NHS RACE & HEALTH OBSERVATORY

SEPTEMBER 2026

BRIEFING FROM NHS RHO ROUNDTABLE:

TOWARDS TRUSTWORTHY, EQUITABLE AI IN THE NHS

FOREWORD

The promise of Artificial Intelligence (AI) in health and care is considerable. Used well, AI can support earlier diagnosis, more personalised treatment, improved productivity and better use of the NHS's precious resources. But technology does not arrive in a vacuum. It is developed from data, implemented through systems and experienced by people whose care has too often been shaped by unequal access, poorer communication, lower confidence and worse outcomes. Unless racial inequality is addressed deliberately and from the outset, AI will not simply reflect these patterns; it risks embedding and scaling them.

As a summary of our July 2026 roundtable on this subject, this briefing is both timely and necessary. It asks a fundamental question for the NHS: how can AI be developed, procured, deployed and governed in ways that strengthen trust, improve care and actively reduce racial inequalities, rather than reproducing the harms already present in healthcare pathways, datasets and decision-making?

The answer cannot be left to technical assurance alone. Equity must be designed into every stage of the AI lifecycle, from the choice of problem a tool is intended to solve, to the data used to build it, the standards applied at procurement, the transparency offered to patients and the accountability mechanisms in place when things go wrong.

The NHS Race and Health Observatory exists to identify and tackle ethnic and racial inequalities in health and healthcare. Our role is to bring together evidence, lived experience and system leadership so that the causes of inequality are understood, named and acted upon. In relation to AI, it means supporting the NHS and its partners to move beyond broad commitments to fairness and towards inclusion: better ethnicity data, more representative evidence, meaningful community involvement, stronger procurement requirements, transparent monitoring and clear routes for escalation or withdrawal where harm is identified.

This briefing sets out a practical agenda for action. It reflects the shared responsibility of policymakers, regulators, NHS leaders, researchers, technology developers and community partners to ensure that digital transformation advances, rather than undermines, health equity. If the NHS is to become a world leader in AI-enabled care, it must also become a world leader in using innovation to close the gaps that have persisted for too long.

PROFESSOR CHAAND NAGPAUL CBE
GP PARTNER AND NHS RACE AND HEALTH OBSERVATORY BOARD MEMBER

PROFESSOR HABIB NAQVI MBE
CHIEF EXECUTIVE, NHS RACE AND HEALTH OBSERVATORY

AUTHORS

SARINDI ARYASINGHE
HEAD OF IMPLEMENTATION, NHS RACE AND HEALTH OBSERVATORY

MARSHA LA-PIERRE
IMPLEMENTATION INTERN, NHS RACE AND HEALTH OBSERVATORY

EDITORS

PROF MARK WONG, UNIVERSITY OF GLASGOW, FORMER CO-CHAIR OF DESIGN ADVISORY GROUP, ANTI-RACISM OBSERVATORY FOR SCOTLAND
DR NANDI SIMPSON, DIRECTOR, IMPLEMENTATION, NHS RACE AND HEALTH OBSERVATORY
RINI JONES, ASSISTANT DIRECTOR, POLICY, NHS RACE AND HEALTH OBSERVATORY
ROSE OBIANWU, ASSISTANT DIRECTOR, COMMUNICATIONS AND MEDIA, NHS RACE AND HEALTH OBSERVATORY

INTRODUCTION AND EXECUTIVE SUMMARY

AI-enabled healthcare offers significant opportunities to improve access, efficiency and outcomes for the NHS, but also raises unresolved questions about patient safety, accountability and equity. Without deliberate leadership, these technologies risk embedding and scaling the racial inequities already present in healthcare data, pathways and decision-making. Building trustworthy, transparent and anti-racist AI is therefore essential to patient safety, public confidence and equitable care.[1]

The roundtable identified challenges across the AI lifecycle, including data quality and representativeness, procurement, community involvement, transparency, regulation, monitoring and accountability when harm occurs. These risks are heightened by the pace of AI adoption, while mechanisms to assure equity and accountability remain underdeveloped.

> Building **trustworthy, transparent and anti-racist AI** is therefore **essential to patient safety**, public confidence and equitable care.

