THE UNIVERSITY of EDINBURGH

GOVERNING THE FUTURE

Recommendations from the Edinburgh Data and AI Exchange

Edinburgh Futures Institute
26 March 2026

Foreword

Artificial Intelligence is rapidly becoming part of the fabric of public life, and will increasingly influence how public services are delivered, how resources are allocated across society, and how people engage with policy and politics. This transformation poses a challenging task for Scotland and the UK: how can we build the capability to use AI well, with the confidence to govern it openly, ensuring it strengthens prosperity, public services and democratic life?

At the University of Edinburgh, the oldest civic university in the English-speaking world, we embrace our responsibilities to contribute to these conversations. Since 1963, we have been at the heart of pioneering AI research and its teaching.

If AI is the engine of the future, then data is the fuel that powers it. Our computer and data infrastructure, data assets, and trusted relationships with industry, government, and civic society have developed over the past decade through our Data-Driven Innovation initiative. This has allowed us to build a unique ecosystem bringing disciplines together to advance AI, data and compute capabilities, and address questions around ethics, regulation, responsible use and impact.

We are also deeply aware that answering these questions requires a breadth of perspectives that goes beyond ours. This is why the Data and AI Exchange was designed as a deliberative event where participants from local and central government, the healthcare sector, finance, academia, civic society and the wider public worked through questions on AI infrastructure, data use and storage, adoption, applications, ethics and

governance. Rather than repeat familiar positions, the Exchange tested what practical priorities emerge when different forms of expertise and experience come together.

A great deal was achieved in a single day. Concrete priorities were identified for national AI infrastructure, data and AI skills, AI adoption in healthcare, public benefit from data, children's online safety, and environmental and democratic accountability. This report does not cover every dimension of AI policy - questions around national security, for instance, are of growing importance and warrant their own dedicated attention. What the Exchange did surface is captured here.

The Exchange sits alongside the University's *AI for Government* initiative, an ambitious programme supporting public bodies in navigating AI adoption, innovation, training, policy and governance. This ambition is practical as well as analytical: to help governments at all levels move beyond exploratory pilots and fragmented governance towards beneficial, scalable and publicly legitimate use of AI.

Scotland and the UK have the assets to shape an AI future, supported by responsible use of data that strengthens economic prosperity and human flourishing through innovative, trusted and socially responsible approaches.

Doing so will require investment, openness, technical excellence and democratic imagination. Our aspiration is to be a partner in that task; helping build the knowledge, infrastructure, skills and public confidence needed for AI and data to serve society.

Professor Michael Rovatsos
Chair of Artificial Intelligence and Dean of Research and Innovation, College of Science and Engineering, University of Edinburgh

Executive summary

99	3	13	16
Cross-sector participants	Themed sessions	Discussion tables	Recommendations

The Data and AI Exchange, held at the Edinburgh Futures Institute on 26 March 2026, brought together 99 representatives from across Scottish and UK public life to work through what both the UK and Scottish governments' AI ambitions require in practice. This report sets out what that process produced - the recommendations that emerged, the consensus that underpinned them, and the questions that remain unresolved.

Three themes. One underlying question.

Across three sessions covering infrastructure, adoption and applications, and ethics and governance, a recurring theme emerged: not whether AI should be developed, but on whose terms, under whose oversight, and to whose benefit.

On infrastructure, the UK's national supercomputer, based at the National Supercomputing Centre operated by EPCC at the University of Edinburgh, is a nationally significant asset already delivering substantial economic returns. Participants called for that recognition to be matched with long-term funding commitments that keep the UK internationally competitive, a significantly longer-term planning horizon for AI infrastructure investment, and genuine public accountability for who benefits from large-scale AI and data infrastructure. On sovereignty, the session concluded that while the longer-term question of domestic infrastructure is worked through, the UK cannot afford to be absent from foundational AI development happening at a European level.

On adoption and applications, participants were clear that sustained investment in data and AI skills is an industrial strategy necessity, not a discretionary education spend. The Data Skills Gateway at the University of Edinburgh, which has engaged over 130,000 people across Scotland, demonstrates both what is possible and what is at risk if that investment is not maintained. Health was identified as one of the most significant and time-sensitive areas for AI adoption, with participants calling for a national framework distinguishing between administrative, diagnostic, and interventional applications, with proportionate standards for each, and for clearer accountability when AI-assisted decisions go wrong.

