HewardMills
Your Global DPO

“Whether or not you’re using AI — AI is using you.”
Professor Renée Cummings

The regulatory landscape and responsible governance of Artificial Intelligence

Key takeaways from a global Data Protection Officer

Certified B Corporation

Executive Summary

On March 21, 2025, HewardMills gathered 30 Artificial Intelligence (AI) practitioners, experts, advocates, enthusiasts and policymakers from the UK, U.S. and global communities at the UK Houses of Parliament, to explore the expanding role of AI, its impacts, and strategies for responsible development. This inaugural AI roundtable was co-hosted with Protect Black Women, a grassroots organisation committed to advancing the rights, dignity, and visibility of Black women across all sectors of society, bringing an insightful perspective to the discussion on AI's intersectional impacts.

The following provides an overview of key areas examined and some of the conclusions drawn:

01 Global AI regulations are diverging, not converging
The US and UK favour a market-driven, sovereign approach, whereas the EU is implementing a risk-based tiered framework under the AI Act.

02 The EU AI Act sets a new global benchmark
With "unacceptable risk" bans and heavy requirements on high-risk systems (e.g., medical, finance), it phases in through 2025–26.

03 The UK's "pro-innovation" model relies on decentralised oversight
Instead of a new regulator, bodies like the ICO, CMA and Ofcom police AI, supported by regulatory sandboxes – though recent deregulatory bills threaten these guardrails.

04 Self-regulation is growing but remains uneven
Companies are standing up ethics committees and governance programmes, yet without harmonised standards or enforcement they risk competitive disadvantage.

05 AI can reinforce structural inequalities
From facial recognition misidentifying non-white faces to biased recruitment algorithms, intersectional harms run deep in data, design and deployment.

06 Diverse participation is non-negotiable
Embedding voices of those most impacted—across race, gender, socio-economic lines—is key to spotting blind spots and mitigating bias.

07 Privacy-by-design and ethical oversight must be baked in
Robust DPIAs, modular compliance, and clear ethical frameworks throughout the AI lifecycle turn governance from an afterthought into a competitive advantage.

08 Adopt a risk-based governance approach
Align with frameworks like NIST and ISO/IEC 42001, ensure traceability and documentation from the outset to avoid costly retrofits.

09 AI literacy and training underpin responsible innovation
Building awareness at all levels—executives, developers and end-users—helps organisations anticipate risks rather than merely react to them.

10 Data Protection Officers are strategic AI partners
DPOs guide global compliance, lead tailored impact assessments, embed privacy controls and help build the trust that drives adoption.

This white paper further explores these key themes and offers a roadmap for navigating and understanding evolving regulatory landscapes, emerging business innovations and associated risks. It also outlines practical measures for mitigating harm, particularly those rooted in bias and inequality, by integrating ethical governance, privacy principles, and inclusive practices into AI development.

Introduction

Artificial Intelligence (AI) is rapidly transforming every facet of society, from healthcare to finance, education to law enforcement. While the promise of AI lies in its power to drive innovation and improve efficiency, its risks—particularly around privacy, bias, and exclusion—are equally profound.

Synthesising insights from the roundtable discussion, here we explore the practical realities of AI adoption in the UK and globally, examining its profound and often disproportionate impact on different communities, and identifying pathways toward responsible AI development and use. We also explore the growing role of privacy professionals and other critical stakeholders as key actors in shaping ethical and accountable innovation.

Contents

Are global AI regulations converging or colliding?

Global AI regulatory initiatives are becoming increasingly divided, as evidenced by the recent Paris AI Action Summit. The U.S. and the UK declined to sign a declaration supporting open, ethical, and sustainable AI, despite 60 nations—including France, China, and India—signing it.

The UK raised worries about national security and unclear global governance, while the U.S. stressed the need of putting innovation ahead of regulation and cautioned against choking the AI sector. The discrepancy is indicative of a larger conflict: the U.S. and UK favour a more market-driven, sovereign approach, while Europe advocates for stricter, cooperative regulation.

The EU's comprehensive AI Act, which entered into force in August 2024, is being implemented in phases through 2025 and 2026. It sets a global benchmark with a risk-based tiered approach to regulating AI systems. It bans "unacceptable risk" AI (e.g. social scoring, exploitative systems) and imposes heavy requirements on high-risk AI (like in medical or finance), reflecting Europe's emphasis on fundamental rights and safety.

