

SEPTEMBER 2025
JAKOB MÖKANDER
JO PUDDICK
NAMAN GOEL
ALAN WAGER
TIM RHYDDERCH
HELEN MARGETTS
DANIEL CAMERON
BEN ROFF



What the UK Thinks About AI: Building Public Trust to Accelerate Adoption

Contents

- 3 Executive Summary
- 8 Why Healthy Public Attitudes Towards AI Matter
- 13 What the Data Tells Us
- 42 Recommendations to Build Public Trust in AI
- 56 Conclusion
- 57 Guest Authors
- 58 Acknowledgements
- 59 Appendix

Contributors: Laura Gilbert, Kevin Zandermann

Executive Summary

Artificial intelligence is at the heart of the UK government's plans to improve public services and boost economic growth. In January 2025, Prime Minister Keir Starmer announced the AI Opportunities Action Plan and outlined a bold vision for making the UK a global leader on AI.¹

The government is right to emphasise the upside of accelerated AI research and adoption. If carefully designed and deployed, AI systems can help improve efficiency, solve complex challenges from health care to climate change, and protect national security interests.² As AI capabilities continue to advance,³ getting AI policy right matters more than ever.

However, there remains a gap between the government's opportunities-oriented agenda and public attitudes towards AI. UK adults are more inclined to view AI as a risk for the economy (39 per cent) than an opportunity (20 per cent), according to a new poll conducted by the Tony Blair Institute for Global Change (TBI) and Ipsos. Moreover, 38 per cent of respondents in the same survey cite lack of trust in AI as a barrier to adoption.

This is a serious problem. Public attitudes shape not only how AI is adopted and governed but also the legitimacy of its use. Just like hype can be dangerous, so low trust in AI can lead to significant opportunity cost by slowing the rollout of well-functioning, socially beneficial use cases.⁴

Without broad support, the government will struggle to implement the AI Opportunities Action Plan and deliver on its wider growth agenda.

Understanding and improving public attitudes towards AI is thus an urgent task – as is building AI systems worthy of the public's trust.

This paper examines public attitudes towards AI in the UK, drawing primarily on a new survey of 3,727 UK adults conducted between 30 May and 4 June 2025. The research uses the Ipsos Knowledge Panel, which ensures high-quality sampling and digital inclusion, allowing for detailed analysis across

regions, demographics and sectors. We present findings from this new data set, alongside insights from prior research, to explore a simple but pressing question: how can the UK government improve public attitudes towards AI and build justified trust to accelerate the adoption of use cases that improve social outcomes?

Key Takeaways

The findings from the TBI/Ipsos survey on public attitudes towards AI are presented and analysed in detail in the second chapter ("What the Data Tells Us"). However, the key takeaways can be summarised as follows:

- **AI adoption:** More than half of UK adults say that they have used generative AI tools in the past 12 months. This is an encouraging figure for a technology that is only a few years old. As previous research has noted, the adoption of generative-AI tools appears to be faster than for the internet (World Wide Web) and personal computers.⁵
- **Usage patterns:** Just under a quarter (23 per cent) of UK adults use generative-AI tools weekly in their work. A similar percentage report knowing at least a fair amount about AI. However, nearly half of the respondents say they have never used AI either at work or at home. This suggests a large part of the population is either unaware of AI tools or not actively using them for other reasons.
- **Risk versus opportunity:** Across the UK, people's perceptions of whether AI is more of a risk or an opportunity to them personally are finely balanced. For example, 29 per cent of 45- to 54-year-olds viewed AI mainly as a risk and 27 per cent mainly as an opportunity. Younger people are much more likely to view AI as an opportunity, but they also express explicit concerns about broader ethical and societal risks.
- **Skills and confidence:** There is a clear correlation between people's confidence in using AI and their attitudes towards the technology. Of people confident in their AI skills, 66 per cent expect AI to help with parts of their job while leaving their core responsibilities intact. In contrast, only 45 per cent of people with lower confidence in their skills view AI as a supportive tool..

- **Barriers to adoption:** 38 per cent of respondents cite a lack of trust in AI content as a barrier, making it the biggest single obstacle to AI adoption. Concerns around data privacy and ethical standards are other leading barriers across all age groups, whereas disinterest in using the technology is a distinct issue among those aged 65 or older.
- **Trust and use:** Frequent AI users are more likely to view AI positively. Only 26 per cent of weekly AI users view AI as a societal risk, compared to 56 per cent of non-users. It is a self-reinforcing spiral: people who trust AI more use AI more, and vice versa.
- **Demographic variations:** Public attitudes towards AI vary across professional and demographic groups. Confidence is higher among workers in sectors like tech and professional services, but lower in health care and education. Women are more likely than men to view AI as a risk and prefer a more cautious approach to adopting AI. Ethnic minorities are less likely to cite disinterest in using generative-AI tools but are also more likely to cite affordability barriers.

Despite use and familiarity with AI being far from universal, public attitudes towards AI are starting to take shape in the UK. Across a range of domains – including security, the economy and society – UK adults are viewing AI as both a risk and an opportunity. Even individuals who report positive personal experiences with AI remain concerned about its societal implications.

Yet the picture is not all bleak for those who want the UK to seize the opportunities of AI. Surveys conducted by the Department for Science, Innovation & Technology (DSIT),⁶ The Alan Turing Institute and the Ada Lovelace Institute⁷ have found that people feel more positively about well-defined AI use cases with tangible benefits. For example, people appreciate the use of AI to assess cancer risk. Further, comfort with AI is linked to the presence of safeguards – for instance those related to data security, explanations for AI-driven decisions and human oversight. This suggests that there is much political leaders can do to bring the public on board and improve attitudes towards AI.

Policy Implications and Recommendations

Labour's growth agenda hinges on people trusting AI-enabled services and adopting AI tools to boost productivity. Understanding and addressing public attitudes towards AI will therefore be a key success factor. However, the solution is neither to ignore public opinion nor to pause the AI opportunities agenda until attitudes shift. The right path forward is active public engagement combined with policies that drive rapid, sustainable adoption by building justified trust in AI.

This cannot be achieved simply through marketing or persuasion. Instead, two core principles should guide this effort. First, the public must feel confidence in the intentions guiding the deployment of AI. Second, the technology must be trustworthy and reliable, and function as expected both by different users and by those impacted by the technology.

To build justified trust in AI and accelerate adoption of technologies that improve social outcomes, the UK government should:

1. **Put AI in context by focusing on use cases that matter to people.** AI must be developed as a tool to tackle real-world problems and key public priorities. People are more likely to engage with, use and support AI when they can see a tangible, positive impact on their lives. Political leaders' messaging should focus on specific use cases that people can directly relate to (such as quicker scheduling of medical appointments, improved access to services and shorter commute times due to real-time traffic data). This messaging from leaders will also help set the right development goals for AI.
2. **Evaluate AI systems for public benefit in real-world settings,** using metrics that reflect user experience, not just technical performance. When AI systems are deployed in public services, their benefits should be demonstrated credibly and transparently through a trial-and-evaluate approach. Systems that cause harm or fail to offer benefits should be improved based on the feedback from evaluations and users. Of course, human decision-makers and bureaucratic systems also have limitations. The pragmatic test is therefore not whether an AI system is entirely "free from error", but whether it improves on an imperfect status quo.⁸

3. **Ensure continued responsible AI governance and oversight** that meet the public's expectations around reliability, accountability, data privacy and transparency. The UK's innovation-friendly, sector-specific approach is sound. However, as AI systems evolve, so too must the regulations that govern their design and use. In addition to its ambition to address the safety concerns posed by frontier-AI models, the government should plug gaps in existing legislation, strengthen the capacity of sector-specific regulators and lay the foundations for an assurance ecosystem that covers the entire AI value chain.
4. **Close skills gaps through inclusive training programmes** that equip people from all backgrounds to benefit from AI. The adoption of new tools must be matched by workers' ability to use them. Aimed at building practical AI skills, training programmes should be tailored to different sectors, with a focus on basic AI literacy, real-world applications, accessibility and awareness of risks. This training should be co-developed with employers, unions and educators. Strengthening public understanding will help ensure that AI enhances, rather than disrupts, people's lives and livelihoods.
5. **Initiate a series of public-engagement initiatives** that provide people with accessible opportunities to understand what AI is, how it works and why it matters. These could range from a nationwide AI Open House programme, whereby the public are invited to visit institutions using AI and engage with practitioners, to a publicly broadcasted AI lecture series. Another idea is to establish a National AI Discovery and Participation Centre that runs interactive exhibits and virtual experiences to demystify AI, and to act as one of the platforms through which the public can participate in its development.

The recommendations discussed in this paper are necessary but will not be enough on their own to ensure that AI adoption is efficient, inclusive, legitimate and sustainable. Ultimately, harnessing the benefits of AI will require infrastructure investments and incentives for innovation, as well as political leadership and good, proportionate governance. It is in this spirit that building public trust in AI today will help ensure that the UK's future is one that works for everyone.

01

Why Healthy Public Attitudes Towards AI Matter

This paper uses the latest polling data to understand public attitudes towards AI in the UK and make recommendations for building trust to accelerate beneficial AI adoption. Before diving into the data, this chapter explains what healthy public attitudes to AI look like and why they matter.

When we describe the need for improved public attitudes to AI, this is not about people being persuaded that AI should simply be embraced without question. Instead, healthy public attitudes mean that people:

- Understand the opportunities AI offers and are equipped to harness the benefits
- Are aware of realistic AI risks and limitations – as well as common mitigation strategies
- Feel empowered to influence how AI is developed and deployed

A healthy public attitude towards AI is defined by confident, informed engagement – rather than uncritical acceptance, automatic rejection or passive disengagement. Improving public attitudes means not only replacing fear and apathy with understanding and agency but also ensuring that public concerns are heard and addressed.

Two Pillars of AI Trust

AI is no longer a distant promise – it is an increasingly present reality. From language assistants on our phones to advanced diagnostic tools in hospitals and automation in public services, AI is reshaping how we live, work and govern.⁹

Yet for many people, AI still feels unfamiliar, unsettling or out of reach. Predictions of rogue robots and mass unemployment circulate alongside news stories of advancing capabilities, huge data centres and productivity boom. The result is a public mood that is at best mixed and at worst dominated by fear, apathy or confusion.¹⁰

Addressing these mixed attitudes requires building public trust in AI, which rests on two mutually reinforcing pillars. First, confidence in the intentions guiding its design and deployment; second, confidence in its competent implementation and reliability.¹¹

The first pillar of public trust in AI speaks to purpose, values and accountability. To begin, some scepticism towards AI stems from a general lack of trust in democratic institutions and concerns that the benefits from efficiency gains will not be widely shared. Moreover, people may not know what AI is achieving and why, how it affects them, or what recourse they have if things go wrong. While there is no silver bullet to address these concerns, public participation in the selection and design of AI use cases helps, as does transparency around how data is used.

