



Strategy

2025-28

Contents

- 01 Where we are: a note from our director
- 09 About us
- 14 What we have learned
- 19 What we are facing
- 27 Where we are going: objectives
- 38 How we work
- 43 How we make change

Where we are: a note from our director

Ada was launched in 2018 with a mission to ensure data and AI work for people and society. The process of building this strategy has been an opportunity to step back and examine how the landscape has changed and what has remained the same. We have explored which aspects of Ada's approach have worked best, and what is still needed for the future. And we have looked at how we can help move towards a world where technology is built to support the societies we want to create – rather than intertwined in our lives without examining the knock-on effects.

The current moment can often feel disorienting, with technical updates, advances and new use cases arriving at breakneck speed. This makes it easy to lose sight of the fact that AI is not just about technology. AI is people all the way down. People who are represented in the data. People who clean and annotate the data. People who sell, buy and use AI systems. People who are impacted by AI systems and by the tools built upon them.

AI is people all the way down. People who are represented in the data. People who clean and annotate the data. People who sell, buy and use AI systems. People who are impacted by AI systems and by the tools built upon them.

We have all heard the familiar refrain that the AI 'revolution' will improve people's lives by making public services easier to access, solving societal problems and boosting economic growth. Yet for all the hopes poured into AI technologies, they rely on and are limited by the quality, provenance and partiality of the data they are built on.

AI is sometimes portrayed as a free lunch – despite the hundreds of billions invested in it. But there are no magic beans and no crystal balls, and competing visions are jostling for dominance. We are seeing countries bet on the promise of Artificial General Intelligence even as scientists and other experts cannot agree what the path to get there would look like, and indeed whether it is even achievable.

The complexities of AI and how it is woven into the fabric of our society mean that there is not one single way to make sense of AI, one single set of outcomes from its use, or one single future to look towards.

For some, extraordinary access to and surveillance of our inner lives, root operating systems and most intimate relationships – often to the benefit of a small number of companies – is a fair exchange for a future in which AI is an ever-present companion to whom we can outsource human drudgery and human decision-making. There are others for whom AI is an existential threat, with all current harms overshadowed by potential future peril. Others see AI at best as a clever but ultimately trivial novelty, and at worst as a source of harms and risks which inevitably and disproportionately impacts the most vulnerable and the already disenfranchised.

Taken on their own, each of these narratives will in time prove to be incomplete. The complexities of AI and how it is woven into the fabric of our society mean that there is not one single way to make sense of AI, one single set of outcomes from its use, or one single future to look towards.

So, where are we?

Ada has published this strategy at a time of extraordinary geopolitical change and uncertainty, which is both shaping and being shaped by rapid technological progress and adoption. We are seeing deregulatory headwinds worldwide, with remarkable convergence between the stated interests of some of the world's largest economies and exceptionally powerful technology companies, who hold an apparent shared belief that they are in a race, with the finish line largely unknown or ill-defined.

These upheavals risk undoing a period of slow but meaningful steps towards ensuring data and AI work for public benefit. We have seen new and important attempts to govern new technologies, with the EU introducing the world's first piece of comprehensive legislation to regulate AI and the UK bringing forward laws on online safety and regulating digital markets. The global network of institutes focusing on AI safety and security has grown to include more than a dozen countries, and AI regulation and governance has been enacted or proposed in jurisdictions across the globe.

The last seven years have also seen major shifts in data and AI deployment, capabilities and markets – from the increased use of predictive policing and facial recognition, to the explosion of generative AI and foundation models, to 'ambient' AI being integrated into the platforms and tools we use every day, without much in the way of consensus-building or sometimes even consent.

This pace of change looks set to continue, with the widespread integration of AI into new contexts. Advancements in machine learning have led to more accurate weather forecasting and improvements in speech-to-text and translation tools. AI systems are being used in drug development and health research, as intermediaries to knowledge and content, and are positioned as being on the verge of acting as autonomous 'agents'.

There is potential for unprecedented and likely uneven acceleration in AI capabilities – whether in the next years or the next decades – which could lead to sweeping changes in society. And yet, despite political expectations that increasing investment in and deployment of AI will yield public benefits, there is no clear inclusive and democratic political vision for what technology harnessed in the public interest could mean, and where various publics want to see the benefits accrue. While we hear much about the futures big technology companies want to see, we rarely hear about the hopes, aspirations and fears of diverse publics as AI becomes intertwined with their everyday lives.

And so, where to?

