

RAi CONNECT

Empowering and Connecting the Responsible AI Ecosystem

Welcome to RAI Connect, our monthly newsletter.

Six steps to building trust in AI in Healthcare

A new study from [Responsible Ai UK Health and Social Care Working Group](#) and the [Policy Institute at King's College London](#) sets out how academic health partnerships and national research communities can support the trusted deployment of AI in the NHS. **[How much do I need to know about AI in the NHS? Six steps to building trust](#)** reports on a policy roundtable that held last month convening fifty representatives from the National Commission into the Regulation of AI in Healthcare, the NHS, academia, industry, policy experts, the voluntary sector, and people with lived experience. Together, they co-produced six trust-building steps:

1. Telling AI's story from the patient and service-user perspective
2. Building an AI-educated and AI-enabled workforce through peer networks such as the [RAi NHS Champions](#)
3. Prioritising AI that tackles health inequalities and treats data quality as foundational
4. Deployment embedding user-centred input across the full AI development pathway
5. Establishing clear guidance on accountability when things go wrong
6. Standardising the assessment and governance of AI tools to reduce variation in adoption decisions.

This is urgent work. The sustainability of health systems globally is severely threatened... supply no longer meets demand and if AI isn't the solution to the supply side of the equation, then what is?

Another study from RAI UK, [King's Health Partners](#) and the Policy Institute at KCL '**[The use of AI in UK Healthcare](#)**' finds one in seven people have used AI instead of seeing a GP. *The public are turning to AI chatbots for health advice, and many say it has shaped their decisions about whether to seek professional care.* **[Read the full study and our recommendations.](#)**

The use of AI in UK healthcare

Public perceptions and healthcare priorities

We're hiring a Head of Ecosystem Engagement

We are recruiting a [Head of Ecosystem Engagement](#) to lead operational delivery of its national AI and Robotics Accelerator, run in collaboration with Future Worlds. The post sits within the University of Southampton, and spans research administration, stakeholder engagement, and strategic programme management across RAI UK's Research, Advocacy, and Engagement. The role is part-time (0.5 FTE), fixed-term to October 2027, and for someone who combines high-level organisational capability with the relationship-building instincts needed to work across academia, industry, and investment communities. **Applications closes on 03 June 2026.**

RAi JOIN OUR TEAM

HEAD OF ECOSYSTEM ENGAGEMENT

CLOSING DATE:
3 JUNE 2026

www.rai.ac.uk

Upcoming Webinar RAI UK Spotlight: Impact in Action, 8 June 2026

How does responsible AI research translate into real-world change? Join us on [Monday 8 June, 2:00–3:00 pm](#), for the next in our new [Spotlight webinar series](#), where RAI UK Enterprise fellows, will explore how they have translated their research projects into real-world products and service. It's a chance to hear directly from researchers about what impact looks like when responsible AI meets healthcare, sustainable agriculture and autonomous systems, and to ask the questions that don't usually make it into the published outputs. Three further Spotlight sessions follow through June and July. **Register here. Visit our website to view all upcoming events.**

RAi

RAi UK Spotlight IMPACT IN ACTION

LIVE WEBINAR | 8 June 2026 @2PM BST

Prof Zoe Kourtzi
University of Cambridge

Prof Lars Kunze
University of the West of England (UWE) Bristol

Dr Caitlin Bentley
King's College London

REGISTER NOW

rai.ac.uk

Policy Brief: Don't let AI do the policing: a framework for evaluating LLMs in law enforcement

Before a police force deploys a large language model, how should it know whether it is actually fit for purpose? The question drove a [new policy brief from PROBAbLE Futures](#), produced with the National Police Chief's Council and the Home Office. Drawing on a November 2025 multi-stakeholder workshop, which included a live hackathon testing real systems, it emphasises rigorous testing, representative datasets, operational oversight, and continuous AI assurance to support accountable and evidence-based adoption in policing. *By Nneoma Ogbonna, Mackenzie Jorgensen, Muffy Calder and Marion Oswald, (PROBAbLE Futures, May 2026. [Read the full brief here.](#)*

PROBAbLE Futures Policy Brief

Don't Let AI Do the Policing: Findings from a Multi-Stakeholder Workshop on Evaluating Large Language Models for Policing