The second pillar speaks to quality, safety and governance, i.e. assurance that AI systems will work reliably, protect personal data and remain under human oversight. Put differently, even when intentions are well-meaning, technical failures like inaccurate and biased outputs can quickly undermine public trust in AI. Preventing such failures requires a combination of robust engineering practices, voluntary testing and certification schemes, and proportionate regulation.

When either pillar is weak, public trust in AI erodes. In both cases, however, scepticism towards AI is not simply about the technology itself but about whether it is embedded in systems that people feel represent them, protect them and work as intended in their interest.

The Role of Public Attitudes in Shaping AI's Future

Public attitudes not only reflect AI's trajectory but are one of its key *drivers*.

ADOPTION AND USE

The way society feels about AI influences whether, how and where it is adopted. Moreover, public attitudes also shape investment decisions and how AI is developed in the first place. Finally, a lack of awareness and trust can prevent the public from adopting and benefiting from even safe and useful AI use cases.

Consider the history of genetically modified organisms (GMOs) as an example. Due to negative public perceptions, the market for GMOs remains small in many countries, despite the role these technologies can play in tackling food scarcity and malnutrition.^{12,13} At the same time, overly optimistic narratives and misplaced trust can also lead to harms and backlash. For example, history has witnessed several “AI winters”, where overblown expectations led to cuts in research funding, and AI failures that have caused real-world harm.^{14,15}

TRUST IN PUBLIC SERVICES

Where AI is integrated into government services such as education and health, public trust is essential. If AI systems are unreliable, intrusive, inscrutable or unfair (or perceived to be so), people may resist their use or disengage from vital services.

The scandal surrounding SyRI – an automated decision-making system deployed by the Dutch government to detect possible benefit fraud – provides a case in point. Due to poor testing and oversight, SyRI wrongfully accused around 26,000 people of making fraudulent benefit claims. SyRI was ultimately shut down, with the court ruling stipulating that it had it unlawfully combined data from different sources.^{16,17} The court warned of a “chilling effect”: without confidence in due process and adequate protections, the public will be less willing to share data, and to trust AI and other types of automated decision-making systems.

AI GOVERNANCE

Public attitudes also affect how AI is governed. Well-informed public attitudes towards AI are crucial, as misinformed attitudes can lead to either unnecessarily restrictive laws or dangerously lax oversight that will eventually cause backlash.

For example, after the March 2011 Fukushima nuclear disaster and subsequent anti-nuclear protests, the German government announced it would be closing all of its nuclear power plants by 2022.¹⁸ The result was increased reliance on fossil fuels, higher electricity prices and slower progress towards climate goals.^{19,20} Preventing potential “Fukushima moments” for AI should be a top priority for anyone interested in maximising its long-term benefits.

Not One AI, Not One Public

To effectively address public attitudes toward AI, it is essential to first recognise that “AI” is an umbrella term referring to a variety of computing systems.²¹ While generative-AI tools like ChatGPT have dominated headlines in recent years, AI systems based on logical reasoning and machine learning have supported public-sector use cases for decades.

Further, public attitudes towards a technology are shaped in part by people’s experience interacting with it. Simplified, people encounter AI in at least two different ways:

1. **Active use (AI as a chosen tool):** In this mode, individuals actively choose to use different AI tools to enhance their personal or professional lives.²² Adoption is voluntary, and people typically maintain a sense of control over when, how and why they engage with these tools. Examples include generative-AI tools like ChatGPT or Copilot, personal fitness and health-tracking apps, language-translation or learning tools, smart assistants like Alexa or self-driving cars. Public attitudes in this context may be linked to factors like awareness of available tools, adoption and experimentation rates, perceived usefulness, ease of use and safety, as well as comfort with privacy trade-offs and data sharing.

- 2. Passive exposure (AI as a system around you):** Here, individuals do not engage with AI directly but are affected by AI-driven decisions indirectly. While this kind of exposure can have significant consequences for people's lives, it comes with limited visibility or control. Examples span AI systems used in government benefit determination and health-care prioritisation, workplace hiring algorithms and monitoring systems, financial-services applications and predictive policing. Beyond the utility of the specific AI use cases, public attitudes in this context are shaped by institutional trust, vendor credibility, perceptions of fairness and transparency, and individuals' sense of agency within AI-influenced systems.

Finally, we must recognise the diversity of human experiences and circumstances that shape public attitudes towards AI. Relevant dimensions include but are not limited to age, gender, race and socioeconomic status. Context and expectations also matter. A person's opinions about AI vary across different use cases and evolve over time. Understanding public attitudes and designing effective interventions mean engaging with this complexity head-on.

02

What the Data Tells Us

AI sparks intense debates. Proponents argue that AI will usher in a new era of growth, rapid scientific advances and human flourishing. Meanwhile, sceptics point out the limitations of AI systems and warn that their premature deployment could end up doing more harm than good. Yet this debate between initiated experts often remains abstract. What does the UK public think?

This chapter presents the findings from the TBI/Ipsos survey on public attitudes towards AI and discusses their implications. The aim of the survey is to get around the noise. How much do UK adults feel they know about AI? How often are they using generative-AI tools? What adoption barriers are people facing? And how do attitudes vary across different demographic groups? The answers to these and other similar empirical questions are key to designing effective AI policies.

The data is primarily based on a new survey of 3,727 UK adults conducted between 30 May and 4 June 2025, using the Ipsos Knowledge Panel.²³ A discussion of survey methodology and limitations can be found in the Appendix. When helpful to shed light on specific dynamics, we highlight complementary insights from past research on public attitudes towards AI in the UK. Also included is a regression analysis conducted by TBI to isolate the role of AI usage and trust as a driver of optimism in the technology from other demographic factors.

Who Is Using AI – and Why Does It Matter?

Just under a quarter of UK workers use generative AI at least once a week in their work. But nearly half have not used it all in the past 12 months.


Just under a quarter (23 per cent) of UK workers report using generative-AI tools at least once a week in their work. A slightly lower proportion of adults say they use AI regularly in their personal lives (19 per cent). These numbers

indicate rapid adoption of generative-AI tools, given that OpenAI only released the first version of ChatGPT in November 2022, less than three years ago.²⁴

That said, nearly half the public says that they have not used any generative-AI tools in the past 12 months, whether at home or in the workplace. While generative-AI tools may not be equally useful to all members of society, the fact is that adoption remains highly uneven.

FIGURE 1

Nearly half of UK adults never use generative AI – whether at home or work



Q: AI tools like ChatGPT, DALL-E, Midjourney and others are now used to generate various types of content, including text, images, videos and films. In the past 12 months, how often, if at all, have you used such generative AI tools in your work / personal life?

Source: Ipsos

Note: Base for personal life: all UK adults 16+ (n 3,727); base for work: all UK adults in work (n 1,868). Fieldwork: 30 May to 4 June 2025.

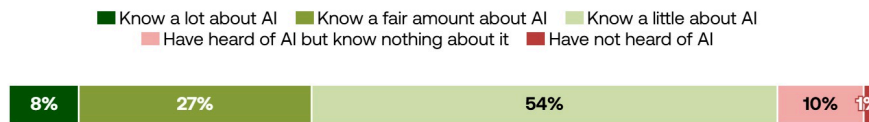
More than a third of UK adults say that they know at least a fair amount about AI.

Knowledge levels follow a similar pattern. Of those surveyed, 34 per cent say that they either know a lot about AI or know a fair amount about AI. On the other hand, 11 per cent say that they have either not heard of AI or have heard of AI but know nothing about it.

This leaves the majority – 54 per cent – somewhere in the middle, knowing only a little about AI. The speed at which these people learn more about AI and how to use it will shape the prospects for productivity gains across society, as well as public attitudes towards AI in public services.

FIGURE 2

The majority of UK adults have some knowledge of AI in 2025 – but only 8 per cent report knowing a lot



Q: Which of the following best describes your knowledge of AI?

Source: Ipsos

Note: Base: all UK adults aged 16+ (n=3,727), unless otherwise specified. Fieldwork: 30 May to 4 June 2025.

The knowledge distribution from our survey mirrors past findings. In 2024, DSIT found that despite the growing popularity of AI systems, there remains a significant knowledge gap about how they work, with 70 per cent of the public reporting they know either nothing or only a little about them.²⁵

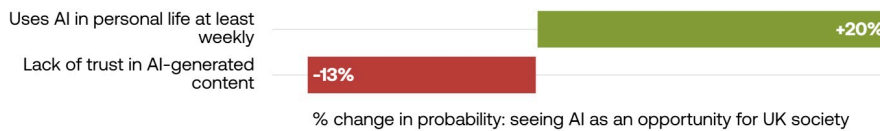
These baseline usage and knowledge levels matter because familiarity and public trust in AI's wider deployment are closely linked. Put simply: the more familiar people are with AI in personal or work contexts, the more likely they are to broadly view AI as an opportunity.

People who use AI regularly are more likely to view AI as an opportunity for society at large.

The statistical evidence for this relationship is compelling. Figure 3 presents the results of a linear-regression model that estimates the change in likelihood that someone views AI more as an opportunity than a risk, compared to the average member of the UK population. Importantly, this statistical method controls for a broad range of demographic factors.

FIGURE 3

People who use AI at least once a week are, controlling for other factors, 20 per cent more likely to see AI as an opportunity than those who use it less often



Q: On balance, do you see AI as more of an opportunity or a risk for UK society?

Source: TBI analysis

Note: Results are from a linear-regression model controlling for demographic factors (age, gender, education, income, occupation and ethnicity). Base: all UK adults 16+ (n 3,727). Fieldwork: 30 May to 4 June 2025.