In 2025 we have seen how quickly policy, regulation and industry priorities can change in light of new geopolitical pressures. We are seeing old problems – like regulatory capture, fragile institutions, dependency on big technology companies and a lack of accountability for harms – come into much sharper focus.

Compounding these problems is a weakening of many tools we have traditionally used to guide innovation and maintain public trust. There is a disconnect between the pace at which technology changes and the time it takes to respond via our civic and democratic institutions and governance processes.

While the world continues to change, Ada's core mission has not. As ever, amid the narratives of hype and hope, we want to ask: do these new technologies work? Do they work well in context – from hospitals to schools? Do they work for everyone? And how do the power and value systems embedded in them impact on both the futures we are seeing and the futures we want to see?

These questions have shaped our strategy. It seeks to understand how people are affected by AI and how they want it to be integrated into their lives. And it will explore how these technologies are built into the public services and products we use every day.

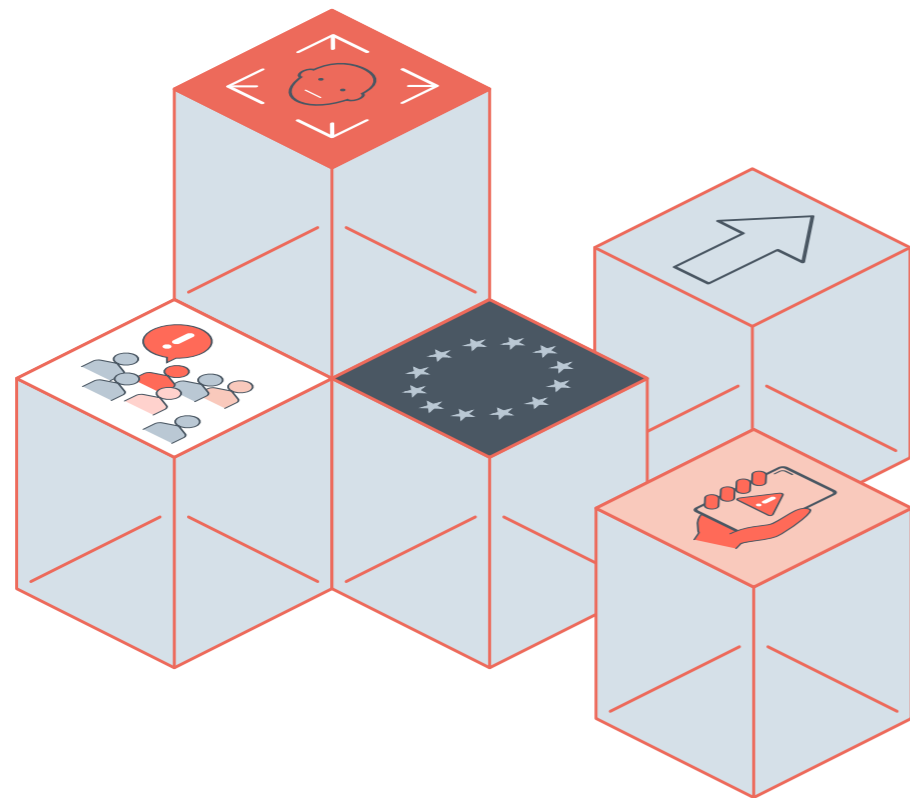
Through our research and convening, we will consider the often messy and complex reality of how technologies interact with real people and real services. We will examine which models of governance will work in the public interest, and identify and challenge power imbalances and inequalities. In a time of changing and weakening institutions, we will seek to rigorously answer all these questions – holding ourselves and others accountable to the evidence – and to place people and society at the centre of every decision about AI.

Gaia Marcus

September 2025

About us

Ada was established by the Nuffield Foundation in 2018, in partnership with the Alan Turing Institute, the Royal Society, the British Academy, the Royal Statistical Society, the Nuffield Council on Bioethics, the Wellcome Trust, techUK and Luminate.



We aim to...

- Ensure data and AI work for people and society, and the opportunities, benefits and privileges generated by data and AI are justly and equitably distributed and experienced.

We do this by...

- **Convening diverse voices** to create an inclusive understanding of the ethical issues arising from data and AI.
- **Building evidence on social impacts** to support rigorous research and foster informed debate on how data and AI affect people and society.
- **Influencing policy and practice** to prioritise societal benefits in the design and deployment of data and AI.