Europe's stringent GDPR also applies to AI systems – e.g. using personal data for AI must have legal basis and safeguards. In 2023, Italy even banned ChatGPT temporarily over privacy concerns, forcing OpenAI to improve consent and age checks before reinstatement. This underscores the EU regulators' vigilance on AI privacy.

The U.S., by contrast, has no single AI law; it favours a flexible, sector-specific approach often underpinned by voluntary frameworks. Guidelines exist, with the White House's voluntary AI Bill of Rights outlining fairness, transparency and protection of civil rights, and agencies like the Federal Trade Commission (FTC) enforce privacy and consumer protection in AI, but regulation is patchy. This pro-innovation stance avoids early over-regulation but risks a lack of consistency.

With the lack of a U.S. federal privacy law, AI deployments rely on a patchwork of state laws and company policies. As a result, controversies like social media data being used to train AI have fewer immediate legal repercussions in the U.S., though public pressure is rising.

China pursues AI dominance by 2030 under strong state oversight. Its approach is "vertical" – implementing discrete laws on specific AI issues. For example, rules on recommendation algorithms and deepfakes require algorithmic transparency and labelling of AI-generated content. It also restricts the use of facial recognition by private actors (it must be for justified purposes), even as the government itself makes prolific use of AI in surveillance.

With a newer Personal Information Protection Law, China's regime permits extensive data use by the government and companies, raising different privacy implications, particularly regarding state access to data and the potential for mass surveillance.

By **January 2024**, over **850 national AI policy initiatives** had been reported across **69 countries**, underscoring a patchwork of evolving regulatory approaches worldwide (OECD, 2024)

For multinational organisations, the divergence in legal frameworks creates a complex compliance environment. Instead of moving toward a globally aligned AI governance model, jurisdictions are prioritising their own strategic, ethical, or economic interests, forcing stakeholders to adapt to localised compliance strategies rather than a single global standard. With the rapid expansion of AI, this global outlook highlights the critical need for adaptive governance frameworks that protect data, uphold individual rights, and foster responsible innovation.

BACK TO TOP

Fostering innovation with guardrails: the UK's approach

The inherent tension between harnessing AI's efficiency and mitigating its potential harms stands out as a recurring theme in relation to innovation, with many examples of success and failure to be found. For instance:

- **Efficiency gains**: use cases in financial services (e.g., AI-driven compliance audits reducing workload by 60-70%) and logistics (e.g., Ocado's AI-optimised warehouses) demonstrate significant potential cost and time savings.

- **Risks and biases**: real-world examples, such as flawed AI implementations in the NHS lacking proper consent and data protection impact assessments, and biased grading systems in education disproportionately harming students from disadvantaged backgrounds, underscore the dangers of unchecked algorithmic decision-making and inadequate governance. In the pursuit of progress, issues may also be overlooked or ignored, as with the roll-out of a Police app in London, despite its AI failures in facial recognition in Wales and ongoing concerns.

Government strategy, including the National AI Strategy and recent AI pro-innovation white paper, aims to cement the UK as a global AI leader while managing risks.

The UK is indeed betting on innovation—but the guardrails are light, decentralised, and at risk of political dilution. For stakeholders, this means navigating a governance environment with evolving guidance, patchy enforcement, and an increasing need to internally champion privacy, ethics, and risk mitigation—especially where statutory safeguards fall short.

On paper, the UK positions itself as a champion of a “pro-innovation” regulatory approach, one that favours context specific oversight over sweeping, centralised legislation like the EU’s AI act. Rather than creating a new AI regulator, the UK assigns responsibility to existing sectoral bodies such as the Information Commissioner’s Office (ICO), the Competition and Markets Authority (CMA), and Ofcom on the basis that they are best placed to understand the unique risks AI poses within their respective domains. This decentralised model is designed to support the gradual development of an AI assurance ecosystem, enabling agility, fostering innovation, and encouraging experimentation through mechanisms like regulatory sandboxes.