Nneoma Ogbonna, Mackenzie Jorgensen, Muffy Calder, and Marion Oswald

1. Introduction: The Rationale for Our LLM Evaluation Workshop
The Home Office has announced reforms in UK policing, including the creation of the National Centre for AI in Policing (PoliceAI), driving a major expansion in AI adoption across UK policing.¹ Police forces are already increasingly adopting a range of AI tools, including large language models (LLMs), across multiple stages of the criminal justice process,² and the national-level rollout is expected to accelerate and expand this adoption. In policing, LLMs could be used to transcribe and translate interviews, generate witness statements, prioritise and triage information and intelligence, and generate summaries to support investigations and charging decisions.³

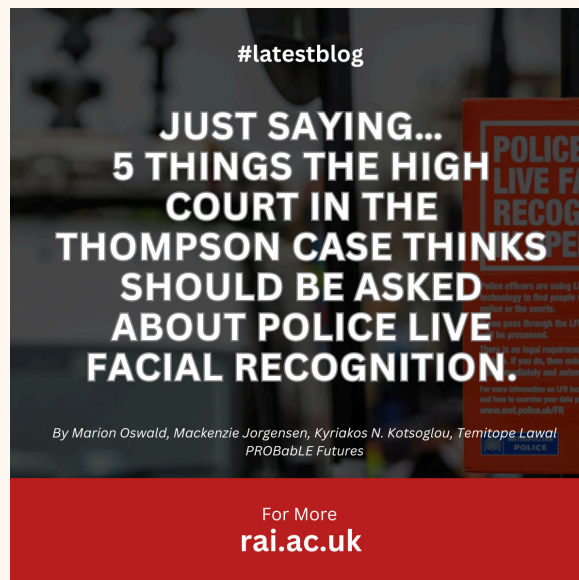
The adoption of LLMs in policing introduces several risks, such as omissions, misinterpretations and erroneous outputs.⁴ These risks can affect case progression, compromise the integrity of evidence and ultimately undermine procedural fairness and public trust in the criminal justice system.⁵ The high-stakes nature of policing, where errors can have profound consequences, requires that AI adoption be approached responsibly,⁶ and an integral pathway to this is through continuous and holistic evaluation of LLM design, development and adoption within policing workflows.

In November 2025, PROBAbLE Futures, the National Police Chief's Council, and the Home Office ran a 1.5 day workshop with stakeholders from academia, government, industry, and policing to discuss what is needed in a LLM evaluation framework for policing use cases, and to test the actual systems in practice through a hackathon. LLM evaluation here refers to the process of assuring LLMs are fit for purpose and testing how well they work within policing workflows to support responsible adoption. The aim of the workshop was to develop and gather the information required for practical guidance to support police forces to evaluate LLM systems, drawing on the insights and recommendations of stakeholders. This guidance will inform the rigorous testing of LLMs used in policing, in the interest of justice. Rigorous testing is required for evidence-based decisions about AI adoption.

We proceed by highlighting the LLM use cases discussed and tested during the workshop in Part 2. In Part 3 we present the substantive insights, grouped into five themes: (1) Importance of Robust and Representative Data, (2) Operational Scenarios and Consent, (3)


'Just saying': what the High Court's facial recognition ruling didn't decide

Shaun Thompson, a youth worker misidentified by police live facial recognition (LFR), lost his [judicial review case against the Metropolitan Police](#) last month. The court ruled the Met's deployment policy met the standard of legality, but left five significant questions conspicuously unresolved: the absence of systematic data on racial misidentification rates; the adequacy of testing under real-world conditions; the repeated use of identical large watchlists; the targeting logic of those lists; and who will independently scrutinise operational practice, not just policy on paper. The authors' examines the wider accountability and governance challenges surrounding the policy. *Our latest blog by Marion Oswald, Mackenzie Jorgensen, Kyriakos N. Kotsoglou and Temitope Lawal (PROBable Futures), [Read in full.](#)*



Assessing the Impact of Artificial Intelligence on Gender Disparities in the Labour Market

This [report from the Open University and Fawcett Society](#), examines how AI may affect gender disparities in the labour market. The study finds that female-dominated occupations show greater exposure to LLMs, while broader AI innovations remain concentrated in male-dominated roles. The report also highlights risks for lower-paid women workers and calls for governance safeguards, targeted upskilling, gender-disaggregated AI monitoring, and stronger AI literacy initiatives to prevent widening labour-market inequalities. Lead author Miriam Fernandez, discussed the findings in a [RAI UK webinar](#), recording available on [this page](#). *By Miriam Fernandez, Ángel Pavón Pérez, Davide Ghia, Damiano Giallongo, Daniele Quercia, Tania Cerquitelli and Maryan Yaqub, Open University / Fawcett Society, 2026.*




How can we best evaluate agentic AI?