People and society are at the centre of our work and vision. Our research covers a wide range of important issues, including data and AI regulation, how technologies can be developed and deployed to support and protect people, and the use of technologies in public services. We operate independently of government and the tech industry, and we ensure that any claim or conclusion we put forward is based on robust and rigorous evidence.

We have shaped and influenced consequential areas of the AI and data landscape, including the governance of biometrics, the adoption and design of COVID-19 contact tracing and vaccine passports, and the EU AI Act.

Our strategy 2025-28

What we have learned

Over the last seven years, Ada has established itself as a trusted, evidence-based institution.

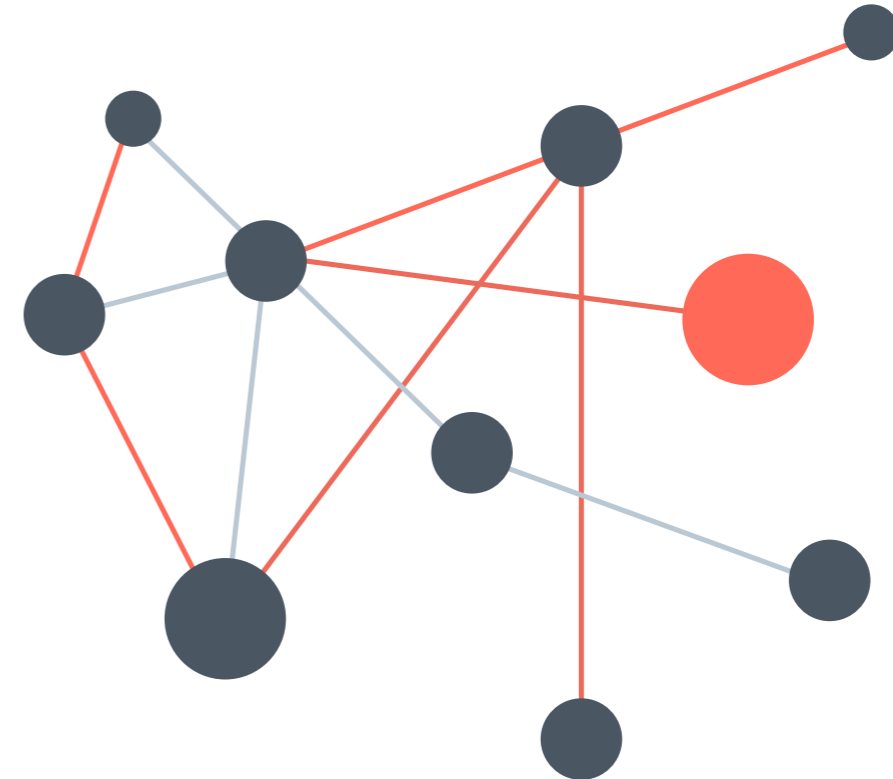
Independently funded and curious, we have sought diverse perspectives on data and AI and successfully brought together evidence and conversations from across civil society, academia, governments, the technology sector and the public. Our measured and balanced voice provides clarity and rigour in a landscape often dominated by hype, hope or fear.

Our work exploring the intersection between AI, data, people and society has surfaced the following lessons:

- **AI and data-driven technologies are ‘sociotechnical’.** This means that technologies do not exist in a vacuum: they influence and are influenced by the social contexts in which they are deployed. This sociotechnical framing guides our approach, by addressing the two-way relationship between technologies and the people who are affected by them.

- **Data is increasingly used, collected and processed in ways that were previously unimaginable**, for example, biometrics capabilities which evolved from verification, identification and categorisation systems to emerging systems on cognitive and biometric inference. This means that regulatory and governance regimes can fall out of step with practice on the ground.
- **In many cases, there is not enough evidence to know whether or in which contexts AI and data-driven systems work as intended.** This is compounded by low levels of transparency and a lack of publicly available evaluations. Society and public services are often over-reliant on industry accounts of innovation and opportunity.
- **People hold diverse and nuanced views about AI that are not always taken into account by governments, companies and other powerful actors.** The tools that have historically emerged to allow people and communities to steer innovation – regulatory checks and balances, civil society and trade union mobilisation, and political leadership from elected representatives – are not always working well enough to protect people and society from the negative impacts of rapid technological change. New participatory mechanisms and deliberative approaches hold promise but require institutional backing and political buy-in to succeed.