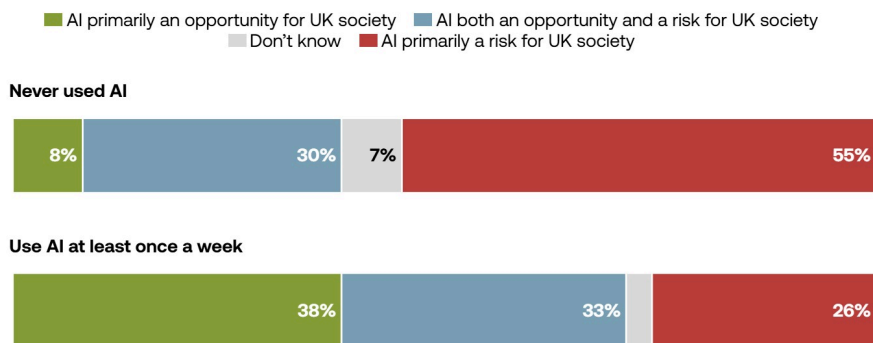
The conclusion is clear: experience with AI is a strong predictor of optimism about its wider societal effects, controlling for factors such as respondents' age, gender, income, whether they are in a white- or blue-collar job and levels of trust in AI. Those who use AI weekly are significantly – 20 per cent – more likely to see AI as a societal opportunity.

This relationship is potentially circular. Those who view AI positively may be more inclined to use it in the first place, but also, those who use the technology regularly and see its benefits may be more likely to develop positive attitudes. Regardless of causality, the relationship has important implications for how AI adoption might create its own momentum.

The data brings this experience-trust relationship into sharp relief: 56 per cent of those who never used generative AI in the past 12 months see it as a risk to society, compared to just 26 per cent of those who already use AI themselves at least once a week in their work. This represents more than a doubling in risk perception between non-users and regular users.

FIGURE 4

More than half of those who have never used AI see it primarily as a social risk, compared to a quarter who use the technology regularly



Q: AI tools like ChatGPT, DALL-E, Midjourney and others are now used to generate various types of content, including text, images, videos and films. In the past 12 months, how often, if at all, have you used such generative AI tools in your work?

Source: Ipsos

Note: Base: all UK adults in work (n 1,868). Fieldwork: 30 May to 4 June 2025.

This pattern aligns with established research findings. DSIT observed a similar trend: that digitally disengaged individuals are more anxious about the impact of AI.²⁶

This trend is not exclusive to adults. A recent survey by The Alan Turing Institute found that 68 per cent of children who use generative AI say they find the technology exciting, as opposed to just 22 per cent of those who don't use it. Similarly, 63 per cent of children who use generative AI say they don't find the technology scary or confusing, as opposed to only 23 per cent of those who don't use it.²⁷

Consistency across different studies suggests that, for all ages, this is a robust relationship rather than a survey artefact.

Building on this insight, the regression analysis also reveals that a lack of trust in AI outputs acts as a drag on people's faith in the technology having positive societal impact – even among people who use the technology regularly. This suggests that experience alone is not sufficient; the quality and reliability of that experience matter.

That is why building digital inclusion and skills is essential, especially among older, lower-income, and underrepresented groups who our data show are currently less confident or feel they do not have access to the technology. Equally critical is ensuring that new users' first experiences with AI are positive and produce reliable results. These early impressions are decisive in shaping whether confidence in AI grows, or whether doubt hardens.

Young adults are much more likely to view AI as an opportunity than older people.

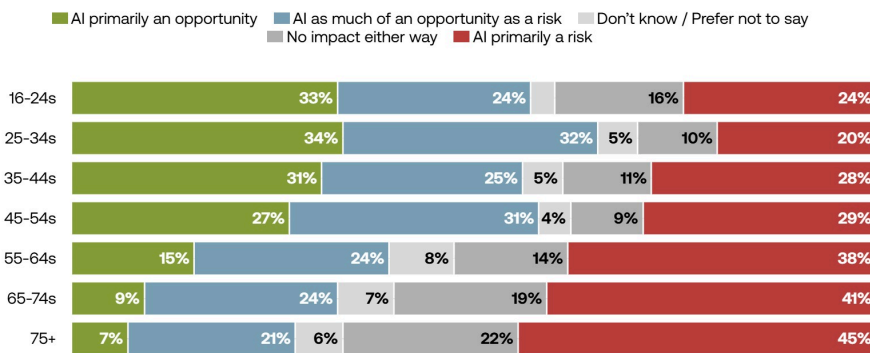
Beyond usage patterns, demographic factors add another layer to these dynamics. To begin, TBI's regression analysis shows that women are 6 per cent less likely than men to view AI as an opportunity for society, even when accounting for other factors like age, income and usage patterns.

In contrast, university graduates are 5 per cent more likely to be positive about AI’s societal effects, again controlling for other factors. Similarly, those in white-collar occupations are more likely to be positive about AI at a statistically significant level. These education and occupation effects point to how social and economic advantages may translate into AI optimism.

However, the strongest predictor of people’s attitudes towards AI is age. Only 16 per cent of people in the 16–24 age group see AI as primarily a risk. However, this perception grows steadily with age, reaching 41 per cent among those aged 65 to 74. This suggests that as younger, more AI-optimistic cohorts age, overall public attitudes may naturally shift towards greater acceptance.

FIGURE 5

More than 30 per cent of people aged 16 to 24 view AI primarily as an opportunity; among people aged older than 55, that number is only 15 per cent



Q: On balance, do you see AI as more of an opportunity or a risk for you personally?

Source: Ipsos

The process and pace at which AI moves from a minority- to a majority-use technology will be decisive in shaping public attitudes over time. From a policy perspective, however, we need to understand and address the adoption barriers people are facing today.

What Is Limiting Public Use of AI?

If wider adoption of AI is key to building confidence in its role across public services, the natural next question is: what is stopping some people from using it?

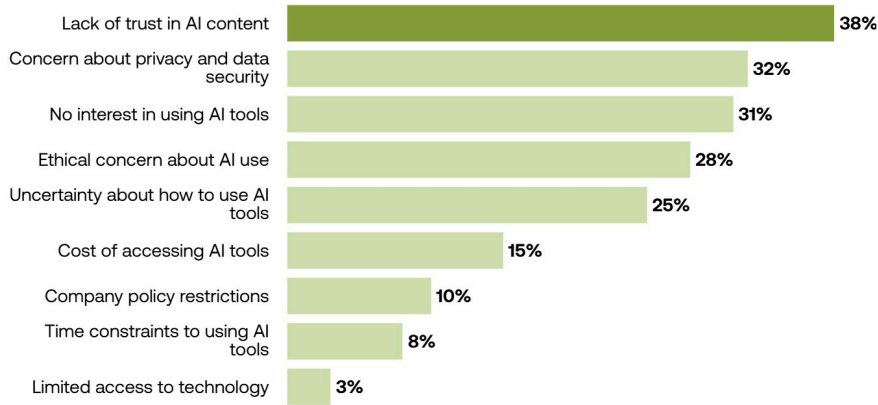
When asked what they felt were barriers to the use of generative AI, respondents most frequently pointed to three factors: a lack of trust in AI content, concerns about privacy and ethics, and a general disinterest in using the technology – especially among people aged 55 and older.

A lack of trust in AI-generated content is currently the biggest barrier to adoption and use of generative AI.

A lack of trust in AI content was cited as a barrier to adoption by 38 per cent of respondents, more than any other factor. This speaks to an inherent limitation of generative-AI systems. When generating text, they do not necessarily produce true statements but rather answers that sound plausible.^{28,29} The risk of “hallucinations” makes generative AI ill-suited for some use cases in high-impact settings like health care and legal analysis, limiting the adoption of such tools.

FIGURE 6

More than 30 per cent of UK adults view a lack of trust in AI-generated content and concerns about privacy and data security as the biggest barriers to adoption



Q: Which of the following barriers, if any, do you face when using or considering the use of generative AI tools (such as ChatGPT, DALL-E, or Midjourney)?

Source: Ipsos

TBI's regression analysis shows how significant this barrier is. Respondents who cite a lack of trust in AI content as a barrier to adoption are 13 per cent less likely to view AI as an opportunity for society at large, while controlling for other factors like age, income, gender and user patterns. Notably, trust in AI content is a significant concern even among those who use AI at least once a week. This suggests that research on truthfulness in AI is important not just for improving specific AI use cases but also for shaping public attitudes towards AI in general.

Other major barriers to adoption are concerns around data privacy and the ethics of AI use, cited by 32 per cent and 28 per cent of respondents respectively. Again, these barriers have technical underpinnings. For

example, researchers have found that generative-AI systems can sometimes expose sensitive information from their training data.³⁰ There are also concerns around how AI developers will treat the confidentiality of conversations users have with chatbots.³¹

Of course, concerns around data privacy when using digital tools and platforms are nothing new.³² In the UK context, however, recent public debates about copyright³³ – and whether artists and rights-holders are being fairly compensated for the data AI labs use to train their models – may have exacerbated the ethical concerns people cite as a barrier to adoption.

A relatively small share of respondents cited financial constraints (15 per cent) and a lack of access to tools (3 per cent) as barriers to adoption. One reason for this may be that many AI labs offer free AI tools that are accessible on any digital device. While this is encouraging, it does not mean that the problem of digital exclusion is solved. More than 1.7 million UK households still lack internet access, according to Ofcom.³⁴ Fixing this should be a top priority for the government.

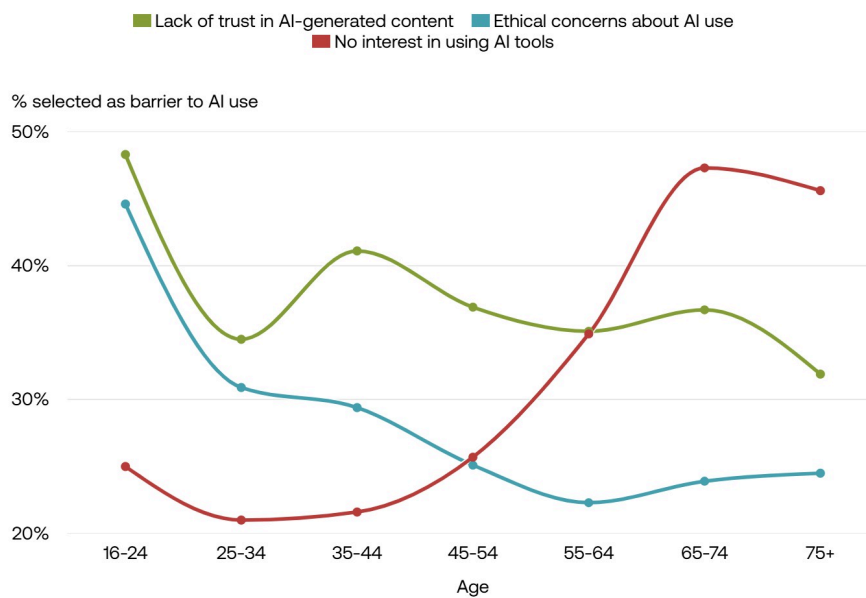
The reasons people cite for not using generative-AI tools vary a lot depending on age and other demographic factors.