On ethics and governance, participants called for a standing citizens' assembly on AI and society - not a one-off event, but a sustained, resourced process through which the people most affected by AI decisions, including children, have a genuine voice and ongoing role in shaping them. Equally consistent was the demand that access to public data, particularly NHS Scotland health data, be governed by clear public benefit requirements, with consent standards that are specific, informed, and freely given.

Key recommendations

The following recommendations, and those throughout this report, are presented by the University of Edinburgh, shaped by the discussions and conclusions of the Data and AI Exchange, and do not constitute a verbatim record of participant outputs.

UK Government

1. Establish a standing citizens' assembly on AI and society, sustained and properly resourced, with the University of Edinburgh as its institutional partner
2. Match the designation of the UK's national supercomputer, based at the National Supercomputing Centre operated by EPCC at the University of Edinburgh, as critical national infrastructure with long-term funding commitments that reflect that status and keep the UK internationally competitive
3. Engage as a co-funder in any EU foundational model development, ensuring UK and Scottish values, culture, and priorities are represented in its design

Scottish Government

1. Commission a long-term AI infrastructure strategy on a significantly longer-term planning horizon, modelled on the approach taken to energy and transport
2. Provide sustained investment in data and AI skills as an industrial strategy commitment, not a discretionary education spend
3. Develop improved mechanisms and guidance for commercial access to Scottish health data that will ensure a triple bottom line of societal benefit and contribution to scientific discovery alongside commercial benefit.

Introduction: Why this conversation, why now?

The UK Government's Industrial Strategy names AI as an economic priority, with its £500 million Sovereign AI programme signalling a clear intent that the UK should be an AI maker, not an AI taker. In Scotland, the Scottish Government has this year (2026) published its new comprehensive AI Strategy.

Taken together, these commitments represent a genuine and significant moment - a recognition that **AI will shape Scotland and the UK's economic future, its public services, and its place in the world**; with decisions made now determining the country's future prosperity, security and global role in the technology sector.

The University of Edinburgh is already working to make that ambition real. Its *AI for Government* initiative, which brings together over seventy researchers across disciplines to collaborate with and support government and public sector organisations, reflects a long-standing conviction that AI's potential is only unlocked when those who develop it work closely with those who govern and apply it.

Translating policy ambition into practice, however, requires answers to hard questions. How will these strategies be resourced in ways that keep the UK internationally competitive? Who will govern them? How will the public understand what is being built in their name, and on whose terms? These are not technical questions. They are political ones. And they are precisely the questions that the Data and AI Exchange was convened to work through.

The Stakeholder Assembly format was chosen deliberately. Designed to create space for genuine thinking across different perspectives and sectors, it draws on the methodology of structured public deliberation rather than standard consultation. Participants were not asked to validate any single institution's position. They were asked to bring their own expertise and experience to bear on questions that will shape society now and for generations to come.

What follows is an account of what the process surfaced - the consensus that emerged, the tensions that remained, and the unresolved questions that governments now have the opportunity to act on.

Professor Oliver Escobar
Chair of Public Policy and Democratic Innovation at the University of Edinburgh

"The Data and AI Exchange was deliberately slowed down, analogue, face-to-face, human-paced - and that is not a footnote to this report. It is central to its argument.

When things move at AI speed, when everything is optimised for efficiency and output, something essential gets lost; genuine human deliberation about what matters. That is where real thinking happens. And it is irreplaceable."

Theme One: Infrastructure

> "The human brain, computing at extraordinary scale, consumes just 20 watts of energy. The AI systems being deployed at national scale consume vastly more - and that gap is growing."

Professor Themis Prodromakis
Regius Chair of Engineering at the University of Edinburgh

The infrastructure session examined three interconnected questions: how the UK and Scotland should invest in and govern large-scale AI and data infrastructure; what sovereign AI means in practice and whether it requires sovereign infrastructure; and how to build the public legitimacy that large-scale data centre development will require.