- **The use of AI and data-driven technologies in public services is complex and requires care, investment and expertise.** The success and acceptance of technological tools depend on their interaction with existing social systems, values and trust. The cost of proceeding without caution is too great – risking harm to people and communities, and significant financial costs.
- **Governments' focus on growth may lead to short-sighted decisions about the development and deployment of new technologies.** In the UK and beyond, ending economic stagnation and returning to growth has become a priority for policymakers. Many see AI as an economic opportunity, but promises of widespread benefits are often poorly defined and evidenced. Too often, the quest for growth at all costs risks undermining the incentives and mechanisms that are necessary to develop and use AI safely and effectively.



What we are facing

Our work to date and analysis of what we are likely to see emerge in the next few years has highlighted the following as core obstacles to ensuring data and AI work for people and society.

- **A global AI ‘arms race’ narrative:** Major nation-states are seeking national security and economic advantage through AI development, framed as a zero-sum competition between nations for investment and a small pool of talented researchers. This endeavour is often predicated on the assumption that the economic benefits of AI will be transformative, with the benefits primarily accruing to the country or region that ‘achieves’ a certain level of capability or deployment first.

This ‘arms race’ narrative has reshaped foreign policy, competition, economic growth and regulatory debates; it has led to an increase in public and private investment in data centres and energy infrastructure; and it has brought technology companies into a more entangled relationship with governments and the national security state.

- **Extreme concentration of market power:** A small number of powerful companies are disproportionately shaping the technology ecosystem (from design to regulation to research), with profound impacts across public, social, professional and personal spheres. In most cases, this concentration of market power is a continuation of the digital economic trends in past decades, where the same companies control the data, talent, platforms, hardware, market share and vertical integration necessary to rapidly develop and deploy AI systems at scale.
- **Slow governance responses to risks:** Governments around the world have been slow to regulate around the multitude of issues AI raises, leaving no mechanism to protect vulnerable groups, balance power and ensure broadly distributed social and economic benefit. From the bias inherent in facial recognition technologies to the opacity of automated decision-making tools, new vulnerabilities continue to be created, and some groups and people have been harmed.
- **The risk of existing protections being traded off against uncertain economic promises:** Legislative progress has been under threat from priorities around economic growth, competitiveness and trade. We are seeing legislation and protections that were previously agreed being brought up for debate, including debates around online safety, data protection and AI regulation laws. These changes (that would largely benefit major AI developers) are being proposed against a context of uncertain – but heavily marketed – future AI benefits.
- **Inadequate tools and incentives for evaluating the impact of AI:** There is a lack of systematic evidence for the efficacy of many AI systems or proper incentives to ensure this evidence is generated, as well as a lack of mechanisms to anticipate and monitor social, economic or environmental impacts. This has let hype and hope dominate the debate rather than a strong evidence base.



- **Less individual and collective control over consequential AI decisions and data use:** Those affected have little say over how their data is used, or how technologies mediate or disrupt important aspects of their lives and relationships – and there is little consensus or knowledge among policymakers on how to include them. Collective mechanisms like consumer groups or unions have also experienced challenges in helping people have greater control. People lack the mechanisms to seek redress from harms caused by technology and this may come to drive a sense of disempowerment.
- **Lack of transparency across the data and AI lifecycles:** There is little clarity about how information is being used or shaped by technology, and there are major barriers to assessing how and whether AI systems work safely, compounding the difficulty policymakers and regulators have in responding to the impacts of AI.
- **Declining trust in information and institutions:** Social goods like trust in information are being undermined. Twentieth-century institutions, policy instruments and services are also struggling to meet people's needs.

Our work over the next three years will be focused on confronting these obstacles – examining and offering solutions for these problems through building evidence, influencing policy and practice, and convening diverse voices.

Where we are going: objectives

Over the next three years, we will focus our research, practice and policy engagement on the following outcomes to help secure a future where data and AI work for people and society.

Our strategic objectives are explicit about the need to understand how AI and data-driven technologies could support a positive vision for society, and the policy choices and institutions this will require. Taken together they allow us to understand:

- how people and society are affected by data and AI, how they want to be affected, and what positive visions for the future look like

and seek to deliver on that future through:

- evidencing and shaping how AI is adopted into essential services and by the state
- exploring mechanisms to ensure appropriate governance of AI and data systems
- explicitly identifying power imbalances and inequalities that underpin the current development and deployment of AI and data-driven technologies, and amplifying and investigating methods to address these.

Objective 1

Explore how people and society are being affected by AI, and what 'AI in the public interest' could mean.