Participants also identified practical opportunities for action. Leaders can strengthen trust by embedding anti-racism requirements into procurement, involving racially minoritised communities from the outset, publishing AI performance by ethnic group, evaluating harms, and establishing clear routes for escalation, suspension or withdrawal.

Acting now would enable the NHS to shape AI in ways that are safer, more accountable and more equitable for the communities most likely to be harmed by poorly designed or governed systems. Done well, trustworthy and anti-racist AI can support better care for those least well served, while making health and care more inclusive and trustworthy for everyone.

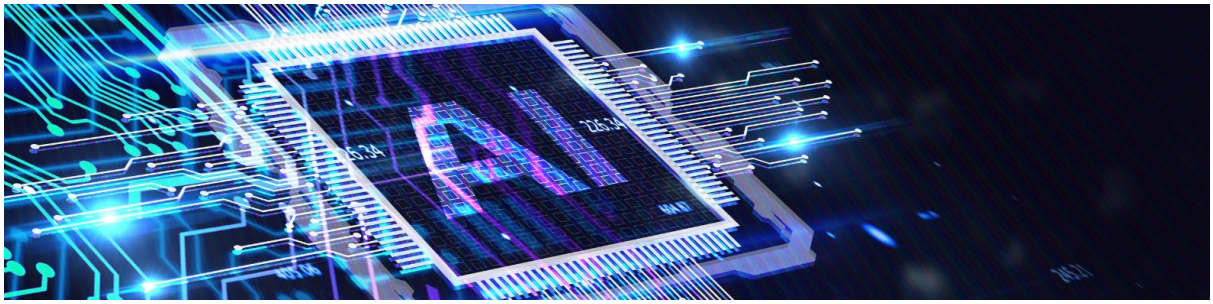

Current use of AI within healthcare settings

AI is already embedded in NHS clinical practice. Machine learning assists radiologists in detecting conditions such as breast cancer, stroke and lung disease by identifying abnormalities on CT, MRI and mammogram images that are difficult to detect quickly or at scale. In acute settings, predictive algorithms analyse vital signs and electronic health record data continuously identifying patients at risk of deterioration, including sepsis and cardiac arrest. Pathology, imaging and genomics services use AI to analyse large datasets and support more personalised treatment decisions.

AI also supports operational functions, including predicting admissions and discharges, optimising bed allocation and prioritising emergency department cases by clinical urgency. In primary and outpatient care, digital triage tools route patients to the most appropriate level of care. For clinicians, the fastest growing application is ambient voice technology, in which AI systems listen to consultations and generate structured notes, referral letters and discharge summaries in real time, potentially reducing documentation burden.

The effectiveness of each of these applications of AI depends on data quality, integration into clinical workflows, robust human oversight and ethical governance mechanisms.

The risk of amplifying racial inequities

The NHS has an opportunity to use AI to improve care for all communities. However, as with previous waves of innovation, AI also risks reproducing or amplifying existing racial inequalities where systems are designed, implemented and governed without sufficient attention to ethnicity recording, fairness, inclusion and trust. Bias enters AI systems through unrepresentative datasets, through variables that act as proxies for race and ethnicity, and through data shaped by longstanding inequalities in access, treatment and healthcare expenditure.[2] Once trained, tested and deployed, AI systems can embed and scale this bias across every decision they support.

The NHS therefore has a limited window in which to embed equity, transparency and accountability across the design, procurement, implementation and monitoring and governance of AI, before inequitable practice becomes established and expensive to reverse.

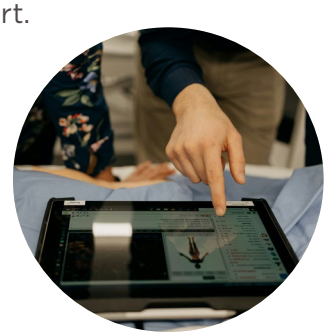

> AI is not creating new inequalities from scratch. It is **encoding and scaling the inequalities already present** in healthcare datasets, clinical pathways, service access, and decision-making systems.