When analysing the data, we found stark demographic divides in response to the question about barriers to adoption. Younger people more frequently cite a lack of trust in AI content and ethical concerns as key barriers. This might seem counterintuitive given that they are more positively disposed to AI overall. Most likely it suggests that younger users are more discerning about AI's limitations – and know how to work around them rather than simply dismissing the technology.

By contrast, those aged older than 55 are more likely than younger people to say they are simply disinterested in the technology. As illustrated in Figure 7, people above retirement age have the least interest in generative-AI tools, which are often marketed as productivity boosters. This is another example of how context and incentives shape public attitudes towards AI.

FIGURE 7

For people aged 55+, low interest in using AI is the biggest barrier to adoption



Q: Which of the following barriers, if any, do you face when using or considering the use of generative AI tools (such as ChatGPT, DALL-E, or Midjourney)?

Source: Ipsos


However, age is not the only demographic factor impacting attitudes towards AI. For example, research from the Ada Lovelace Institute and The Alan Turing Institute in March 2025 found that black and Asian people in the UK are more likely than the national average to see certain AI applications (LLMs and mental-health chatbots, and robotics applications such as driverless cars and robotic care assistants) as beneficial, while citing a different set of concerns around their usage.³⁵

Our survey data corroborate this finding, revealing two key differences between white and ethnic-minority respondents. First, ethnic-minority respondents were much less likely to be “disinterested” in using AI,

suggesting greater openness to the technology. However, they were much more likely to cite cost as a barrier to adoption. This points to structural inequalities, with cost potentially excluding groups who are otherwise eager to engage with AI. According to The Alan Turing Institute’s analysis, when all other variables are held constant, those on low incomes still have significantly lower net benefit scores than those with higher incomes.³⁶

FIGURE 8

Ethnic-minority respondents were far less likely to cite being “disinterested” as a barrier to using AI, and much more likely to cite cost



Q: Which of the following barriers, if any, do you face when using or considering the use of generative AI tools (such as ChatGPT, DALL-E, or Midjourney)?

Source: Ipsos

The AI Confidence Gap

Another key factor shaping people's attitudes towards AI is how they feel about AI's potential impact on jobs. Workforce readiness at a structural level and individual confidence in AI skills will therefore have important implications for how AI is adopted across the economy.

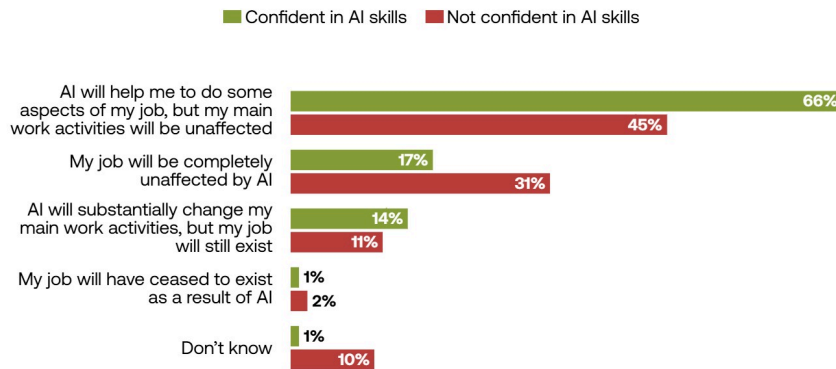
UK adults tend to view AI as a supportive tool – less than 2 per cent believe that AI will take their job in the next 12 months.

According to our survey data, a clear majority of those who feel confident in their own AI skills believe the technology will augment their work, not replace it. Two-thirds (66 per cent) expect AI to help with parts of their job in the next 12 months while leaving their core responsibilities intact. Only 1 per cent believe AI will eliminate their role entirely in the same timeframe. This suggests that confidence breeds a collaborative rather than replacement mindset about AI at work.

But among those not confident in their AI skills, the picture is more uncertain. Fewer than half (45 per cent) see AI as a supportive tool. Nearly a third (31 per cent) believe their jobs will be entirely unaffected and one in ten say they simply don't know what to expect. This may be more problematic than outright opposition, as it leaves workers unprepared for future disruptions.

FIGURE 9

Confidence in AI skills is correlated with seeing AI as a job enabler, rather than a job replacer, over the next 12 months



Q: Which of the following statements best describes how improvements in AI will affect your job 12 months from now?

Source: Ipsos

Note: Base: all UK adults in work (n 1,868). Fieldwork: 30 May to 4 June 2025.

Understanding who feels ready and confident enough in their skills to prepare for the upcoming impact of AI on their jobs, and who does not, is another central challenge for policymakers. Of course, it is not easy to predict how different professions will be affected by AI or other forms of automation. Until recently, a common fear was that manual labour would be replaced by robots. However, recent studies suggest that AI's greatest impact on the labour market may be on white-collar jobs like accounting, software development and knowledge work.^{37,38}

To analyse our data with respect to AI's potential labour-market impact, we built upon recent work by the UK government to model the exposure of different industries to AI. Specifically, the Department for Education has created an AI Occupational Exposure (AIOE) score and applied it across UK industries.³⁹ The AIOE score identifies the following sectors as having a

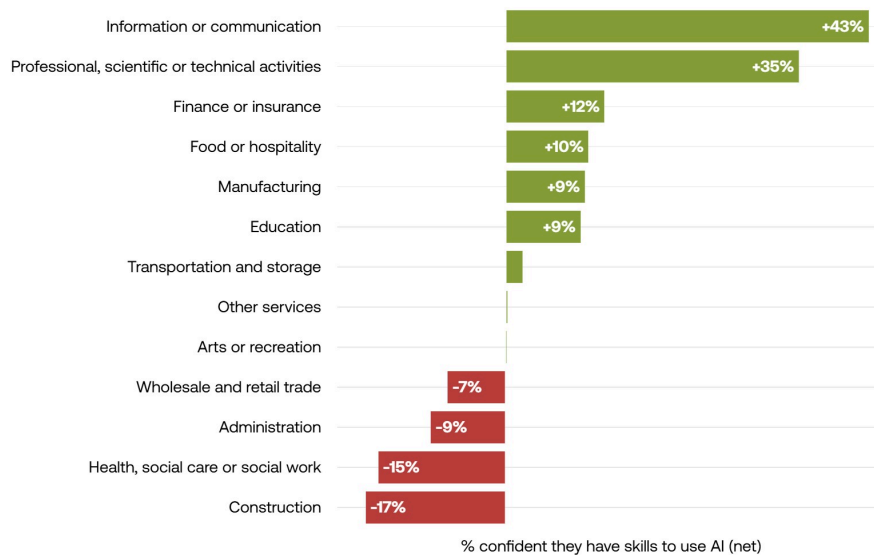
relatively high level of exposure to AI: finance and insurance; information and communication; scientific and technical; property; public administration and defence; education; and health.

High-income earners report higher confidence in AI skills, even when their industries are more exposed to disruption by AI.

Our survey data show that, broadly, people in highly exposed sectors such as “information and communication” and “professional, scientific and technical” fields are among the most confident that they have the requisite skills to use AI and handle potential workflow disruptions. At the same time, in some other highly exposed sectors – education and health and social care, for example – there is a clear “confidence gap” between exposure and skills readiness.

FIGURE 10

Confidence in AI skills is highest among professionals in communication, consulting and research – and low in sectors like health, social care or social work and education



Q: How confident, if at all, are you that you have the skills you might need in order to use AI in your current job role / in the job you want to do?

Source: Ipsos

Note: Net confidence is calculated as the share “confident” minus the share “not confident”, excluding “don’t know” and neutral responses. Data from Wave 2 (2024) and Wave 3 (2025). Sectors with a base under 100 were removed from the analysis. Base: all UK adults in work 16+ May to June 2025 (n 1,868) and all UK adults in work 16+ March 2024 (n 2,506). Combined n 4,374.

This discrepancy in AI-skills confidence between workers in different highly exposed sectors might be caused by socioeconomic factors. To begin, generative-AI tools may simply be easier to integrate into desk-based jobs than manual ones. However, when comparing our findings with data from previous surveys conducted by Ipsos, we also found that recent increases in AI-skills confidence have been concentrated among higher-income respondents, as shown by Figure 11.

FIGURE 11

Recent increases in confidence regarding AI skills have been concentrated among higher-income respondents

| Income | Net confident in AI skills (2024) | Net confident in AI skills (2025) | Change in net AI skills confidence |
|-------------------|-----------------------------------|-----------------------------------|------------------------------------|
| Up to £25,999 | -6% | 0% | +6% |
| £26,000 – £51,999 | -13% | +10% | +23% |
| £52,000 – £99,999 | +8% | +30% | +22% |
| £100,000+ | +3% | +36% | +33% |

Q: How confident, if at all, are you that you have the skills you might need in order to use AI in your current job role / in the job you want to do?

Source: Ipsos

Note: Base: all UK adults in work 16+ May to June 2025 (n 1,868) and all UK adults in work 16+ March 2024 (n 2,506). Combined n 4,374.

Different Contexts Elicit Different Perceptions

Moving beyond general attitudes and workplace concerns, examining where and how the public wants AI deployed reveals crucial insights about acceptance. The evidence consistently shows that attitudes towards AI are highly dependent on context and use cases. Blanket assessments of attitudes thus miss important nuances that policymakers need to understand.


Views on using AI in schools are mixed – but people are more comfortable with AI use in secondary schools than in primary schools.

The education sector provides an example of how context shapes acceptance. While attitudes are mixed, the public expresses relatively more comfort about the use of AI in secondary than in primary schools. This age-based distinction suggests that concerns about AI interact with perceptions

of children's vulnerability and developmental needs. In another survey conducted by The Alan Turing Institute, the majority (76 per cent) of parents or carers whose children use generative AI feel positively about their children's use of the technology.⁴⁰

FIGURE 12

People are more comfortable with having AI tutors in secondary schools (37 per cent) than in primary schools (30 per cent)



Q1: AI tutors are computer programs that can provide personalised learning support, adapt to a student's pace and learning style, and could offer real-time feedback and assistance. How comfortable or uncomfortable are you with using AI to create personalised learning materials and resources, for pupils in primary schools / secondary schools?

Q2: It has been suggested that AI could handle some tasks currently done by teachers (like personalising material and giving feedback). This could free up teachers to give students more direct support. How comfortable or uncomfortable would you be with this use of AI in primary schools / secondary schools?