Ada will study the gap between people's experiences of technologies and what 'AI in the public interest' could mean for different communities. Phrases like 'public benefit' and 'public interest' AI are used across policy, practice and public discourse but lack specificity, legitimacy and a clear vision. In some cases, technology companies have exploited the language of 'public interest' to justify deployments and decisions. Our research will highlight how AI technologies could be adopted, used and deployed with public legitimacy – and where they should not be – interrogating the public's 'red lines' and their priorities.

To ensure decisions about the deployment of these technologies are legitimate and informed by public evidence, Ada will engage with diverse publics, those working in public services, policymakers and academics. We will:

- explore and make explicit what a positive vision for AI means for different use cases and for different communities
- demonstrate where and how public views can be built into technology decisions and processes relating to AI development, adoption and governance, using a range of grounded participatory methods to incorporate people's perspectives and experiences
- set out suggestions and examples of the policies and practices needed to deliver a positive vision, to inform those seeking to serve affected communities.

Objective 2

Evidence and shape the use of AI and data-driven technologies where they most impact people and society, with a focus on the public sector.



AI and data-driven technologies are being deployed across the public sector and their use is often framed around expectations of increased efficiency, lower costs and better outcomes for sectors such as healthcare and education. There is a pressing need for evidence to inform and influence decisions about the adoption of AI in the public sector, particularly to identify conditions which balance the needs of people, communities, frontline professionals and services. Priority areas include sectors where data and AI are making or are likely to make profound changes to frontline practice, such as health and care, education, justice and welfare.

Ada will deeply explore these changes, with a focus on workforce transition, the experience of the most vulnerable, and the changing role of institutions and traditional power dynamics. We will:

- engage with different groups of people and users to critically examine the choices embedded in AI implementation, considering what might be needed for policymakers and practitioners to use data and AI well in specific domains, and where there may be red lines or unsuitable use cases
- identify conditions for beneficial deployment of AI across the public sector using these sectoral case studies, and synthesise evidence on impacts and harms.
- continue to work on transparency and evaluation requirements for AI technologies adopted in the public sector to ensure they work, as well as to identify use cases where the adoption of AI improves the public sector and public services
- advise on, pilot and co-develop methods for evaluation and assessment of the impacts of new technologies in the public sector.

Objective 3

Evaluate and inform incentives for managing AI risks.



Achieving positive social and economic outcomes from AI technologies fundamentally involves managing their risks. However, we are seeing a wave of deregulatory attitudes among national and international policymakers. Our work on AI governance has been driven by a simple principle: those best able to manage risks and harms at each point in the AI value chain should be credibly incentivised and empowered to do so.

Ada will advance the design and implementation of incentives and mechanisms that meaningfully address AI risks and harms, particularly through AI governance. We will:

- continue to describe and evaluate approaches for how to responsibly develop, deploy and govern AI and data-driven technologies
- seek to ensure proposals by policymakers in the UK, EU and internationally are as effective as possible in minimising harm to people and society and provide meaningful routes to redress
- study commercial and economic drivers for governance demand, for example by assessing the incentives of actors in the investment, procurement and assurance ecosystems and their common interest in AI risk management, and build evidence on how jurisdictions already manage risk for other consequential technologies
- support policymakers to understand the social, economic and political costs of leaving AI risks unmanaged, including the impact on public trust in AI
- anticipate the impacts of emerging technologies and help policymakers prepare with respect to both specific technologies (e.g. advanced AI assistants and agentic capabilities) and wider systemic impacts or disruption (e.g. on the labour market).

Objective 4

Examine vehicles for the redistribution of power and opportunity across society.



A small number of big technology companies currently dominate the resources and infrastructure upon which modern AI depends, creating new forms of reliance and vulnerability. Increasingly, the interests of these companies shape the trajectory not only of AI development but of public services and government policy, diminishing our collective ability to imagine alternative technological futures.

Ada will interrogate asymmetries of power over AI and develop credible proposals for how policymakers can support alternative political and economic options and priorities for AI development and use. We will:

- convene international experts from governments, academia and industry, as well as from trade unions, consumer groups and other organisations representing communities affected by AI development and use
- work towards a better understanding of public expectations around sovereignty and democratic choice in a context of increasing technological interdependence and growing public sector reliance on big technology companies
- spotlight the best emerging thinking on the political economy and political theory of AI, and evaluate existing proposals for 'public interest' AI for their impacts on power and equality
- work to develop and advance new, credible proposals for how policymakers can support alternative political economies for AI that disperse power and opportunity more widely.