Policy and regulatory context

The policy and regulatory landscape is evolving rapidly, with growing recognition of the need for safe, trustworthy and equitable AI. The Government's 10 Year Health Plan for England identifies the shift from analogue to digital as one of three fundamental changes required to reform the NHS, and commits to making the NHS the world's most AI-enabled health system, with AI integrated across all hospitals and clinical pathways within the Plan's lifetime.[3] While AI transformation offers considerable potential to improve patient outcomes and system productivity, it also increases the urgency of action to prevent it from reproducing or worsening existing inequalities.

Significant challenges remain around data quality, digital maturity, transparency, oversight and public confidence. As AI-enabled tools become embedded in clinical practice, questions arise about who is responsible for monitoring performance, identifying inaccuracies and responding when systems do not perform as intended. There is also continuing ambiguity about where responsibility and liability lie when AI-enabled systems contribute to harm, including the respective responsibilities of developers, suppliers, healthcare organisations and clinicians. The rapid growth of ambient voice technologies has highlighted the need for clearer approaches to ongoing assurance, so that responsibility for detecting errors does not fall disproportionately on frontline clinicians.

In response, the Government established the National Commission into the Regulation of AI in Healthcare to advise on a new regulatory framework. The Medicines and Healthcare products Regulatory Agency (MHRA) subsequently issued a call for evidence, bringing together patients, healthcare professionals, academics, industry and other stakeholders to examine the opportunities, risks and practical challenges associated with AI-enabled technologies.[4] Together these developments expose the tension at the centre of current policy: the NHS is expected to adopt AI at pace, while the systems required to assure its safety, effectiveness, equity and accountability remain underdeveloped.

ROUNDTABLE DISCUSSION

The NHS RHO convened a national roundtable in July, bringing together expert stakeholders from the Government, national health system bodies, regulators, NHS providers, industry, academia, patient organisations and community researchers involved in the use, development and oversight of healthcare AI. The following organisations were represented:

- Ada Love Lace Institute
- Data, Tech and Black Communities
- Department of Health and Social Care (DHSC)
- Health Innovation Yorkshire & Humber
- Imperial College London
- National Institute for Health and Care Excellence (NICE)
- National Institute for Health and Care Research (NIHR)
- National Voices
- Newtons Tree
- NHS England
- NHS Race and Health Observatory
- Medicines and Healthcare Products Regulatory Agency (MHRA)
- The Patients Association
- University of Glasgow

The group heard from expert speakers before taking part in a conversation café-style roundtable discussion. This section summarises the discussion, which explored the following questions.

Overarching Question:

What needs to be in place to identify racism in AI used in the NHS, take anti-racist action, and ensure accountability to the communities most affected?

The discussion was structured around three key questions:

1. Patient Trust, Safety and Transparency: what do racially minoritised patients and communities need to see from AI-enabled healthcare systems in order to trust and use them, and what should organisations prioritise to achieve this?

2. Anti-racism in AI Development and Implementation: how can the DHSC, NIHR, NHS organisations, researchers and industry ensure that the development and implementation of AI-enabled healthcare actively advances racial equity and improves outcomes for ethnically minoritised communities?

3. Governance, Accountability and Monitoring: what practical actions should the DHSC, regulators, NHS organisations and industry prioritise to identify, prevent and respond to racial bias and inequitable outcomes in AI-enabled healthcare?

For each question, participants were invited to consider:

- The strength and quality of the available evidence.
- Which recommendations should be prioritised for short, medium and long-term implementation.
- The expected benefits, of AI enabled healthcare including earlier and more accurate diagnosis, improved treatment, more equitable patient outcomes and potential cost savings for the NHS.

THEMES

The discussion was wide-ranging, reflecting both the complexity of the issue and the range of lived and learned experience in the room. The following themes aggregate views from several contributors rather than representing a final consensus position; where participants reached different conclusions, the range of views is set out.

Trust cannot be built on undisclosed use of AI, least of all with communities the NHS has already failed

AI is already used across NHS services, and patients are often not told how. The use of AI scribes in general practice provides an example of a technology adopted at pace, without patients understanding what it is or what is happening to their information. Rather than framing it as 'good practice', transparency must be recognised as an existing obligation under information governance and data protection principles.[5] Several participants cautioned against assuming that patients are comfortable with AI in their care and AI use should be a choice without compromising access to and quality of care.