Most existing tools for evaluating AI were built for models that stay still. Agentic AI systems, that act autonomously over time, pursue open-ended goals, and interact with users and environments, it demand something entirely different. **This [workshop recap from Brookings and Carnegie Mellon University identifies four critical gaps:](#)** no shared definition of what an "AI agent" even is; benchmarks that bear little resemblance to real-world deployment; the need for domain-specific, sociotechnical evaluation frameworks; and the absence

of credible, independent safety assessments. The authors make the case that evaluation reform is not a technical afterthought, it is a prerequisite for governance.

By *Elham Tabassi, Ramayya Krishnan (RAI International Ambassadors Council Member) and colleagues, Brookings Institution, April 2026.*



Podcast: The Age of AI: From UK to Silicon Valley : The AI Gap Explained

Fresh from San Francisco, Ben Clark joins Lara Lewington to ask a pointed question: while billions pour into OpenAI, Anthropic, NVIDIA and Scale AI, is the UK quietly watching from the sidelines? This week's episode of The Age of AI unpacks the gap between a Valley culture wired for speed and scale, and a European ecosystem that produces world-class researchers but struggles to back them at frontier scale. From ElevenLabs to longevity research, the conversation asks whether Britain's culture holds founders back, and whether this moment could dwarf even the internet boom. ***New episodes every week, tune in here.***



UKAIRS 2026: The national AI research community to convene in Edinburgh

The UK AI Research Symposium returns this November, and the call for submissions closes on 08 June 2026. Organised in collaboration with UKRI, RAI UK, the national AI Hubs and CDT centres, UKAIRS 2026 is the premier annual gathering for the UK's AI research community, a space to share work, cross-disciplinary collaborations, and help shape national research priorities in AI for societal and economic impact. Submissions are welcomed across a range of themes, including AI for science, health, engineering, sustainability and human-AI collaboration. *The symposium takes place on 24-25 November at the John McIntyre Conference Centre, University of Edinburgh. Registration opens 1 June 2026. For submissions visit the webpage.*



Opportunities

An update on latest opportunities, funding, projects or activities.

- **Edge AI Pump Priming 2026** — The EPSRC National Edge AI Hub is offering grants of up to £10,000 for up to 12-month projects exploring Edge AI from non-technical perspectives. Closes 1 June 2026.
- **Sovereign AI Strategic Assets Grants Programme** — Government grants supporting strategic AI assets and sovereign capability in the UK. Closes 5 June 2026.
- **Frontier Artificial Intelligence Discovery** — Funding for frontier AI research pushing the boundaries of discovery and scientific understanding. Closes 10 June 2026.
- **Talos Network: Policy Leaders Programme** — A programme for professionals ready to shape European AI governance, combining a virtual bootcamp, a week-long Brussels summit, and an optional 12-month paid placement. Applications open for Autumn 2026 cohort. Closes 14 June 2026.
- **FOAK26: Contracts for Innovation** — First of a Kind 2026 programme funds innovative procurement contracts helping public sector bodies adopt novel technologies and solutions for the first time. Closes 24 June 2026.
- **CfP: Responsible Use of AI in Art Creation and Archival Practice** — AI & Society journal seeks submissions exploring ethical and responsible AI applications across arts and archival contexts. Deadline 30 June 2026.
- **ARIA: Safeguarded AI – Mathematics for Safe AI** — ARIA's Safeguarded AI programme funds foundational mathematical research aimed at building AI systems whose behaviour can be formally verified and guaranteed. Closes 01 July 2026.

We'd love to hear your thoughts, stories or views.

Get in touch

Follow us



Copyright (C) 2024 Responsible Ai UK. All rights reserved.

This email was sent to K.Mittal@soton.ac.uk

[why did I get this?](#) [unsubscribe from this list](#) [update subscription preferences](#)