The economic case for sustained investment is already being made. An independent evaluation found that ARCHER2, the UK's current national supercomputer based at the National Supercomputing Centre at the University of Edinburgh, has delivered £4.3 billion to the UK economy against a total project spend of approximately £105 million - a return of £8.30 for every pound invested.[1] The debate the session surfaced was not whether to invest, but how, on whose terms, and who gets to decide.

Image Credit: Chris Duguid

What sovereign AI means in practice

One of the central questions in the infrastructure session was whether Scotland and the UK need sovereign AI infrastructure, and if so, what that actually looks like. The discussion did not produce a settled answer. Participants recognised that sovereignty does not necessarily require owning every component of the technology required to build and run AI systems at national scale. You can have AI that reflects domestic values and is accountable to domestic law without building every piece of AI infrastructure from scratch. But the implication was equally clear that the current position, in which the UK relies heavily on infrastructure and models built around the assumptions and priorities of other nations, is not adequate for a country with serious AI ambitions.

Where the session did reach agreement was on an immediate and practical step. While the longer-term question of sovereign infrastructure is worked through, the UK should engage as a co-funder in any EU foundational AI model development, ensuring that UK and Scottish values, culture, and priorities are represented in models based on shared European cultural heritage. Models such as this will increasingly underpin public services and civic life across Europe and beyond. Participation now is the means of influence later.

Planning for decades, not spending cycles

Several contributions independently arrived at the same conclusion, that infrastructure planning must operate on long time horizons. The Scottish rail electrification programme was cited as a model - not because rail and AI are comparable technologies, but because it represents the kind of sustained, committed thinking that transformational infrastructure requires.

The consensus was that investment in supercomputing should be treated on a par with investment in roads or energy grids, as essential national infrastructure planned across decades and not subject to the cyclical pressures of short-term spending reviews. Tables consistently called for a significantly longer-term planning horizon, with the fifty-year commitments made to transport infrastructure cited as the benchmark.

Theme One: Infrastructure

The environmental case, and why it needs to be made publicly

There was scepticism among participants about whether environmental concerns would drive meaningful behaviour change from either government or industry. The discussion also surfaced a distinction that supercomputers and data centres are not the same thing. A data centre is a building designed to house many different types of computing equipment; supercomputers are purpose-built for maximum computational power, with energy efficiency central to their design. Historically, the same has been less true of cloud computing infrastructure, though both face the same fundamental environmental challenges.

What was clear across multiple tables was an appetite for transparency, alongside a recognition by participants that Scotland's access to renewable energy and natural water sources makes it well placed to host data infrastructure more sustainably than many alternatives. Publishing clear, accessible data on the environmental performance of Scotland's AI and data infrastructure was among the most consistently repeated actions across the infrastructure tables, and was seen as a prerequisite for building the public legitimacy that large-scale investment will require. Participants also called for any new commercial data centre development in Scotland to establish a community benefit trust, on the model already well established alongside renewable energy developments such as wind and solar farms, ensuring that the regions bearing the burden of large-scale infrastructure see tangible local benefit in return.

Image Credit: University of Edinburgh

Recommendations

UK Government

1. Match the designation of the UK's national supercomputer, based at the National Supercomputing Centre at the University of Edinburgh, as critical national infrastructure with long-term funding commitments that reflect that status and keep the UK internationally competitive

2. Engage as a co-funder in any EU foundational AI model development, ensuring UK and Scottish values, culture, and priorities are represented in its design

Scottish Government

1. Commission a long-term AI infrastructure strategy on a significantly longer-term planning horizon, modelled on the approach taken to energy and transport

2. Require any new commercial data centre development in Scotland to establish a community benefit trust, ensuring tangible returns for the communities that host it

3. Publish clear, accessible information on the environmental performance of Scottish AI and data infrastructure, including water usage, energy sourcing, and carbon footprint, to build the public legitimacy that large-scale investment requires

[1] London Economics for UKRI report - The scientific and economic impacts of ARCHER2 (November 2025)

Theme Two: Adoption and applications

> "The real issue is whether the NHS and other health systems have the agility, infrastructure, workflow capacity, staffing, training, and governance to actually turn what AI tells us into integrated care."