National public engagement has confirmed that transparency, accountability and human oversight are widely expected of AI in healthcare, and these expectations carry particular weight for racially minoritised communities. Mistrust in this context is not the result of a general wariness of new technology; it is the accumulated product of discrimination, poor communication and concerns being dismissed in healthcare interactions, which the Observatory has evidenced and documented across NHS services, including maternity[6], mental health[7] and primary care[8]. These experiences have led racially minoritised communities to perceive use of data and AI in health as untrustworthy and likely to exacerbate discrimination[9]. Participants emphasised that patients are likely to encounter AI-enabled technologies through NHS services, even though the NHS does not control every stage of their development, supply and regulation. Building trust therefore requires a golden thread running throughout the pathway from developers and suppliers to regulators, NHS organisations and the public. Participants also distinguished trust from clinical safety: meeting regulatory standards does not automatically make a tool trusted by patients or professionals and a lack of trust may limit its uptake.

Transparency also determines whether the safeguards offered alongside AI are meaningful. Patients should retain the right to know, question and object to AI use, supported by an accessible human alternative[10]. Some participants emphasised that this right is fundamentally contingent on the human alternative being available within a comparable timeframe. Being presented with AI immediately or a clinician weeks later is not a genuine choice, particularly for patients and communities who continue to be disproportionately poorly served by the health system.

Involving racially minoritised communities requires resource, not invitation

Racially minoritised communities are rarely involved in decisions about the AI tools used in their care, and where involvement does happen it usually comes after significant decisions have been taken. Participants brought a range of perspectives on where responsibility for addressing this should sit. Some emphasised the importance of communities feeling equipped and confident to advocate for themselves, while others focused on the changes organisations must make so that communities are not expected to adapt to systems already designed without them. Both perspectives point to the same underlying requirement: ensuring that patient voice and influence can be exercised in practice and not only recognised in principle. Where the emphasis falls on individual advocacy alone, initiatives tend to work best for patients already confident in navigating the system, a pattern the NHS RHO has previously identified in efforts to strengthen patient voice[11].

Proposals from participants included:

- People-led panels convened before implementation decisions are taken[12].
- Local AI ambassadors, community champions and trusted community organisations and peer researchers as routes into communities that formal engagement does not reach.
- Public involvement in deciding what AI should and should not be used for, rather than decisions taken at executive level and consulted on afterwards.

Participants observed that appetite among racially minoritised communities to take part in AI research and innovation is strong when engagement is credible and meaningful. The main barrier to participation is that involvement is usually unfunded, which determines who takes part and for how long and reinforces racialised inequalities and exclusion. The length and depth of community involvement is especially important in AI-related research, as AI products continue to evolve after deployment, making sustained, long-term involvement essential. It was also stressed that early community involvement matters as much as the resource underpinning it, as many of the decisions that shape AI tools' performance for different ethnic groups have already been taken or neglected b developers and decisionmakers before they reach the procurement stage.

Procurement is where anti-racism requirements become binding

Participants argued that the decision about what AI is used for matters as much as the terms on which it is purchased. Priority areas and use cases should be based on needs and gaps determined by population health data, commissioned in line with national and regional priorities, and evaluated through a race equity lens that covers patient experience as well as expected clinical outcomes.

Both decisions, what to buy and on what terms, are settled at procurement. Once a product is deployed across an NHS organisation, changing it is slow, expensive and dependent on supplier cooperation. Embedding anti-racism at the point of procurement was the strongest point of agreement across the roundtable as it is the stage at which requirements can be enforced rather than encouraged, provided they are expressed as conditions that suppliers must meet.

Participants proposed procurement frameworks designed through a race equity lens, with patient and community engagement built into the process rather than added at implementation. Suppliers should be required to demonstrate, before adoption, that an AI system has been assessed for bias and performs equitably across ethnic groups and these should be monitored and evaluated regularly in the context of its use. Attendees argued that these requirements should be mandated nationally and applied to products already in use as well as to future procurements. Without a national standard, suppliers have no consistent expectation to build towards, and the equity of AI entering NHS services depends on the varying terms that NHS organisations can negotiate individually.