The UK now hosts **3,000+ AI companies** generating over **£10 billion** in revenue and employing **60,000** people (Department for Science, Innovation & Technology, 2024)

Nevertheless, increasing deregulatory trends threaten this approach. Through the Digital Information and Data Protection Bill, the UK government is attempting to decrease data protection regulations. Potentially reducing the ICO’s independence and imposing a growth-and-innovation responsibility. Furthermore, the Centre for Data Ethics and Innovation’s (CDEI) mission has changed from independent monitoring to partnership-driven advising, which lessens the pressure on business and government to embrace strict ethical norms. Without mandatory replies or regulatory authorisation, these “guardrails” might not be enforceable. The lack of legislative authority in new organisations like the Digital Markets Unit also suggests reluctance to give regulators significant enforcement authority.

80% of UK adults believe that **AI regulation** is required, and **91 %** want laws specifically to combat **AI-generated misinformation** (KPMG, 2025)

BACK TO TOP

Companies are taking responsibility, but is self-regulation enough?

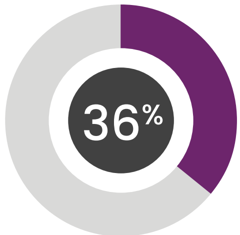

CEOs who express high confidence in their leadership's understanding of AI governance needs (Avanade, 2023)

The global AI landscape is rapidly evolving, with organisations across sectors, from emerging startups to established multinationals, embedding AI technologies into their operations. Alongside these expansions, there is growing investment in responsible AI practices. Companies are increasingly establishing ethics committees, internal governance structures, and compliance programmes to guide the development and deployment of AI. This momentum is contributing to a diverse and maturing ecosystem of voluntary standards aimed at promoting transparency, fairness and accountability.

Government agencies often lack the expertise and resources to keep pace with rapidly evolving AI capabilities, and effective industry self-regulation may reduce pressure for potentially innovation-stifling government intervention.

Recent industry-led initiatives demonstrate this trend[1]. Nevertheless, increasing competitive pressures threaten

For stakeholders, this means preparing for an evolving governance environment with increasing regulatory requirements, while continuing to develop robust internal ethics programmes that anticipate rather than merely react to regulatory developments.

the effectiveness of this approach. With the race to deploy generative AI technologies accelerating, companies face strong market incentives that may conflict with responsible development principles. Organisational strain in managing complex ethical considerations alongside commercial objectives is also suggested by data showing that 55% of privacy professionals now handle AI governance tasks.

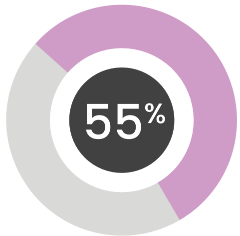

Privacy professionals now handling AI governance tasks (International Association of Privacy Professionals, 2024)

Furthermore, self-regulation lacks consistent enforcement mechanisms or standardised requirements across companies, creating an uneven landscape where responsible actors may face competitive disadvantages against those prioritising speed to market.

The global trend is clearly moving toward a hybrid model, combining corporate responsibility frameworks with proportionate regulation and independent oversight. However, the current trust levels remain concerning, with less than half of employees completely trusting AI-produced results and only 36%[2] of CEOs expressing high confidence in their leadership's understanding of AI governance needs.

1 The U.S Artifical intelligence Safety Institute Consortium (AISC) brings together over 200 firms to promote safe AI deployment through knowledge sharing and testing environments. The trustworthy and Responsible AI Network (TRAIN) combines 16 healthcare leaders with technology organizations to improve AI safety in healthcare settings. The AI enabled information and Communication Technology workforce consortium focuses on upskilling workers in roles likely to be impacted by AI.