Professor Julie Jacko
Professor of Health Informatics and Data Science, Dean of Innovation and Engagement, College of Medicine and Veterinary Medicine, University of Edinburgh

The session on adoption and applications explored two interconnected themes: the implications of AI for education and workforce development, and the opportunities - alongside the accountability challenges - of adopting AI across public services. Underpinning both was a shared and unresolved question. Rather than focusing solely on what AI systems can do, participants returned consistently to what people and institutions will need in order to use these tools well. The organisational capability, the workforce capacity, and the governance structures required to integrate AI thoughtfully and responsibly, rather than simply deploy it, were recurring preoccupations throughout the day.

Health emerged as one of the dominant practical focal points of the discussion - not because it was prescribed in the agenda, but because participants consistently gravitated towards it as the context in which the consequences of poor AI adoption are most significant, and where the need to move beyond technical implementation to sustained, responsible adoption is most pressing.

Image Credit: Sam Ingram-Sills (Whitedog Photography)

The scaffolding problem

For generations, young people have developed analytical capability for senior roles by doing entry-level work; reading, summarising, drafting, checking, revising. Those are precisely the tasks AI now performs first. Several tables identified this as a structural challenge rather than a disciplinary one. The question is not whether AI is useful in education. It is how the developmental scaffolding works when the lower rungs of the ladder have been removed.

The responses that emerged were practical rather than prescriptive. Marking frameworks should be designed to reward visible improvement in pupil progress and independent thinking, ensuring that the value of original work is not undermined by easy access to AI-generated output. Equally, critical thinking and the capacity to interrogate, challenge and contextualise AI outputs should be treated as core educational objectives in their own right, as foundational skills for a generation that will live and work alongside these tools.

Building a workforce fit for an AI economy

The most urgent and consistently repeated finding across the adoption tables concerned the workforce. The Data Skills Gateway at the University of Edinburgh has engaged over 130,000 people across Scotland, spanning schools, colleges, universities, and the workplace. Programmes of this scale and reach demonstrate both what is possible and what is at risk. Participants were clear that sustained investment in data and AI skills is not an education budget question, but an industrial strategy one. A data and AI-literate workforce is the enabling condition for everything else in both governments' AI agendas. Without it, the infrastructure cannot be used well, applications cannot be governed properly, and governance frameworks cannot be understood or scrutinised by the public they are meant to serve.

Equally, participants stressed that investment in skills must be matched by investment in how those skills are developed. Explicit offline learning opportunities should be built into curricula as deliberate practice, and teachers must be adequately supported and trained to deliver AI education effectively. Scotland's curriculum review cycle presents a direct opportunity to get this right.

Theme Two: Adoption and applications

AI in health: the system matters more than the algorithm

Discussions on health consistently returned the clear and practical conclusion that effective AI adoption in clinical settings is not primarily a technology challenge, but a systems one. The central issue is less whether an algorithm performs well, and more whether health systems have the organisational agility, workforce capability, and service-design capacity to act meaningfully on what those tools can offer. Across the tables, there was strong agreement that AI should support and inform clinical judgment rather than replace it, and that patients and the public must be able to trust that a human remains meaningfully in the loop when consequential decisions are made. Participants stressed that frameworks for AI in health must reflect this distinction explicitly, with differing standards applied to administrative, diagnostic, and interventional uses, in line with their respective risks and responsibilities.

With Public Health Scotland forecasting an increasing burden of disease across Scotland, for example a 63% increase in COPD prevalence by 2044[2], and markedly higher mortality rates in Scotland's most deprived communities, participants pointed to the significant potential for AI-enabled care pathway redesign to improve outcomes at scale. The question of how far AI should go, including in settings such as surgery, did not produce a settled answer. Tables converged instead on proportionality as the relevant principle, with an approach grounded in context, risk, and impact, and careful attention to what meaningful human oversight requires in each specific clinical situation.