Participants prioritised and recommended how these requirements should be advanced, including:

- An independent AI regulatory authority providing transparent oversight and public assurance.
- A national kite mark signalling ethical procurement, supported by agile post-market surveillance.
- Review and stronger enforcement of existing equality legislation before further policy is created.

Attendees noted the EU has adopted AI-specific legislation[13] while the UK has adopted a more flexible approach[14], and that assessment will need to be iterative, since point-in-time approval is poorly matched to products that change after deployment.

Unrepresentative data embeds racial inequities into AI tools

AI tools perform according to the data on which they were trained and tested. Where those datasets do not adequately represent the populations on which a tool will be used, the resulting differences in performance are a feature of the tool itself, rather than a deployment issue that can be resolved through better implementation. Participants emphasised that AI systems perform differently across environments, and that fairness measured across an average population does not equate to fairness for any group within it. A tool can report strong headline accuracy while significantly failing a specific community, and yet headline accuracy is the metric made most visible by suppliers for purchasers to assess. Participants highlighted the importance of:

- Ensuring that AI tools are developed using diverse datasets that are representative of the populations a tool is intended to serve, alongside system-wide efforts to improve the routine collection and quality of ethnicity data within healthcare and research.
- Transparent publishing of any market testing and validation conducted during AI tool development, so that limitations are visible to adopting organisations rather than known only to suppliers.

Participants also contended that researchers and developers should be required to assess whether a design, model, or dataset holds racist assumptions or perpetuate "race" ideologies, rather than waiting for bias to surface downstream. The evidence base already demonstrates the consequences of failing to do so[15]. NHS RHO's review with the University of Nottingham found that Black, Asian and ethnic minority communities are poorly represented in genomics and precision medicine research, with distinct ethnic inequities embedded within the AI-based risk prediction tools used to calculate the chance of developing future disease[16]. In pulse oximetry, devices calibrated predominantly on light skin tones overestimate oxygen saturation in patients with dark skin tones, reflecting calibration data rather than any biological difference[17].

> "The Observatory's review with the University of Nottingham found that Black, Asian and ethnic minority communities are **poorly represented in genetic medicine research and in AI-based risk prediction** tools used to calculate the chance of developing future disease."

Participants noted that existing guidance describes what good looks like but does not establish clear thresholds for when a product should be refused. They proposed mandatory but proportionate equity assessment for AI before AI tools are adopted in NHS settings, including minimum testing requirements that ensure adequate representation of minoritised ethnic communities.

Anti-racism considerations cannot be sequenced for later

Health systems under delivery pressure tend to deal with core function first and equity later. However, participants were clear that this sequencing does not work for AI, as the decisions determining how a tool performs for different ethnic groups are made early. Training data, testing population and performance thresholds are all fixed before a product reaches the NHS. An assessment carried out after deployment can identify differences in performance, but they cannot undo the design decisions that created them. In the absence of requirements to address these issues earlier in the development process, deferral becomes the easier option under pressure.

Participants took two positions on the pace of adoption. Some held that adoption should proceed more slowly while development continues, since assurance mechanisms are not yet in place. Others held that trustworthiness should set the pace of adoption, because uptake depends on confidence and previous examples of rushed rollout have undermined trust. While framed differently, both positions are compatible, reflecting the same position that adoption should be determined by the strength of assurance and governance, not the maturity of technology alone.

Participants were also clear that the barriers to equitable AI are organisational, cultural and resource-related rather than solely technological. Treating racial equity as a purely technical assurance issue would overlook the organisational, cultural and structural factors that produce unequal outcomes. These outcomes therefore need to be evaluated through socio-technical assurance, using a race equity lens.

There is no agreed route for withdrawing an AI tool that causes racial harm

Where a medication or device causes harm, the NHS has established routes for reporting and escalating concerns and withdrawing the product from use. For AI-enabled tools no equivalent system exists, so in practice, a decision to withdraw a harmful tool would likely fall to whichever body first identified the problem, despite there being no agreed threshold for intervention. To address this gap, proposals from participants included:

- Transparent reporting routes accessible to staff and patients, with the Yellow Card Scheme raised as an existing model[18].
- Proactive post-market surveillance, rather than relying on incident reports reaching a threshold before action is taken.
- A predefined escalation and exit framework for every deployment, with clear criteria and thresholds for review, suspension or withdrawal.