2 https://www.avanade.com/en/insights/articles/rise-of-industry-self-regulation

BACK TO TOP

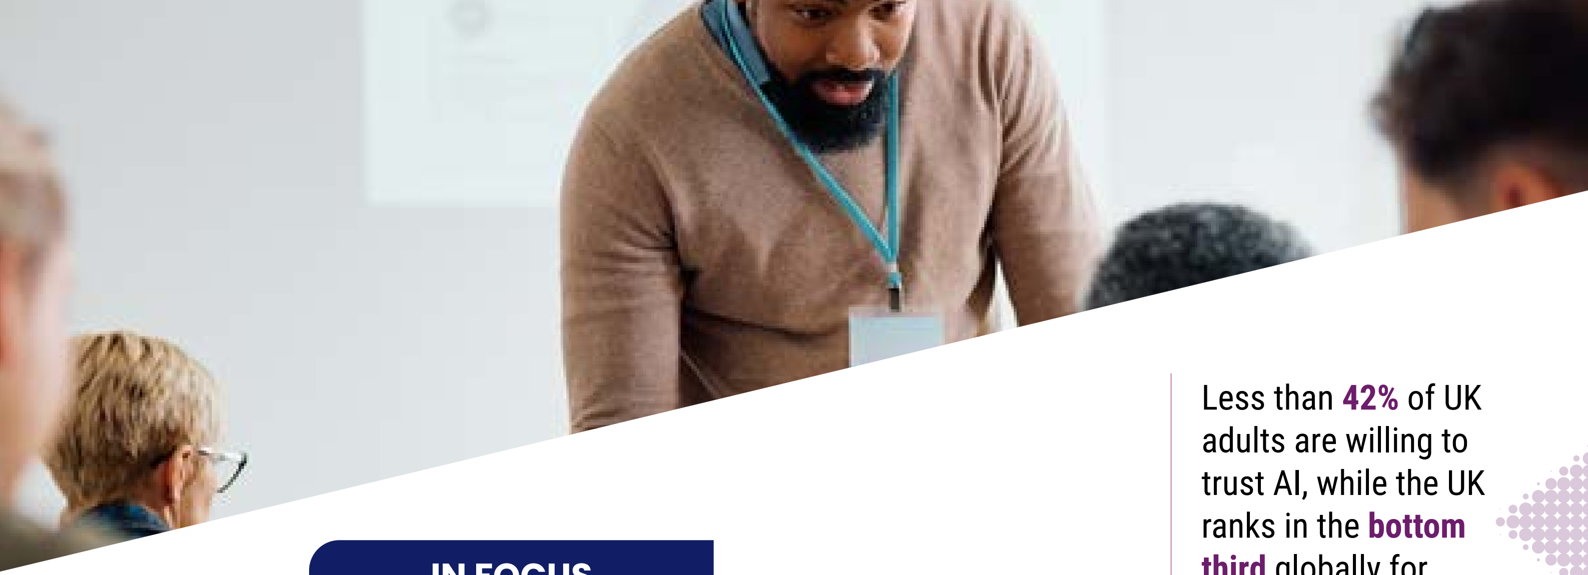

IN FOCUS

AI literacy, training and awareness

Less than **42%** of UK adults are willing to trust AI, while the UK ranks in the **bottom third** globally for formal or informal **AI training** (KPMG International, 2025)

Ensuring stakeholders—from executives to developers and end users—have a solid grasp of AI is essential for responsible innovation and risk mitigation. Experience shows us that a comprehensive and effective approach should:

Start by clarifying core concepts

Provide brief, role-tailored overviews of machine learning, data governance, bias, and ethics. Technical teams learn algorithmic mechanics; legal and privacy teams focus on compliance; business leaders understand strategic risks and opportunities.

Build in role-specific training

Governance frameworks, risk management and regulatory updates should be the focus for Executives, whilst Privacy/DPO teams should be trained in privacy-by-design, impact assessments and audit trails.

For technical staff, model development, bias detection and explainability are key areas. Last but not least, front-line employees should be trained to recognise AI-driven decisions and flag potential issues.

Emphasise AI limitations and ethical awareness

Use concise case studies to illustrate how AI can fail or be fooled (e.g., hallucination, adversarial content in CVs) and how data skew can produce unfair outcomes (e.g., biased hiring algorithms, facial recognition failures), alongside hands-on exercises, such as reviewing sample datasets, to help participants recognise and address bias.

Foster continuous learning

Host periodic "AI forums" or lunch-and-learns where cross-functional teams (technology, legal, privacy, business) share insights on new tools, risks, and changing regulations. Maintain an internal AI resource hub, including e-learning modules, recorded webinars and policy updates, to keep training current.

Raise organisation-wide awareness

Distribute brief newsletters on recent AI incidents, best practices, and regulatory developments. For example the EU AI Act, UK guidance. Offer executive briefings so leadership stays informed about strategic benefits and compliance obligations.