Image Credit: Callum Bennetts

Accountability gaps in health AI

Two structural gaps were identified that current policy does not adequately address. The first concerns the flow of knowledge between NHS health services and academic research institutions. While data is shared between both and research outputs are produced, findings do not always reliably return to inform clinical practice. Several tables characterised this dynamic as extractive, and suggested that data-sharing agreements should more clearly require translation of research outcomes back into service improvement. The second gap relates to accountability. When an AI-informed clinical decision contributes to harm, responsibility can be difficult to locate. Although Article 22 of GDPR establishes a requirement for human

involvement in significant automated decisions, participants highlighted the lack of clarity about what this means in practice in clinical environments where human and algorithmic judgment are increasingly intertwined.

Existing legislation was not designed with clinical AI in mind. Several groups therefore called for clearer interpretation and guidance; not to weaken human oversight, but to define what it should meaningfully involve in contemporary clinical settings.

Recommendations

UK Government

1. Review the regulatory approval pathway for AI-enabled health applications through the Medicines and Healthcare Products Regulatory Agency (MHRA), ensuring the framework keeps pace with the speed at which consumer-facing health tools are being deployed

2. Review the practical application of Article 22, GDPR's human-in-the-loop requirement, in healthcare and other high-stakes AI settings, and establish clear accountability for when AI-informed clinical decisions cause or contribute to harm

Scottish Government

1. Provide sustained investment in data and AI skills as an industrial strategy commitment, not a discretionary education spend

2. Integrate AI education into the school curriculum and provide adequate training and support for teachers to deliver it effectively

3. Establish a national framework for AI in health that explicitly distinguishes between administrative, diagnostic, and interventional uses, with proportionate standards for each, in line with their respective risks and responsibilities

4. Invest in ways of ensuring relevant findings from research using NHS Scotland data can flow back into the NHS to inform clinical practice

[2] Public Health Scotland, Scottish Burden of Disease: Future prevalence and burden of chronic obstructive pulmonary disease (March 2025)

Theme Three: Ethics and governance

> "Without data, there is no AI"

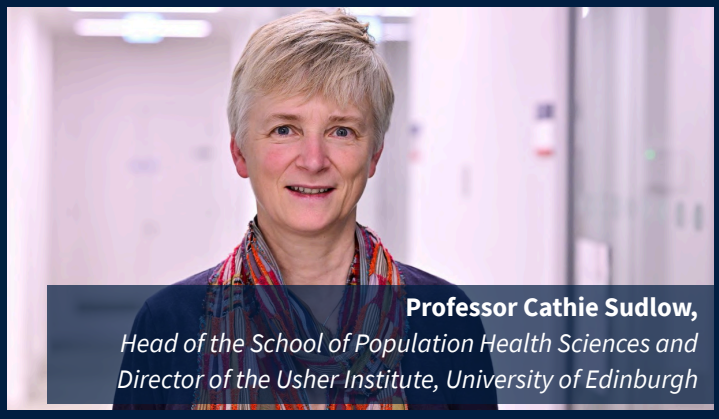

Professor Cathie Sudlow,
Head of the School of Population Health Sciences and Director of the Usher Institute, University of Edinburgh

The ethics and governance session opened out the broadest questions of the day: on what terms should powerful institutions be trusted with data about people's lives, and what role should the public have in answering that? Where the adoption session examined how AI tools are deployed in health settings, the governance session turned to a prior question: who controls the data those tools depend on, and under what conditions should access to it be granted?

Three sub-themes emerged from the structured discussions: the governance of health data, children's online safety, and the use of data analytics in public sector decision-making. Across all three, the same demand surfaced consistently that governance of AI be driven by values and public accountability rather than by the financial incentives that currently dominate.