How we will get there

The ways we work and drive change are vital to the success of our new strategy.

How we work

We have six principles that guide our work:

Independence

We are independent of government and industry. This means we can determine the focus and content of our work, take a long-term view and critically examine complex systems. Our independence allows us to bring together multiple lenses and points of view, without external influence or obligations.

Quality, rigour and credibility

We begin our research from a position of empirical curiosity and critical awareness of power dynamics. Our work is grounded in robust evidence, expert analysis and interdisciplinary commentary. We actively incorporate public perspectives and diverse expert knowledge into our evidence and understanding.

Collaboration, interdisciplinarity and openness

Our endeavour is an inclusive and collaborative one. We believe that interdisciplinary approaches produce better outcomes, and so we seek to build a diverse team and convene partners from across sectors, disciplines and lived experiences. We are transparent about relationships and funding.

Connectivity and diversity

We situate our work at the intersection of technology and society, drawing on a nuanced understanding of national and international developments. We employ comparative approaches, recognise and celebrate difference, engage in international debates and participate

in discussions about global governance. We believe in fostering a diverse community of scholarship around the impacts of AI and data, prioritising public participation and amplifying the voices of groups that are traditionally marginalised.

Timeliness, relevance and impact

Our ambition is that our work will be consequential now and in the future. Our mission is to create positive change and this shapes our areas of focus, who we engage with and the speed at which we produce outputs.

Monitoring, evaluation and learning

We recognise our work exists in an environment of emergent change. We encourage evaluative thinking, open sharing of successes and failures, and thoughtful reflection and learning.

How we make change

As we work towards achieving the objectives in our strategy, we are guided by our theory of impact. We take a relational and realistic approach to impact, acknowledging our role in a complex and committed network of people and organisations. We recognise how our work, alongside the work of others seeking to create positive change, can contribute to significant impact. As well as producing rigorous, evidence-based research, we recognise the need to turn our research into knowledge that can be applied by decision-makers to real-world situations.

We seek to make change and understand our impact in five areas:

Policy and law: How our activities (for example, policy advice, evidence, briefings) can lead to specific changes in legal mechanisms, governance and policy.

Practice: How we can influence practices in regulatory bodies, the public sector or technology companies and how this can lead to behavioural change, for example in procurement or assessment.

Understanding and awareness: How we contribute to a better understanding of the scale or urgency of a problem or raise awareness of an issue that has had low research or media attention.

Attitudes and perceptions: How our work can change attitudes or perceptions, in the way people in e.g. media and politics think and talk about issues at a local and/or societal scale, or lead to new appreciations of underrepresented views.

Capability and preparedness: How we can enable changes in the specific capacities of individuals, groups or communities, or to the development of new institutions or infrastructure.

We have learned we are most impactful when we are:

- **Producing timely and anticipatory work:** spotting emerging trends and informing policy and practice through high-quality, evidence-based responses to the most pressing challenges and questions in data and AI.

- **Convening and amplifying diverse voices:** using our independent position to bring together perspectives from industry, public service, civil society, academia, and the people and communities affected by uses of data and AI.
- **Driving clarity:** providing rigorous research, demystifying and translating technical complexity, offering clarity in contested debates, and making all our work accessible to a wide range of audiences.
- **Surfacing public attitudes:** connecting public perspectives to those designing, deploying and governing technologies, to explore and explain complex and contested questions that do not have easy answers.
- **Diversifying inputs into data and AI governance:** working collaboratively with policymakers, regulators, governments, academics and industry to develop solutions to challenging issues in data and AI governance.
- **Identifying and filling sociotechnical evidence gaps:** bringing diverse expertise and a range of methodological approaches to questions about how and where data and AI should be used, incorporating arguments from social sciences and policy; privacy and rights; and participation and social justice.

We know that technology moves fast, and we have built this strategy with that in mind. As we drive our objectives forward – and strive for a world where data and AI work for people and society – we will ensure we have the flexibility and capacity to meet the moment and respond quickly to new developments.

Find out more



[Adalovelaceinstitute.org](https://adalovelaceinstitute.org)



[@adalovelaceinst.bsky.social](https://adalovelaceinst.bsky.social)



[Ada Lovelace Institute](https://www.linkedin.com/company/adalovelaceinstitute)



hello@adalovelaceinstitute.org



Sign up for our fortnightly [newsletter](#)