Measure and iterate

Use pre- and post-training assessments to gauge improvements in AI risk awareness. Track referrals to ethics or data protection teams as an indicator of vigilance and gather feedback to refine content, especially in areas like algorithmic explainability or cross-border data flows.

Ultimately, embedding AI literacy and training across the organisation helps stakeholders identify risks early, make informed decisions, and cultivate a culture where responsible AI is a shared priority.

BACK TO TOP

Are AI systems reinforcing structural inequalities?

Intersectional harms are structural in nature rather than incidental or uncommon. In AI, bias is not limited to algorithmic results; it is also ingrained in training data, system architecture, and the makeup of the development teams.

These prejudices are particularly noticeable at the intersections of race, gender, class, and socioeconomic status, with the scope and gravity of intersectional harms clearly demonstrated by real-world examples:

- **Racial bias:** facial recognition software still misidentifies non-white faces, speech recognition systems disproportionately fail Black users in fraud detection, and criminal justice historical datasets continue to exhibit biased trends

- **Gender bias:** the majority of deepfake abuse targets women, with Black women receiving even less institutional protection. Additionally, gender prejudice is still present in recruitment algorithms that are educated on exclusionary hiring histories

- **Socio-economic bias**: poorer populations are disadvantaged in areas like education and credit scoring due to socioeconomic prejudice ingrained in postcode-based algorithmic judgments

- **Exploited labour**: much of AI's backend operations depend on exploited labour, with data labelling frequently outsourced to low-paid workers in vulnerable situations who face exposure to traumatic content with inadequate support

- **Funding bias**: significant hurdles are faced by diverse founders, particularly Black and Brown women, in securing investment for AI ventures, limiting the potential for more inclusive technologies to emerge

- **Intellectual property & erasure**: concerns about AI models being trained on copyrighted works without permission, disproportionately impacting creators, artists and authors whose contributions risk erasure

More than just technical solutions are required to address these problems; a governance model based on accountability, openness, and fairness is needed, together with inclusive design principles and privacy frameworks that acknowledge and address intersectional risk.

Exploring these issues at length during HewardMills' AI roundtable, several panellists stressed the urgent need for regulatory and governance frameworks to:

1. **Recognise historical data biases**: acknowledge that datasets often reflect and perpetuate long-standing societal inequalities

Addressing the deeply embedded nature of AI bias requires a concerted and multi-faceted effort from all impacted stakeholders. Critically, Policymakers and regulators have a key role in establishing and enforcing robust frameworks that mandate transparency, accountability, and fairness in AI systems. These regulations should explicitly address the potential for intersectional harms and ensure effective mechanisms for redress are available to those adversely affected. Concurrently, organisations procuring and deploying AI tools, must exercise thorough due diligence. This includes conducting comprehensive impact assessments, focusing on bias and fairness, and establishing clear ethical guidelines for AI's use within their operations to protect users and mitigate risks.

2. **Ensure transparency, fairness, and accountability**: implement mechanisms to scrutinise AI decision-making and provide recourse

3. **Facilitate diverse participation**: actively include the voices and perspectives of those most likely to be negatively affected by AI in its design, development, and oversight

Organisations may go beyond compliance and create AI systems that benefit and do not harm the most vulnerable by including these principles from the beginning.

A recent University of Washington study found that large language models favoured **white-associated names 85%** of the time in résumé screenings and **never preferred typically Black male names** (Wilson, K., & Caliskan, A., 2024)

BACK TO TOP

How can businesses build responsible and compliant AI systems?

Mitigating AI risk starts with strong governance and data protection and privacy controls. Embedding ethical oversight throughout the development lifecycle through dedicated committees, robust privacy programmes, accountable leadership, and alignment with frameworks, like the OECD AI Principles and ISO/IEC 42001, lays the groundwork for responsible and compliant AI systems.

The fragmented regulatory landscape requires organisations to monitor legal developments across all jurisdictions of operation and build flexibility into their AI systems. By developing modular components that can be updated as requirements evolve and creating non-AI fallback processes, businesses can maintain continuity during compliance updates. This adaptability is crucial in an environment where regulations are rapidly evolving at different rates across regions.