Participants also raised product recall mechanisms and probationary periods during which clinical benefit must be demonstrated, with reference to NICE Early-use HealthTech guidance assessments[19]. Taken together, these proposals constitute a workable framework for identifying, monitoring, and responding to harm. However, responsibility for delivering such a framework is currently dispersed across manufacturers, providers and regulators with no single body accountable for the system overall.

Underlying the gap is a broader and unresolved question about responsibility and accountability. Participants expressed differing views on where responsibility should sit, including clinicians retaining accountability for AI-enabled decisions, named accountable officers within provider organisations, and shared responsibility across manufacturers, NHS organisations and regulators. Participants highlighted that, at present, the risk sits with the clinician using the tool, not the manufacturer that developed it.

RECOMMENDATIONS

The following recommendations were developed by roundtable participants and prioritised through voting for inclusion in this briefing. They are ordered by the stage of the AI lifecycle at which they take effect, beginning with procurement, where participants agreed requirements are most likely to become binding. Given the rapid deployment of AI in healthcare, the proposed delivery timeframes reflect the urgency of action.

Recommendation	Responsible Owner	Timeframe
1. Make anti-racism requirements determinative in AI procurement. Mandate national requirements across the AI lifecycle, expressed as procurement conditions rather than guidance and applied to products already in use as well as future purchases. To change what is bought rather than what is documented, requirements should specify performance data disaggregated by ethnic group; a defined standard below which a product does not proceed; and sufficient weight in scoring to affect the selection outcome or product withdrawal.	DHSC and NHSE, with regulators, supported by the NHS RHO	Within 36 months
2. Direct AI procurement and development to respond to population health need. Identify priority areas and use cases from needs and gaps identified through population health data, commission in line with national and regional priorities, and require evaluation against both clinical outcomes and patient experience through a race equity lens during development and implementation.	National and regional commissioners	Within 24 months
3. Embed communities in the co-design of procurement, evaluation and implementation of AI. Involve racially minoritised patients and communities so that public voice, digital exclusion, historical distrust and data concerns are addressed from the outset. Involvement should be funded as a delivery cost rather than absorbed by community organisations and sustained through continuous national and local conversation rather than convened once.	Cross-system, led by the NHS RHO with ICBs, regulators, patient and community groups, and industry	Within 12 months

4. Require disclosure of AI use and support it with patient-facing education to make informed choices. Set a consistent national expectation that patients are told when AI is used in their care, what it does and what its limitations are, delivered in culturally competent formats through trusted messengers, with accessible feedback routes for patients from diverse backgrounds.	NHSE and NHS organisations, with CQC	Within 12 months
5. Publish AI performance by ethnic group and set the metrics with communities. Extend the existing expectation that NHS bodies routinely publish health inequalities information to cover the performance of AI-supported decisions, with assurance metrics agreed between national regulators and community organisations, drawing on academic evidence and lived experience.	National regulators with community organisations, supported by NHSE	Within 36 months
6. Require a predefined escalation and exit route for every deployment and resolve where accountability sits. Every AI system deployed in NHS settings should carry an agreed framework setting out the outcomes that trigger review, suspension or withdrawal, including inequitable performance and discriminatory process/ outcome by ethnic group, and who holds the decision at each stage. The allocation of accountability between clinicians, provider organisations, manufacturers and regulators should be settled as part of this.	DHSC and NHSE, with regulators	Within 36 months

