clearly requires an evolution of workforce roles and responsibilities. In this context the "*explainability paradox*" was raised: clinicians may not need to understand the precise inner workings of an algorithm, yet they must meaningfully contextualise its output and be able to explain this to their patients. They need to understand how the AI performs for subgroups of patients and so may need to understand both the training data set limitations and the way in which the AI decides on what it labels positive vs negative. This may create an additional cognitive work and time. AI explanations are inherently incomplete and differ from human reasoning, in turn maybe making it hard for clinicians to grasp and explain to their patients.

4.1 Public trust and the patient journey

Explainability is also an issue for public trust - being transparent with patients about where AI is used and the datasets used in AI that affects their care is important. Furthermore, if their data will be used for ongoing AI training and updating will they be consulted, or will consent be assumed?

The patient journey is changing, direct to consumer tests are now commonly available, and patients are increasingly using online AI tools to interpret their results.

"There has been a shift from reliance on health care professionals, to driving your own personal healthcare"

The ability to explain risk probability of an AI interpretation/diagnosis vs that of a clinician, particularly if there is disagreement, will be important. To maintain trust and credibility, clinicians need to have this information available, and it needs to be easy for patients to understand.

The social contract between the public and the NHS is also shifting, AI must not deepen cracks in trust and confidence. Part of the reassurance needed includes maintaining, and showing, where the human is in the loop.

The fear of cyber-attacks was also raised and considered a fundamental concern that would need to be handled clearly and openly.

4.2 Health equity

The data used to train AI will determine how accurate and so how applicable it is for any particular population. AI risks exacerbating health equity by performing poorly for under-represented groups.

"Health equity, who is driving this part of the conversation? Zero commercial argument to look at those who will be left behind."

The other side to the AI and equity argument is that AI has the biases it is given through the training data that is used. It does not, however, hold prejudices as we often do and so has the potential to democratise decision making. Alongside this, given the potential for greater stratification in patient populations as larger data sets are collated and analysed, greater equity may be created as access and treatment are personalised.

Further inequity may be created given that the infrastructure needed to implement AI systems is not evenly distributed e.g. digital pathology is only fully up and running in 30% of labs across England. Where you live currently dictates whether you will have access to AI solutions – another postcode lottery.

Overall utility will always be a trade-off, as with previous technologies such as facial recognition and fingerprint applications.

5. Implementation, infrastructure and evidence

"If we are going to layer it on top of a failing system then how will it help?"

Underpinning all the above points is the need to address implementation barriers and evidence concerns. Implementation is "not a linear experience" and participants saw the need to learn from past implementation experience.

5.1 System readiness

Unsurprisingly, it was unanimously agreed that the NHS IT infrastructure remains largely poor and often outdated, causing anxiety when considering AI implementation. Without foundational investment in "unsexy back-end infrastructure", AI risks adding fragility to an already strained system. For example, whilst there has been adoption of AI in pathology in some places, widespread adoption is hampered by the lack of digital pathology in many labs across the country, making large-scale AI deployment currently unfeasible.

Local implementation and evaluations were seen as important, with practical support for training and infrastructure being essential. Also considered important was the need to have some idea of what post-pilot costs might be. An example quoted suggested that an AI licence might cost £65,000/month – participants agreed it was very hard to see where the money to support this type of expenditure might come from, even if productivity was demonstrated.

As AI is entering society at speed, regulations are struggling to keep up. This is compounded by the breadth of what AI is and can potentially do. The fear of under-regulation leading to unsafe deployment was expressed by participants.

5.2 Politics and economics

The 10 Year Plan highlights the potential for AI to transform healthcare, banking on it to both save time and money and be a solution to the workforce crisis. All pressing problems. This policy imperative runs alongside the increased linkage between the health and policies to support economic growth. Bringing together the life sciences and health agendas brings challenges: a shift in the social contract for health and debate about who is driving decisions. AI companies, and their private equity backers, want (and need) financial returns quickly. Perhaps at a speed that outstrips the generation of robust evidence and implementation. These returns, and government policy that wants international recognition and reduced cost of the welfare state, may create unrealistic hope. For example, using AI to read images in radiology.

"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."

Properly assessing time saving and the consequent impact on the numbers of radiologists needed requires a full understanding of the day-to-day reality of being a radiologist, and a proper determination of which bits of time are saved. Alongside this the full cost of implementation needs to be considered. This needs to include procurement costs, licences, ongoing quality assurance, dataset updates and infrastructure upgrades.

Economic models of the full cost of AI, including any disbenefits such as equity impact, are needed before claims of efficiency or cost saving can be substantiated.

5.3 Mind the 'evidence' gap

"There is lots of top end research on can it work – but less on does it work."

There is not yet generalisable, trustworthy evidence that AI will deliver on its stated promises. Economic evaluations and cost-effectiveness are weak and limited "real-world" implementation means there is a lack of robust clinical outcome data. Participants compared AI to any other new clinical intervention, which would not be widely rolled out without a strong evidence base. AI should be subject to the same processes determining can it work, does it work and is it worth it?

6. Implications

"Introducing something into an already troubled sector, creates more issues for workforce, patients, and interprofessional relationships."

- Stop and think – what do we need and where can AI help, it is not a "silver bullet". It should not be used to cover deficits in the workforce or the system. Time will probably be saved, but only if integrated thoughtfully, given existing workflows and infrastructure.
- Evidence needs to be more robust considering impact in the round including equity, relationships, skills and morale as well as full economic costs.
- Developers, patients and clinicians need to work together to create explainability resources, alongside training materials.
- The continued importance of relational care should continue to be highlighted and embedded in training and implementation.