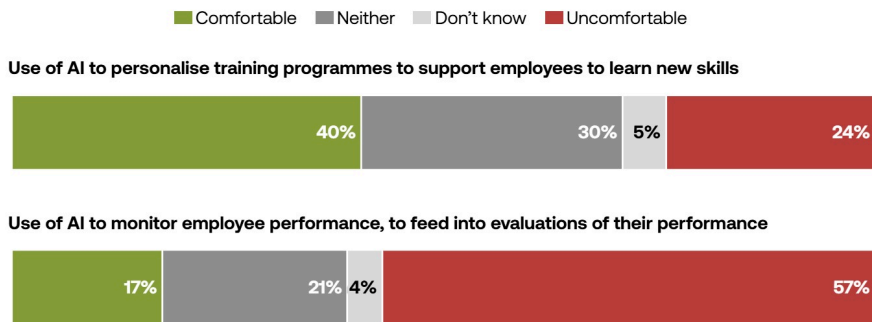
Source: Ipsos

Comfort with AI is highest for use cases that are perceived as socially benevolent.

The pattern whereby attitudes towards AI depend on the specific use cases extends to the workplace too. People are generally comfortable with AI being used to personalise work training – a supportive, developmental application – but express discomfort about the use of AI to monitor employee performance, which implies surveillance and potentially negative personal consequences.

FIGURE 13

40 per cent of UK adults are comfortable with the use of AI to personalise training programmes – but only 17 per cent are comfortable with AI monitoring workers



Q: Now thinking about the current or future use of AI technologies in workplaces, how comfortable or uncomfortable are you with the following?

Source: Ipsos

These examples illustrate a broader point: similar AI systems can elicit very different public responses depending on their application. The variation suggests that public concerns centre not just on AI capabilities, but on how those capabilities might be used and by whom.

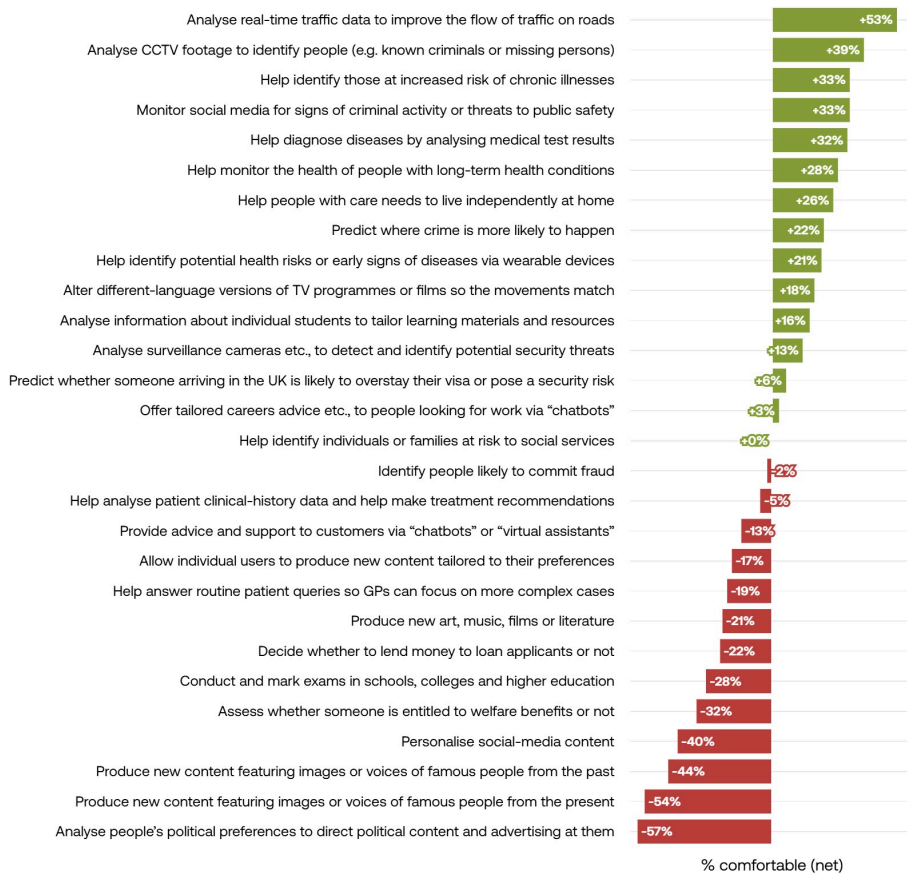
People are most accepting of AI that has tangible and visible public benefits – but reject use cases whose utility is unclear or perceived to only benefit a few.

Another survey by Ipsos in 2024 mapped the variations in public attitudes towards AI across multiple use cases. When asked how comfortable or uncomfortable the public were with current or future use of AI technologies, responses varied significantly – from being largely comfortable with using AI to analyse real-time traffic data to improve traffic flows on roads (+53 per

cent net comfort), to largely uncomfortable with using AI to analyse people’s political preferences to direct political content and advertising at them (-57 per cent net comfort).

FIGURE 14

Public net comfort with AI varies significantly between different use cases



Q: Thinking about the current or future use of AI technologies, how comfortable or uncomfortable are you with the following?

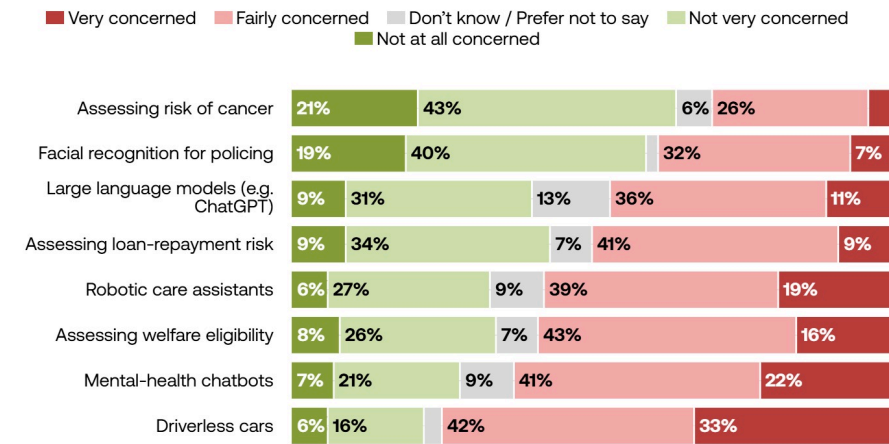
Source: Ipsos Public Trust in AI, 2024

Note: Base: all UK adults 16+ (n 5,150), March 2024; all adults 16+ (n 5,098), September 2023. Note that use cases were asked in either March 2024 or September 2023.

This 56-point swing between the highest and lowest acceptance rates demonstrates how much context influences public attitudes towards AI. Research by the Ada Lovelace Institute and The Alan Turing Institute confirms this insight. Their 2025 study found that while the public sees clear benefits in using AI to assess cancer risk, the benefits of mental-health chatbots remain unclear to many.⁴¹ Similarly, AI use cases that might cause negative effects for individuals – such as assessment of eligibility for jobs or welfare – are met with considerable scepticism.⁴²

FIGURE 15

Concerns about AI depend on context, with people trusting its use for cancer detection but doubting its use for monitoring welfare and work



Q: To what extent are you concerned about the use of this technology?

Source: [Ada Lovelace Institute, Alan Turing Institute](#)

Note: Base: all British adults 18+, November to December 2022 (n 4,010).

This chapter has focused on public attitudes towards AI in the UK. But similar patterns hold true internationally. A 2022 TBI study examined public acceptance of AI across different aspects of public life, drawing on survey data from 26 countries, spanning both developed and emerging economies.⁴³ Respondents expressed strongest support for the use of AI in medical diagnosis and policing while showing more resistance to its deployment in welfare and justice systems.

More than half of UK adults accept the use of AI to diagnose health problems, and support for AI-aided medical diagnosis is highest among people aged older than 55.

The UK survey data from TBI's 2022 study found that 54 per cent of respondents deemed the use of AI to diagnose minor health problems acceptable.⁴⁴ The same number was 49 per cent for fatal diseases. People aged older than 55 appeared to be the most supportive of AI for medical diagnosis, with acceptance rising to 62 per cent and 57 per cent respectively. This contradicts the assumption that older adults are uniformly more sceptical towards AI and highlights the importance of visible, tangible benefits.

This raises important questions about whether the UK's gap in trust with respect to AI in health care might affect the adoption of potentially beneficial innovations by the NHS.

Policy priorities provide another lens for assessing public attitudes towards AI. Research conducted by DSIT has found that despite optimism regarding AI's role in climate monitoring, only 17 per cent of the public consider climate change and the environment to be the most important issues facing the country.⁴⁵ The cost of living (56 per cent), health (35 per cent), the economy (33 per cent) and immigration (31 per cent) are more frequently seen as pressing issues and as more important areas in which to leverage data and AI. This suggests that even when AI applications are viewed positively, they may not receive public support if they don't address immediate concerns.

What Are the Public's Expectations About the Governance of AI?

A key takeaway from our survey data is that members of the UK public, generally speaking, view AI as much of a risk as an opportunity. This naturally leads to questions about what is shaping that perception so that public concerns can best be addressed, including through appropriate levels of oversight and control.


UK adults are more likely to view AI as more of a risk for the economy (39 per cent) than an opportunity (20 per cent).

As highlighted in this chapter, public opinion about AI varies depending on whether the question focuses on personal risks or societal risks. Respondents were more likely, on balance, to say that AI in general is a risk rather than an opportunity for the UK. Even many of those who personally see AI as an opportunity worry about the societal risks.

When asked about the UK economy, our polling data suggest that 39 per cent view AI primarily as a risk (while only 20 per cent see it primarily as an opportunity). Perceptions about AI in public services are similarly risk-oriented. Most concerningly, 59 per cent see AI as a risk for the UK's national security. Combined, these numbers may boost public support for AI-governance initiatives.

FIGURE 16

UK adults are more inclined to view AI as a risk for the economy than an opportunity



Q: On balance, do you see AI as more of an opportunity or a risk for the following?