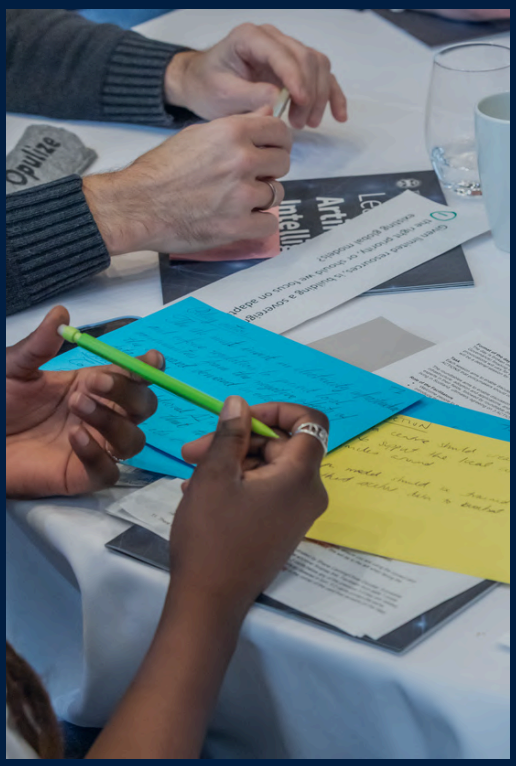

A standing citizens' assembly on AI and society

The most significant finding of the governance session, and of the day as a whole, emerged across multiple tables where participants independently called for a standing citizens' assembly on AI and society. The ask was not for a one-off consultation, but for a sustained, resourced, ongoing mechanism through which the people most affected by AI decisions have a genuine and continuing role in shaping them. The University of Edinburgh, with established expertise in deliberative methodology and a track record of working with government on participatory processes, would be well placed to serve as the institutional partner for such a mechanism.

Health data: defining who can access it and on what terms

Scotland holds a significant research asset in its longitudinal health and care data. The third session surfaced broad agreement on the principles that should govern access to data, as well as recognising ongoing technical and governance challenges that need to be resolved.

The purposes for which data can be used should be within a well-defined transparent framework. Tables discussed issues around differing consent models and the need for engagement through involving the public in discussions about how their personal data is used for research purposes. Some attendees expressed a view that citizens should be able to access their own records and exercise a right to amend them.

Where access by for-profit organisations was discussed, a triple bottom line test emerged from the tables as a practical framework. Any such access must demonstrate broader societal benefit and contribution to scientific discovery alongside commercial benefit. Commercial return alone is not sufficient justification.

Some attendees questioned whether corporate access should be permissible at all, or whether access should be reserved for academic institutions. This is a complex and nuanced issue and, while the session did not reach any firm conclusions on this point, it highlighted the need for transparency, robust discussion and guidance.

Another challenge that was raised was also left unresolved: how do you protect the privacy and confidentiality of service users whose data has been used to train an AI model in a secure data environment when deploying it as an operational model in the real world? This requires sustained research attention, which should then inform the development of access frameworks and guidelines.

Theme Three:
Ethics and governance

Including children in the design of their own online safety

Tables addressing children's online safety converged on a finding that was practical rather than aspirational. Children are digital natives who navigate online spaces that many adults, including policymakers, do not fully understand. They know how platforms work, how to navigate around restrictions, and how social dynamics operate within them. Designing protection without their input produces rules that do not fit the reality they experience.

Multiple tables called for children to be included as genuine participants in the design of online safety frameworks, not consulted after decisions have been made, but involved in the process of designing them. A number of tables also called for a cross-jurisdictional regulatory body, modelled on the Civil Aviation Authority or the MHRA, with agreed standards that apply regardless of where a platform is headquartered and meaningful consequences for non-compliance. Platforms that profit from operating in Scotland and the UK must comply with Scottish and UK child safety law. The view across tables was direct that if they do not, they should face bans or significant fines.

Data should inform public decisions, not make them

Deliberations addressing public sector use of data analytics converged on the principle that large-scale data analysis should help inform public sector decisions; however, it should not produce them. Subject expertise and human judgment remain necessary because context and values matter in ways that models cannot fully capture. Several tables pointed to the importance of government demonstrating trustworthiness through transparency rather than assertion, with Denmark cited as an example of how algorithmic decision-making can be made visible and accountable. The specific mechanisms proposed included mandatory impact assessments for government use of data analytics, public registers of algorithmic tools in use across government, and explicit limits on categories of decision that may not depend primarily on automated analysis.