A risk-based governance approach aligns with emerging regulatory frameworks like the EU AI Act, which imposes stricter requirements for high-risk applications. By implementing the National Institute of Standards and Technology (NIST) framework for technical risk management

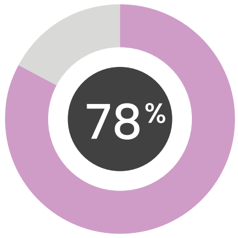

Global companies that use AI in at least one business function (McKinsey & Company, 2025)

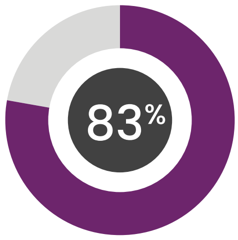

Global companies that say they still want clearer guidance on AI-related regulations before investing (KPMG International, 2025) and ensuring traceability and documentation from the outset, businesses can avoid costly retrofitting later. Building ethics and compliance checks into each development stage also creates multiple opportunities to identify and address potential problems before they manifest in deployed systems.

Transparency and stakeholder engagement strengthen AI systems by incorporating diverse perspectives and building trust. Involving potentially impacted groups, providing clear explanations about system functionality and limitations, establishing contestability mechanisms, and participating in industry consortiums help businesses develop more robust and trusted AI applications. These engagement practices often reveal blind spots that purely technical approaches might miss.

Privacy offers a foundational blueprint for responsible AI governance. Fairness, openness, and risk management are just a few of the AI-specific responsibilities that closely resemble the established data protection guidelines in international privacy regulations like the GDPR. These frameworks can be used by organisations with established privacy programmes to expedite AI governance initiatives. Principle-based, risk-driven, and quickly formalising, the development of AI regulation closely resembles that of the early days of data protection. For privacy experts, who are already essential players in directing morally and legally sound AI development, integrating privacy-by-design, carrying out thorough impact analyses, and making sure accountability procedures are in place are already fundamental practices.

By positioning responsible AI as a value proposition, communicating compliance efforts, integrating compliance into innovation processes, and staying ahead of regulatory trends, businesses can transform what might be seen as regulatory burden into competitive advantage. Ultimately, responsible AI development is about building sustainable systems that create lasting value through earned trust from users, regulators, and society at large.

Ensuring data used by AI systems is handled securely and with appropriate consent also remains paramount, with the increasing use of 'administrative data' gathered implicitly requiring careful consideration of consent and privacy implications.

Ultimately, the most efficient method to guarantee responsible, reliable, and future-proof AI systems is to take a unified approach that views privacy compliance as the operational norm for AI governance.

BACK TO TOP

IN FOCUS

The strategic role of data protection experts

Effectively navigating the complexities of AI requires a deep understanding of privacy regulations, ethical frameworks, and risk mitigation strategies.

Data protection professionals can and should play a critical role, being strategically positioned to:

1. **Ensure legal compliance:** guide organisations through the patchwork of global AI and data protection regulations

2. **Embed privacy-by-design:** integrate data protection principles into the earliest stages of AI system design and development

3. **Conduct robust impact assessments:** Lead Data Protection Impact Assessments (DPIAs) specifically tailored for AI, identifying and mitigating risks related to bias, fairness, transparency, and security

4. **Advise on ethical frameworks**: help organisations develop and implement clear ethical guidelines for AI use

5. **Build trust**: by ensuring compliance and ethical handling of data, Data Protection Officers help build the trust necessary for user and public acceptance of AI technologies

83% of businesses surveyed indicated they want more information about existing or upcoming AI/data regulations, highlighting the **critical need for data protection expertise** (WHITE & CASE, 2025)

BACK TO TOP

A time for action

As AI adoption accelerates, organisations can no longer afford to treat governance as an afterthought. Whether you are building, buying, or simply integrating AI into business processes, now is the time to act.

The following roadmap outlines the practical steps organisations can take to move from reactive compliance toward proactive, responsible AI adoption. Central to this effort is the Data Protection Officer (DPO), whose expertise is critical in designing, developing, and implementing key elements of the framework—ensuring that privacy and data governance are built in from the outset. Structured in phases from **Foundational Actions** to **Strategic Implementation**, the framework helps businesses navigate legal uncertainty, embed ethical practices, and build trust with stakeholders. By following this path, organisations can meet emerging regulatory expectations and gain a competitive edge through transparent, trustworthy AI systems.