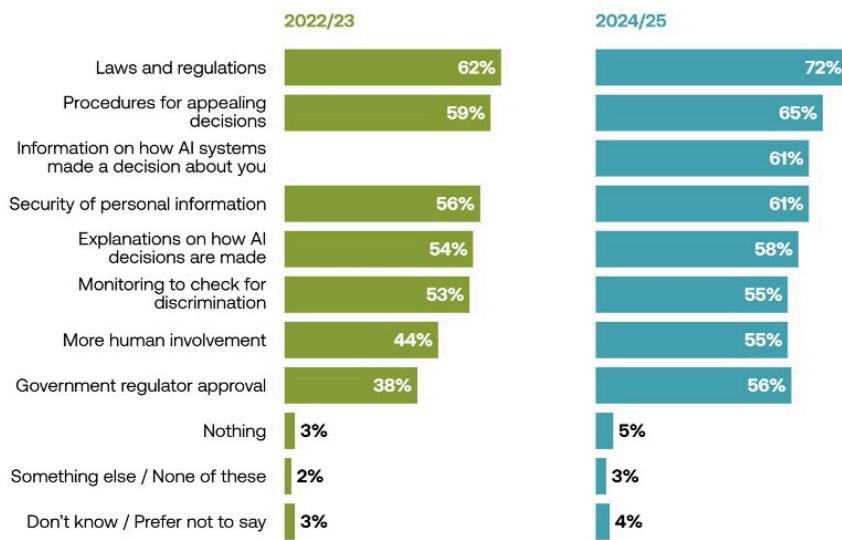
Source: Ipsos

The public expects the UK government to play an active part in managing AI risks.

Against this background, it is not surprising that the public wants the government to address AI risks. Previous research by the Ada Lovelace Institute and The Alan Turing Institute has highlighted that the public expects robust regulations, procedures for appealing decisions, security of personal information, explanations on how decisions are made, monitoring to check for discrimination and human involvement. Crucially, the public also expects the government to play a role in ensuring that AI systems are safe, rather than leaving this task entirely to private companies. It is also worth noting the 10 per cent increase across the two waves of the survey in the percentage of people who say that laws and regulations would make them more comfortable with AI.

FIGURE 17

More than 70 per cent of UK adults say that stricter laws and regulations would make them more comfortable with AI



Q: Which of the following would make you more comfortable with AI being used?

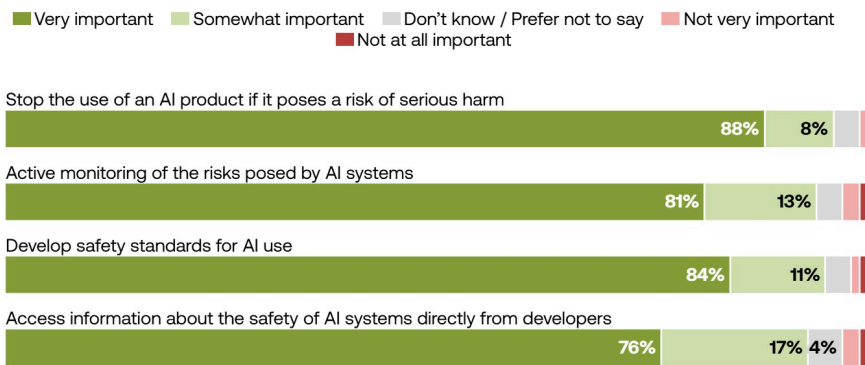
Source: [Ada Lovelace Institute](#), [Alan Turing Institute](#)

Note: "Information on how AI systems made a decision about you" was included in the 2024/2025 wave, but not in 2022/2023.

Base: all British adults 18+, November to December 2022 (n 4,010); all British adults 18+, October to November 2024 (n 3,513).

FIGURE 18

The majority believe it is the government’s responsibility to ensure AI safety by enforcing rules, managing risks and preventing harm



Q: How important is it that the government or regulators, instead of only private companies, have the following powers?

Source: [Ada Lovelace Institute, Alan Turing Institute](#)

Note: Base: all British adults 18+, October to November 2024 (n 3,513).

These expectations are occurring against a backdrop of growing concerns about AI risks. The same study by the Ada Lovelace Institute and The Alan Turing Institute also showed that while perceptions of AI’s beneficial impact have remained stable, concerns around AI uses have increased since 2022/2023. This suggests that familiarity is not automatically breeding acceptance – instead, greater awareness may be highlighting potential problems.

Specific concerns centre on issues of control and privacy. DSIT’s research showed the public is concerned about data security, unauthorised sales of data, surveillance and lack of control over data sharing. Importantly, these issues mirror the themes survey participants recalled hearing about in news

stories.⁴⁶ The study highlighted that overall public associations with AI remain dominated by negative concepts, reflecting persistent fears and concerns.

This alignment between media coverage and public concerns suggests that narratives about AI risks are resonating with, and potentially reinforcing, people's anxieties. This is a problem, as overly restrictive regulations may also incur significant economic and social costs, for instance by slowing down the adoption of socially beneficial AI use cases or by placing unjustifiably high overhead expenses on technology developers. The goal of AI governance should therefore not be to add red tape but rather to support innovation by providing a level playing field and harmonised standards.⁴⁷

Trust in Technology Is Dependent on Trust in Institutions

The UK public's expectations around AI governance highlight an important insight about the relationship between technology acceptance and institutional confidence. The question of who governs AI may be as important as how it is governed.

Past research by DSIT highlighted a crucial finding about how trust operates in practice. Public attitudes towards data sharing are primarily influenced by the organisations that are involved, and survey participants place relatively less importance on how the data is used or the safeguards that are in place. NHS and academic researchers consistently rank high in trust, while social-media companies and the government receive lower levels of trust. This suggests that institutional credibility may trump technical safeguards in shaping public acceptance.

Excitement about AI and support for good AI governance can go hand-in-hand.

These findings about institutional trust have important implications for the debate around AI governance, in which innovation and regulation are often presented as a false dichotomy.⁴⁸

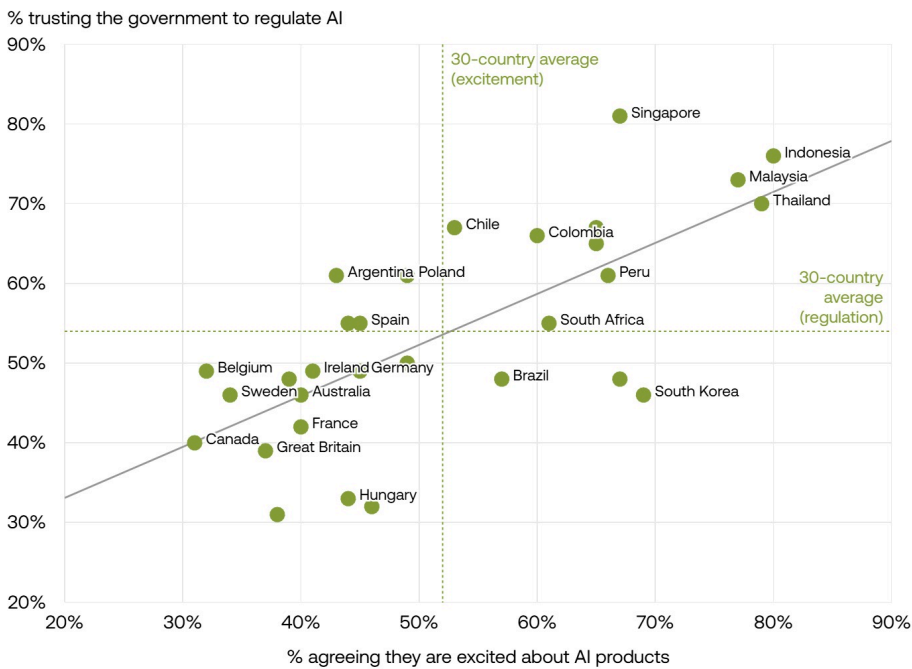
In fact, a careful international comparison shows that excitement for AI and support for good AI governance co-exists in many countries with high institutional trust.

The Stanford Institute for Human-Centered AI's 2025 AI Index Report highlighted global differences in public attitudes.⁴⁹ While participants in countries like China and South Korea report most excitement about AI, people in the UK and USA – where public faith in regulation is lower – are among the most nervous.

Figure 19 shows the correlation between institutional trust and excitement about AI across countries. While many factors may help explain this correlation, the trend runs counter to the idea that AI regulations stand in tension with AI adoption.

FIGURE 19

Excitement about AI is correlated with people's trust in the government's capacity to regulate the technology



Q (AI Regulation): How much do you agree or disagree with the following? *I trust the government of ... to regulate AI responsibly.*
Q (AI Excitement): How much do you agree or disagree with the following? *Products and services using AI make me excited.*

Source: [Ipsos Global AI Monitor 2025](#)

Note: Base: 30-country survey (n 23,216).

The link between institutional trust and excitement about AI has important policy implications. In short, policies that strengthen trust in regulatory institutions are likely to improve public attitudes towards AI and boost excitement about its use.

03

Recommendations to Build Public Trust in AI

The conclusion from our survey data is clear: public attitudes towards AI in the UK are mixed at best. Voters and consumers are right to demand high standards for robustness and transparency. Still, low trust and widespread scepticism pose challenges for the government as it seeks to implement the AI Opportunities Action Plan.

In this chapter we explore what the UK government can do to build justified trust, improve public attitudes towards AI, and accelerate the adoption of safe and beneficial use cases. This will require a combination of sound policies and improved communication.

To build trust in AI, the government should focus on five things:

1. **Strategic communication:** Focus on use cases that matter to people – not efficiency gains.
2. **Real-world evaluation:** Measure AI's impact with human-centric (not technical) benchmarks.
3. **Responsible governance:** Strengthen the UK's sector-specific approach to AI regulation.
4. **Digital upskilling:** Invest in training programmes for safe AI adoption across the population.
5. **Public engagement:** Initiate outreach activities to increase awareness and participation.

In this section, we discuss each of these recommendations and how best to implement them.

1. Strategic Communication

Recommendation 1: AI must be developed as a tool to tackle real-world problems and key public priorities. *People are more likely to engage with, use and support AI when they can see it having a tangible, positive impact on their lives. Political leaders' messaging should focus on specific use cases with benefits that people can recognise (such as quicker scheduling of medical appointments, improved access to public services and shorter commute times due to real-time traffic data). This messaging will also help set the right development goals for AI.*

Today, much of the AI discourse centres on abstract metrics like GDP growth projections, data-centre capacity and global-competitiveness rankings. While economically sound, these narratives often fail to connect with citizens who struggle to see how AI will improve their daily lives. There is also a risk that focusing exclusively on this messaging can deepen scepticism, reinforcing the sense that AI is a project made by elites, for elites rather than an everyday tool that can benefit everyone.

The government should reframe its AI-communications strategy around human outcomes rather than technical capabilities. Instead of announcing new AI initiatives with efficiency metrics, begin with the problems they solve: for example, say “patients will receive a pre-bookable appointment nine days faster on average”, rather than “this will improve scheduling-algorithm performance by 73 per cent”.⁵⁰ Early deployment in areas like health-appointment scheduling, benefit processing and public-transport optimisation creates positive touchpoints that shape broader attitudes.