Recommendations

UK Government

1. Establish a standing citizens' assembly on AI and society, sustained and properly resourced, with the University of Edinburgh as its institutional partner
2. Work with international partners to develop and enforce cross-jurisdictional consequences for platforms that profit from operating in the UK while failing to comply with child safety standards, treating regulatory compliance as a condition of market access
3. Commission research into the privacy risks of training AI models using secure data environments, and ensure the findings inform data access policy

Scottish Government

1. Develop improved mechanisms and guidance for commercial access to Scottish health data that will ensure a triple bottom line of societal benefit and contribution to scientific discovery alongside commercial benefit
2. Build explicit frameworks distinguishing between categories of public sector decision and the role data analytics may appropriately play in each, with mandatory impact assessments and a public register of algorithmic tools in government use

Conclusions

The Data and AI Exchange produced clear findings across three themes, with a high degree of consistency across tables drawn from very different sectors and perspectives. Where consensus existed, it is reflected in the recommendations throughout this report. Where it did not, the thematic sections say so honestly.

The immediate priorities that emerged are clear:

For UK Government, the asks centre on infrastructure investment that matches stated ambitions, regulatory frameworks for AI in health that keep pace with the speed of deployment, international engagement that ensures UK and Scottish values are represented in the foundational AI development happening at a European level, and a sustained, properly resourced mechanism for public deliberation on AI that gives the people most affected by these decisions a genuine role in shaping them.

For Scottish Government, the priorities are a data and AI-literate workforce treated as an industrial strategy necessity, governance of health data that puts public benefit first, integration of AI education into the curriculum, and clear environmental standards for data infrastructure that ensure large-scale investment commands genuine public legitimacy.

These recommendations are put forward in that spirit. Not as the final word, but as a considered set of asks, grounded in genuine cross-sector deliberation, that governments now have the opportunity to act on. Good policy on AI will not emerge from technical expertise alone. It requires the kind of broad, evidence-informed public conversation that the Data and AI Exchange was designed to start - and that this report is intended to carry forward.


SOVEREIGN
Ensure supercomputer continue to prioritise renewability and remain environmentally conscious

Acknowledgements

This report would not have been possible without the 99 participants who gave their time, expertise, and candour to the Data and AI Exchange on 26 March 2026.

The event was designed and facilitated by Professor Oliver Escobar, whose expertise in deliberative methodology shaped both the format of the day and the quality of what it produced. The facilitation team, drawn from trained student facilitators supervised by Professor Escobar, ensured that every table was heard.

The Data and AI Exchange was organised by University of Edinburgh, DDI, GAIL and Stakeholder Relations teams. The event was hosted at the Edinburgh Futures Institute, with thanks to Professor Marion Thain, Director of the Edinburgh Futures Institute.

Image Credit: Hufton & Crow

The expert speakers who opened each session and provided the evidential foundation for the structured discussions were invaluable in framing the discussions that took place and we thank them for their time. They were:

Infrastructure

Professor Themis Prodromakis, Regius Chair of Engineering and Director, APRIL, University of Edinburgh
Professor Mark Parsons, Director, EPCC, University of Edinburgh
Professor Alexandra Birch, Chair of Multilingual Natural Language Processing, University of Edinburgh

Adoption and Applications

Professor Fiona McNeill, Personal Chair of Computing Education, School of Informatics, University of Edinburgh
Theodore Pengelley, Digital Learning and Development Manager, Qualifications Scotland
Jamie Brogan, Head of Climate Partnerships, Edinburgh Climate Change Institute, Strategy Director, Scottish Climate Intelligence Service
Derek McGowan, Service Director, Housing and Homelessness, City of Edinburgh Council
Professor Julie Jacko, Professor of Health Informatics and Data Science; Chief Academic Officer of the School of Population Health Sciences and the Usher Institute; and CMVM Dean of Innovation and Engagement, University of Edinburgh
Professor Ewen Harrison (OBE, FRSE, FMedSci), Professor of Surgery and Data Science, Co-Director Centre for Medical Informatics, University of Edinburgh
Dr Mohsen Khadem, Director of the Surgical and Interventional Robotics Group, School of Informatics, University of Edinburgh

Ethics and Governance

Professor Cathie Sudlow, Head of the School of Population Health Sciences and Director of the Usher Institute, University of Edinburgh
Professor Deborah Fry, Professor of International Child Protection Research, University of Edinburgh, and Director of Childlight
Dr Morgan Currie, Senior Lecturer in Data and Society, School of Social and Political Science, University of Edinburgh

Appendix: Methodology

The Data and AI Exchange took place on 26 March 2026 at the Edinburgh Futures Institute, which tackles complex global challenges through interdisciplinary research, education and collaboration.