Our AI adoption roadmap

This roadmap outlines the practical steps organisations can take to move from reactive compliance toward proactive, responsible AI adoption.

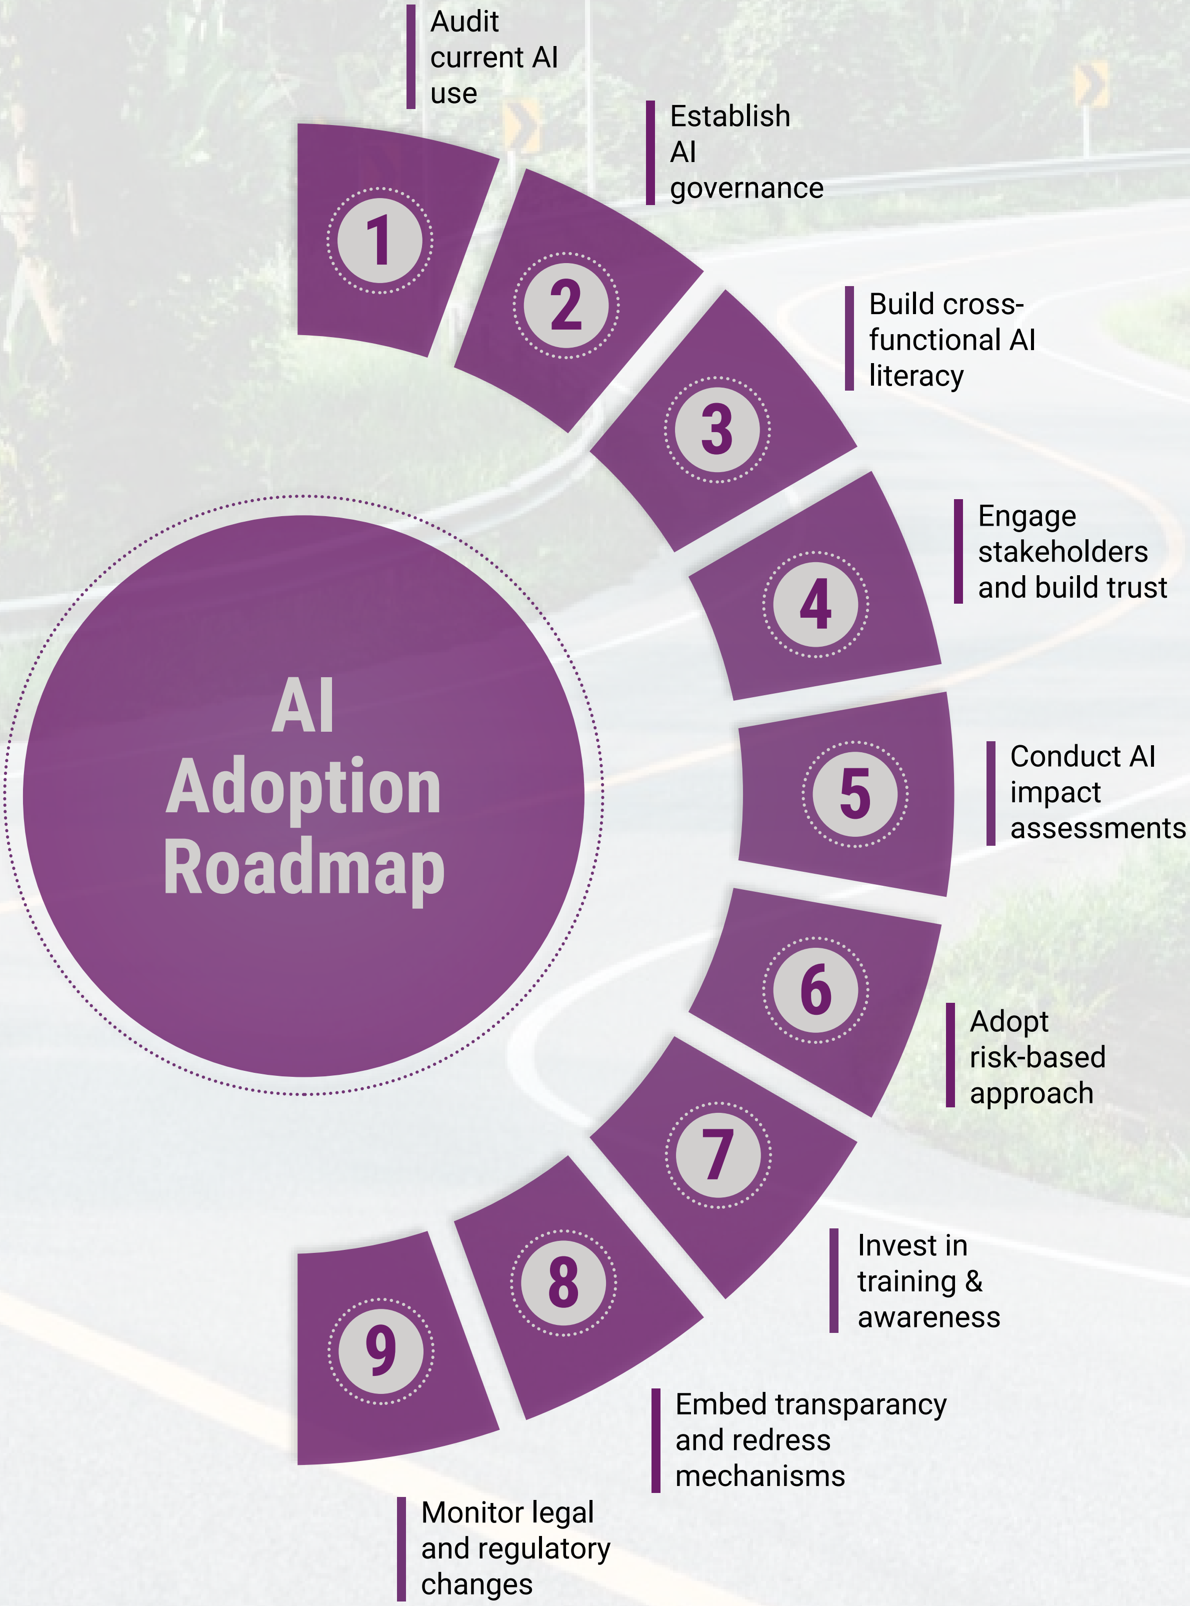

Furthermore, leaders, organisations and communities are recommended to:

1. Integrate intersectional analysis into all AI risk assessments and mitigation strategies
2. Scrutinise data sources and implement processes to detect and address bias, acknowledging data's historical context
3. Promote comprehensive AI and data literacy programmes within organisations/communities and support wider societal initiatives
4. Actively recruit diverse talent for AI roles and support initiatives aimed at broadening participation in the field
5. Strive for transparency in AI systems, explaining decisions where possible and providing users with insight and control
6. Stay informed about international regulatory developments and ensure compliance

BACK TO TOP

Next steps

Building on the success of the inaugural AI roundtable event—which ultimately called for a fundamental rethinking of how AI is designed, deployed and governed, and by whom—further events are planned by HewardMills and Protect Black Women.

By facilitating dialogue and fostering bold, values-driven, and cross-sector collaboration, HewardMills aims to create actionable pathways and actively contribute towards shaping a future where technologies like AI are developed and deployed responsibly, equitably, and for the benefit of all.

Join the conversation by registering your interest in future AI events here >

For more information on how HewardMills can support your organisation, visit hewardmills.com >

BACK TO TOP

About HewardMills

HewardMills is a B Corp certified, global provider of data protection and privacy services, committed to delivering not just regulatory compliance but data dignity. With a deep understanding of the GDPR and other data protection regulations, we partner with responsible businesses to prioritise privacy, build trust, and ensure that data is handled ethically. We offer a comprehensive range of services, including outsourced Data Protection Officer (DPO) solutions and flexible data protection and information security consultancy, designed to help businesses achieve compliance and protect their most valuable assets.

Diverse by design, HewardMills champions the inclusion highlighted as critical throughout the roundtable. We believe harnessing diverse perspectives is essential for developing equitable solutions, both within our organisation and for our clients grappling with AI's societal implications.

Delivering data dignity. Championing diversity.

Certified B Corporation

BACK TO TOP