To build trust, the government should take a data-driven approach and meet the public where they are. The TBI/Ipsos survey data discussed in the previous section revealed that acceptance of AI varies by as much as 56 percentage points between different use cases. While 66 per cent of respondents are comfortable with AI analysing traffic data to reduce commute times, only 10 per cent support its use for political targeting. A good starting point is to centre government communication on use cases

with high levels of public support. Another key point is to focus on use cases that offer not only efficiency gains but new ways of solving real-life problems.

When deploying AI in public services, the government must also make sure that citizens understand four things: why the specific AI system helps address a specific problem, how it delivers for different individuals and groups, what remains under human control, and how they can appeal decisions. Public AI projects should include “counterfactual non-AI” summaries outlining how the same services would be delivered without AI, emphasising both the greater resource requirements and poorer social outcomes this involves.

A final element of this communication strategy should be to foreground UK companies and use cases. The UK government’s current narrative has focused a lot on the importance of attracting big US tech companies. This could undermine trust by suggesting that AI benefits flow primarily to foreign shareholders rather than British citizens. The government should instead highlight UK AI success stories such as Babylon Health’s NHS partnerships, DeepMind’s protein-folding breakthroughs (conducted in London) or Cambridge-based Prowler.io’s work on autonomous systems for UK logistics companies.^{51,52} By celebrating these achievements, this approach demonstrates that AI innovation can emerge from and benefit UK communities directly.

Taken together, this framing and storytelling approach can foster pragmatic optimism and encourage the public to view AI as a tool for solving pressing challenges, not a disruptive force.

Before proceeding, two clarifications are necessary. First, the government should continue to pursue its opportunities agenda, and invest in skills and infrastructure to accelerate safe and beneficial AI adoption. The point is about strategic communication: framing policies based on near-term, demonstrable utility is important to build broad coalitions for positive change.⁵³

Second, the goal is not to make the public uncritically accept AI; people should demand the kind of AI that benefits them. But the government also has a legitimate role in moving public opinion towards beneficial technologies. Just as public-health campaigns increased vaccine uptake by demonstrating clear benefits,⁵⁴ political leaders should actively build support for AI applications that improve social outcomes. A good start would be to tie AI initiatives closer to the government's stated missions and milestones, demonstrating the role that AI can play in delivering concrete progress on voter priorities.⁵⁵

However, good intentions and compelling storytelling alone are not enough. For this approach to work, it must be backed by credible evidence of benefit. This leads to our next recommendation.

Real-World Evaluation

Recommendation 2: Evaluate AI systems for public benefit in real-world settings, using metrics that reflect user experience, not just technical performance. When AI systems are deployed in public services, their benefits should be demonstrated credibly and transparently through a trial-and-evaluate approach. Systems that cause harm or fail to offer benefits should be improved based on the feedback from evaluations and users. Of course, human decision-makers and bureaucratic systems also have limitations. The pragmatic test is therefore not whether an AI system is entirely “free from error”, but whether it improves on an imperfect status quo.

People's experiences with AI shape their attitudes.⁵⁶ The greatest threat to public trust in AI, therefore, is the gap between promised benefits and delivered outcomes. In our survey, 38 per cent of respondents cited a lack of trust in AI content as the main barrier to adoption. This scepticism is rational: claims about AI capabilities are often based on laboratory performance rather than real-world utility, leading to disappointment when they fail to deliver the expected benefits.

Current AI-evaluation practices exacerbate this problem. Developers typically focus on abstract benchmarks that bear little relationship to user experience.^{57,58} An AI system might achieve 95 per cent accuracy on a test data set while proving unreliable or unhelpful in practice. For public services, this mismatch between technical performance and human experience damages both technological and institutional credibility.⁵⁹

To build trust, the government should develop and use human-centric evaluation frameworks for all public AI deployments. These evaluations must measure outcomes that matter to service users rather than system operators. For health-care AI, this means tracking patient outcomes and satisfaction alongside diagnostic accuracy. For benefits processing, it means measuring claimant experience and fairness perceptions alongside administrative efficiency.

This approach requires establishing robust trial-and-evaluate protocols. Before full deployment, AI systems should undergo small-scale pilots with diverse user groups, including vulnerable populations like ethnic minorities who are often excluded from initial testing. These pilots should employ mixed-method evaluations capturing both measurable outcomes (service efficiency, error rates) and human experience (user satisfaction, perceived fairness, accessibility).

The NHS AI Lab's approach to clinical AI provides a model worth scaling.⁶⁰ Their evaluation framework requires demonstration of clinical benefit, not just technical performance, with patient and clinician feedback integrated throughout development. Similarly, the Government Digital Service's user research methodology offers principles for iterative testing that prioritises user needs over system capabilities.⁶¹

Transparent reporting is also essential for maintaining credibility.⁶² Every time AI is deployed in the public sector, the responsible department or agency should produce publicly accessible evaluation reports detailing what worked, what didn't and how outcomes compared to pre-deployment hypotheses. This builds confidence in the government's use of AI while providing learning opportunities for other organisations.

International blueprints already exist. Countries like Finland and the Netherlands have developed public AI registers to increase transparency and engagement.⁶³ The Organisation for Economic Co-operation and Development recommends a common reporting framework for AI incidents, which can be extended into a more general reporting framework for AI outcomes and lessons learned.⁶⁴

To summarise, public acceptance of AI hinges on demonstrable improvements to outcomes. Are patients being diagnosed earlier and more accurately? Are social workers spending less time on paperwork and more with families? Do citizens feel better served and more in control of the decisions and services that affect them directly?

The pragmatic standard for success should be therefore improvement on the status quo.⁶⁵ Human decision-makers and bureaucratic systems also have biases and make mistakes. The test for AI is whether it reduces rather than amplifies these problems while delivering additional benefits. If an AI-assisted benefits-assessment system reduces processing time from eight weeks to three while maintaining accuracy and increasing fairness, it represents progress.

In most cases, however, the best solution is human experts aided by AI tools. For example, a recent study published in *Nature Medicine* showed that physicians using AI performed better than both the control group and AI on its own.⁶⁶

By grounding deployment in human-centred evidence, institutions demonstrate a commitment to public value. This improves the AI systems themselves and strengthens their legitimacy in the eyes of the public, making it easier to scale adoption with confidence and consent. Measuring how AI delivers for people represents good design, good governance and good politics.

Responsible Governance

Recommendation 3: Ensure continued responsible AI governance and oversight that meet the public's expectations around reliability, accountability, data privacy and transparency. The UK's innovation-friendly, sector-specific approach is sound. However, as AI systems evolve, so too must the regulations that govern their design and use. In addition to its ambition to address the safety concerns posed by frontier-AI models, the government should plug gaps in existing legislation, strengthen the capacity of sector-specific regulators and lay the foundations for an assurance ecosystem that covers the entire AI value chain.

Public attitudes towards AI reveal the need for proportionate governance. On the one hand, our survey indicates that even individuals who view AI as a personal opportunity often consider it a societal threat. To build trust in AI, the government must show that mechanisms are in place to address legitimate concerns around trustworthiness, data privacy, algorithmic bias and safety.

On the other hand, 2024 Ipsos data show that citizens believe AI discriminates less than humans in many situations.⁶⁷ Further, economic growth remains a top priority for voters.⁶⁸ This underlines the need for an innovation-friendly approach that balances opportunity with oversight: neither unchecked AI deployment nor overly restrictive regulation will deliver on the public's priorities.

Yet one thing is clear: on AI oversight, citizens expect government leadership, not industry self-regulation. Specifically, the public demands security of personal information, explanations of how AI systems work, monitoring of those systems for false or biased outcomes, human oversight in high-stakes decisions and procedures for appealing AI-driven decisions. These expectations are particularly pronounced among younger demographics – a key Labour constituency.

The UK's sector-specific approach provides the right foundation but requires significant strengthening. Current arrangements leave many regulators under-resourced and lacking AI expertise. For example, the

Financial Conduct Authority has issued guidance on algorithmic trading but lacks the technical capacity required to audit complex AI-driven credit-scoring systems affecting millions of consumers.⁶⁹ Similar capacity gaps exist across health, education and social services.

As TBI's paper [Getting the UK's Legislative Strategy for AI Right](#) argued, the government should prioritise three complementary reforms. First, it should substantially increase AI-related funding for existing regulators, enabling them to recruit technical expertise and develop sector-specific guidance. The current £10 million so far allocated to regulatory capacity building is not enough given the scale of AI adoption across public services.

Second, the government should establish shared regulatory capacity and infrastructure to enable effective coordination without bureaucratic overlap.⁷⁰ This includes common standards for AI-impact assessments, shared databases of algorithmic tools in use across government and joint training programmes to build AI literacy among regulatory staff.

Third, the government should facilitate the development of AI-assurance ecosystems by creating market incentives for responsible development. This includes supporting AI-auditing services, establishing certification schemes for high-risk AI applications, and creating liability frameworks that ensure appropriate risk allocation between developers, deployers and users. The Trusted Third-Party AI Assurance Roadmap published earlier this month is a step in the right direction.⁷¹

While the UK should strengthen existing regulators, it should avoid a big, all-encompassing "AI bill". AI governance covers a wide range of concerns – from copyright to public safety – that are best treated separately.⁷² A sector-specific approach also provides an opportunity to address specific concerns with differentiated policies while avoiding overly restrictive blanket bans.

Estonia's AI Leap 2025 programme – which provides a blueprint for using AI in education – exemplifies this approach. By targeting AI use cases at high-school students while implementing policies to protect primary-school students, Estonia demonstrates how context-specific use and governance can maintain both effectiveness and public comfort.⁷³

Beyond regulation, trust in AI ultimately depends on trust in the institutions responsible for oversight. One way to strengthen trust, therefore, is meaningful public participation in AI development and governance.⁷⁴ However, informed and effective public participation requires investment into AI awareness and literacy. This leads to our next two recommendations.

Digital Upskilling

Recommendation 4: Close skills gaps through inclusive training programmes that equip people from all backgrounds to benefit from AI. The adoption of new tools must be matched by workers' ability to use them. Aimed at building practical AI skills, training programmes should be tailored to different sectors, with a focus on basic AI literacy, real-world applications, accessibility and risk awareness. This training should be co-developed with employers, unions and educators. Strengthening public understanding will help ensure that AI enhances, rather than disrupts, people's lives and livelihoods.