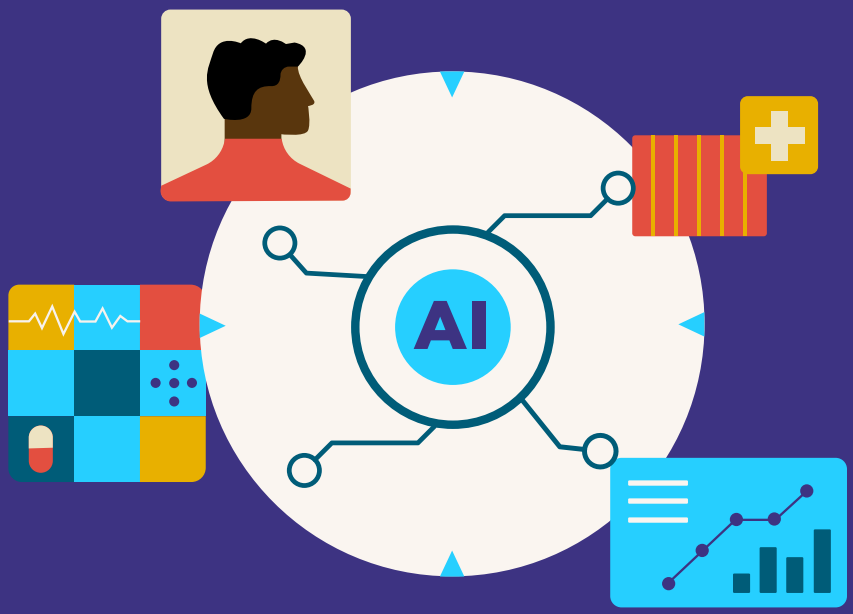

CONCLUSIONS

Without deliberate system leadership, AI-enabled healthcare risks embedding and scaling racial inequities rather than helping to reduce them. Ensuring that AI is trustworthy, transparent and anti-racist is therefore not only an ethical obligation, but a core requirement for patient safety, public confidence and equitable care.

The issues highlighted in this roundtable span the full AI lifecycle, from the quality and representativeness of data to procurement, community involvement throughout the AI lifecycle, transparency, regulation, monitoring and accountability when harm occurs. Taken together, they present a significant challenge for the NHS and its partners, particularly as AI is adopted at pace while the mechanisms needed to assure equity and accountability remain underdeveloped.

However, the roundtable also identified clear and timely opportunities. By placing anti-racism requirements into procurement, involving racially minoritised communities from the outset, publishing performance by ethnic group, evaluating harms, and establishing clear routes for escalation and withdrawal, leaders can shape AI in ways that strengthen rather than weaken trust.

If these opportunities are acted on now, the NHS can build an approach to AI that is not only safer and more accountable, but more equitable for the communities most likely to be harmed by poorly designed or poorly governed systems. Done well, trustworthy and anti-racist AI can support better care for those who have been least well served, while making health and care more inclusive, equitable, and trustworthy for everyone.

REFERENCES

[1] Artificial intelligence in healthcare and medicine: clinical applications, therapeutic advances, and future perspectives | European Journal of Medical Research | Springer Nature Link

[2] Weighing the benefits and risks of collecting race and ethnicity data in clinical settings for medical artificial intelligence - The Lancet Digital Health

[3] 10 Year Health Plan for England: fit for the future - GOV.UK

[4] National Commission into the Regulation of AI in Healthcare: research, engagement and call for evidence findings - GOV.UK

[5] Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data (United Kingdom General Data Protection Regulation) (Text with EEA relevance)

[6] Ethnic Inequalities in Healthcare: A Rapid Evidence Review, NHS RHO

[7] Trauma Informed Care and Racialised Communities, NHS RHO

[8] Patient Experience and Trust in NHS Primary Care, NHS RHO

[9] Minoritised ethnic people's security and privacy concerns and responses towards essential online services

[10] The need for patient rights in AI-driven healthcare – risk-based regulation is not enough.

[11] Briefing From HSSIB RHO Roundtable: Bias and Discrimination in Patient Safety Investigations, NHS RHO

[12] Valuing lived experience and co-design solutions to counter racial inequality in data and algorithmic systems in UK's digital services

[13] http://data.europa.eu/eli/reg/2024/1689/oj

[14] Artificial Intelligence in Healthcare: Current Regulatory Landscape and Future Directions | British Journal of Hospital Medicine

[15] Beyond productivity: a NASSS-informed review of implementation risks and research priorities for ambient AI scribes in healthcare | BMJ Digital Health & AI

[16] Ethnic Inequities in Genomics and Precision Medicine, NHS RHO

[17] Pulse oximetry and racial bias: Recommendations for national healthcare, regulatory and research bodies, NHS RHO

[18] Yellow Card | Making medicines and medical devices safer