The day ran across three themed sessions: Infrastructure; Adoption and Applications; and Ethics and Governance. Each session followed an identical structure. Expert presentations introduced each theme. Participants then moved to thirteen discussion tables, each working from a different provocation - an open-ended prompt designed to surface assumptions, generate inquiry, and reveal where genuine disagreement exists. Tables worked in pairs first, creating the conditions for face-to-face conversation before expanding to full table discussion.

Each table produced three concrete policy actions, written on cards. These were posted to a shared wall and clustered in real time by the facilitation team, with duplicate actions marked to show where consensus was strongest. The day closed with each table nominating a speaker for a one-minute plenary summary, followed by Professor Escobar's synthesis of the main themes from the clustered wall.

Facilitators were trained student facilitators supervised by Professor Escobar, with flipcharts capturing key ideas from each table throughout the day. The action cards, flipchart notes, plenary summaries, and Professor Escobar's closing synthesis form the evidence base for this report.

Further reading on deliberative methodology

Escobar, O. (2017). Public Dialogue and Deliberation: A communication perspective for public engagement practitioners. University of Edinburgh, Edinburgh Beltane.

Escobar, O. and Elstub, S. (eds.) (2024). Climate Assemblies: New civic institutions for a climate-changed world. De Gruyter Brill.

World Health Organization (2023). Citizen engagement in evidence-informed policy-making: A guide to mini-publics. WHO.


THE UNIVERSITY of EDINBURGH
Generative AI Laboratory
A centre for excellence dedicated to researching all aspects of generative artificial intelligence in society.
THE UNIVERSITY of EDINBURGH
EDINBURGH INNOVATIONS
HARNESSING AI FOR BETTER HEALTH
Find out more
eil.ac/eiexpertise


HARNESSING
DATA, DIGITAL
AND AI FOR

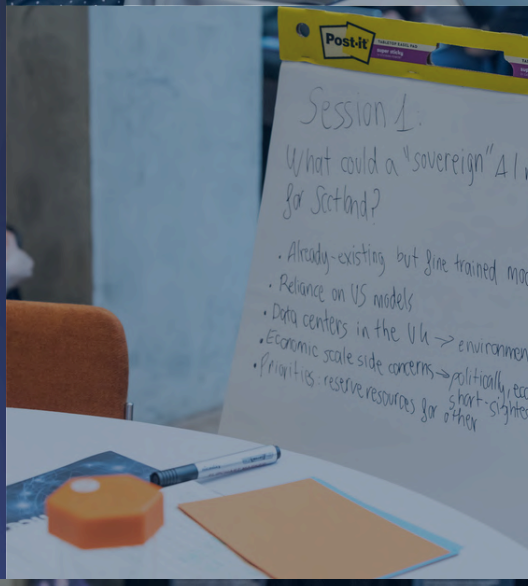
Post-it
Session 1:


HARNESSING DATA, DIGITAL AND AI FOR PUBLIC GOOD
HARNESSING AI FOR BETTER HEALTH
THE UNIVERSITY of EDINBURGH
www.ed.ac.uk
Let's be the Data capital of Europe
Innovation
Partnerships
Skills


Leading Artificial Intelligence

THE UNIVERSITY of EDINBURGH | HERIOT WATT UNIVERSITY

DDI Data-Driven Innovation

Part of the Edinburgh & South East Scotland City Region Deal

THE UNIVERSITY of EDINBURGH
Generative AI Laboratory

Commissioned by the University of Edinburgh with support from the Data Driven Innovation (DDI) initiative and the Generative AI Laboratory (GAIL)

www.ed.ac.uk/ai

Photography, event: Chris Scott
Photography, University of Edinburgh: Andrew Perry, Callum Bennetts, Chris Duguid, Hufton & Crow, Sam Ingram-Sills (Whitedog Photography)
Designed by: Hot Tin Roof

AI tools were used in the production of elements of this report

Published: May 2026