The confidence gap in AI skills presents both a significant barrier to adoption and an opportunity for intervention. Our survey reveals stark sectoral differences: workers in professional, scientific and technical fields feel confident about AI skills, while those in health, social care and education express much lower confidence. This disparity is concerning because health and education sectors face high AI exposure according to government analysis, creating a mismatch between technological deployment and workforce readiness.

The skills challenge extends beyond technical competence to encompass critical evaluation and risk awareness. Research conducted by KPMG and the University of Melbourne found that two-thirds of employees admit

relying on AI output without verification, with more than half making work mistakes due to AI use.⁷⁵ Effective AI skills include prompt engineering, output verification, bias recognition and understanding of system limitations – capabilities that protect both individual users and organisational integrity.⁷⁶ Upskilling should also focus on digital inclusion to ensure that safe AI adoption spreads across the whole population – rather than benefits being concentrated on the better off, and harms concentrated on the most vulnerable.

Our regression analysis demonstrates the importance of addressing skills gaps: people confident in their AI skills are 21 percentage points more likely to view AI as a job enabler rather than a replacement threat. Among those confident in their AI skills, 66 per cent expect AI to augment their work while leaving their core responsibilities intact. Among those lacking confidence in their AI skills, that number is only 45 per cent.

The government should establish a comprehensive national AI skills strategy built around three core principles. First, training should be context-specific and focus on practical use cases within existing workflows rather than technical concepts. A primary-school teacher needs to understand how AI can assist with lesson planning, not machine-learning algorithms. A social worker requires knowledge of bias in risk-assessment tools, not neural-network architecture.

Second, training should focus not only on how to use AI capabilities but also on building risk awareness. This includes skills to recognise critical limitations and manage common AI failure modes like hallucinations, bias and leakage of sensitive information. AI may never be fully error-free, but it can still be useful if people know the risks and how to work around them.⁷⁷

Third, training should be inclusively designed to ensure access for low-income workers, digitally excluded populations and those currently lacking AI skills. This requires multiple delivery mechanisms, including workplace training programmes, community education through libraries and adult-learning centres, and integration with existing professional development frameworks.

To build trust, workplace training programmes should be co-developed with employers, unions and professional bodies. Unions can ensure that trainings address worker concerns about job displacement while building collective bargaining power. Employers can provide workplace context and opportunities for real-world application. Professional bodies can integrate AI competencies into continuing education requirements and certification processes.

The UK already has several good skills programmes under way. For example, the NHS AI Lab's training for health-care professionals and the Department for Education's guidance for school leaders provide models worth scaling in other sectors.⁷⁸

Our recommendations for upskilling and responsible AI governance mirror proposals that the government has signposted in the Technology Adoption Review and the AI Opportunities Action Plan.^{79,80} However, effective implementation will be key. Upskilling makes people feel in control rather than left behind. It enables safer, more effective use of AI tools and fosters a sense of shared responsibility. Crucially, upskilling transforms AI from something done *to* people into something they can work *with* – confidently, critically and creatively.

However, upskilling can only be efficient and effective when there is also a standardisation across different AI tools, and when the tools are user-friendly. There should therefore also be mechanisms that enable the public to participate actively in the design and deployment of public-sector AI use cases.

Public Engagement

Recommendation 5: Initiate public-engagement initiatives that provide people with accessible opportunities to understand what AI is, how it works and why it matters. These could range from a nationwide AI Open House programme, whereby the public is invited to visit institutions that use AI and engage with practitioners, to a publicly broadcasted AI lecture series. Another idea is to establish a National AI Discovery and Participation Centre

that runs interactive exhibits and virtual experiences to demystify AI, and that can act as one of the platforms through which the public can participate in the technology's development.

Public engagement with AI suffers from a disconnect between abstract discourse and concrete experience. While experts debate algorithmic bias, most citizens lack opportunities to see AI systems in action or engage meaningfully in conversations about their development.

Our survey data underscore the importance of direct experience: people who use generative-AI tools at least weekly are nearly 30 percentage points less likely to view AI as a societal risk compared to non-users. However, nearly half the population has never used these AI tools, while 54 per cent know only “a little” about AI.

To build trust, the government should initiate a series of outreach activities to increase public awareness about AI and participation in its design and use. While such activities could take many different forms, we suggest three flagship initiatives to start with.

First, the government should establish recurring nationwide AI Open House events where the public can visit institutions deploying AI technologies and engage directly with practitioners. Partners could include local authorities trialling algorithmic tools, universities conducting AI research, NHS hospitals using diagnostic AI, or transport operators implementing intelligent traffic systems. These events should focus on the demonstration that cannot be experienced online – observing radiologists working with AI-assisted imaging, watching traffic-management centres optimising signal timing or seeing how social services use AI tools responsibly.

Second, the government should create a National AI Discovery and Participation Centre to showcase how AI works and how it can improve everyday life. This facility would offer interactive exhibits, live demonstrations and hands-on experiences tailored for diverse communities. Virtual and mobile components would extend the centre's reach across the country. The centre could build on the UK's strong tradition of public science

engagement, following successful models like the Science Museum or National Space Centre, while serving a distinct function as a national touchpoint for civic engagement and myth-busting.

Third, the government should commission a publicly broadcasted lecture series on AI delivered by expert communicators. Building on the Royal Institution's Christmas Lectures tradition, this programme could reach millions while offering a balanced understanding of how AI works and impacts society. The appetite for such programming has been demonstrated by both Professor Michael Wooldridge's 2023 Christmas Lectures on "The Truth about AI" and DeepMind's 2024 "AI for Science Forum" moderated by Professor Hannah Fry.^{81,82} A similar series of televised, expert-led AI lectures could help counteract the often-sensational portrayals of AI in popular culture.

For maximum impact and inclusion, outreach activities should be designed with accessibility in mind. Reaching diverse communities will require multiple delivery mechanisms (in-person, virtual, mobile), content adapted for different educational backgrounds and active outreach to underrepresented groups. Partnership with community organisations and adult-education providers can ensure broad participation rather than serving already engaged audiences.

Further, the government should leverage the UK's extensive library network as an outreach vector. Libraries are not merely repositories for books but serve as community forums for public education and civic participation.⁸³ The 3,000+ public libraries across the UK already provide digital-inclusion services and could easily extend this to AI-literacy programmes.⁸⁴ Leveraging the UK library network for AI use-case demonstrations or consultations ensures geographic coverage while building on existing infrastructure and community trust.

The goal is to transform AI from something that happens to people into something they can influence. Public engagement creates space for democratic input into AI development priorities, deployment decisions and

governance frameworks. When citizens understand AI's capabilities and limitations, they can make informed choices about how these technologies should be used in their communities and governed by our public institutions.

Conclusion

In an age of rapid technological change, democratic societies must prove they can harness innovation while preserving the values that define them. However, there is currently a gap between the UK government's support for rapid AI adoption and the more cautious attitudes held by the public. Even some technology-friendly MPs have recently pointed out the discrepancy between how ministers speak about AI and how voters view it.⁸⁵

Promisingly, this gap is neither inevitable nor insurmountable. Trust in AI can be built through deliberate action that addresses public concerns and delivers tangible benefits. Two things are needed for the government to improve public attitudes to AI and accelerate adoption. First, robust policies that guide the design and use of safe and beneficial AI systems. Second, a strategic communication plan that centres on real-world AI use cases and reflects voters' priorities.

Trust in AI is inseparable from trust in the institutions that govern it. Public attitudes towards AI reflect deeper concerns about accountability, economic fairness and social inclusion. While technical solutions may help, addressing these concerns requires political leadership committed to ensuring that AI serves everyone, not just those who develop it, deploy it or can access it most easily. Succeeding in this endeavour will be key – not only to the government's ability to implement the AI Opportunities Action Plan, but also to Labour's ability to defeat populists in the next election.

Guest Authors

TBI thanks our external co-authors for their valuable input and contributions to this report:

- Daniel Cameron, Research Director, Ipsos
- Ben Roff, Senior Research Executive, Ipsos
- Helen Margetts, Professor of Society and the Internet, Oxford Internet Institute; Senior Adviser, LSE Data Science Institute

Acknowledgements

The authors would like to thank the following experts for the time they took to review this paper and provide helpful feedback (while noting that reviewing or providing input does not equal endorsement of all the points made in the paper).

- Markus Anderljung (Centre for the Governance of AI)
- Ray Eitel-Porter (Intellectual Forum, Jesus College, Cambridge)
- Tim Flagg (UKAI)
- Chris Jones (Royal Society)
- Daisy McGregor (Anthropic)
- Karen Mcluskie (FCDO)
- Elizabeth Seger (Demos)

Appendix

TBI/Ipsos 2025 Survey Methodology

The Ipsos 2025 findings are based on a survey of 3,727 UK adults aged 16+ via Ipsos's UK Knowledge Panel. Fieldwork was conducted between 30 May and 4 June 2025. Ipsos's UK Knowledge Panel is the UK's largest online random-probability panel. Participants are randomly selected through postal invitations, and devices and data are provided to individuals without internet access to ensure full population coverage. This ensures the sample is truly representative of the UK population when it comes to levels of digital inclusion. Avoiding some of the potential pitfalls of online-only panels is essential when it comes to understanding attitudes to AI. Data are weighted by age, gender, region, Index of Multiple Deprivation quintile, education, ethnicity and number of adults in the household in order to reflect the profile of the population.

By also drawing on the 2023 and 2024 waves of the Ipsos AI tracker we were able to conduct detailed analysis by sectors. Through oversampling, we can analyse reliable regional breakdowns.

Limitations

- In this work, we draw attention to some correlations observed in our data in order to understand and explain the big picture. However, we acknowledge that correlations can be due to confounders, reverse causation, coincidence, etc. Results presented in the paper should not be interpreted as providing evidence about the causal effect of any (policy) interventions.
- Quantitative analysis of survey data is a valuable approach for understanding public attitudes at scale. But we also recognise the importance of qualitative methods – such as interviews, focus groups and ethnographic studies – which offer complementary insights.

Although the primary focus of this work is on quantitative analysis, we hope that our findings will also prompt new questions and lay the groundwork for future research using qualitative approaches.

- All surveys face limitations when it comes to understanding true human preferences and behaviour. To address this limitation to some extent, instead of basing our recommendations entirely on one survey's data, we have drawn from multiple surveys and also cited non-survey literature to support the recommendations.

Endnotes

- 1 <https://www.gov.uk/government/publications/ai-opportunities-action-plan/ai-opportunities-action-plan>
- 2 <https://www.science.org/doi/10.1126/science.aat5991>
- 3 <https://www.gov.uk/government/publications/international-ai-safety-report-2025>
- 4 <https://link.springer.com/article/10.1007/s11023-018-9482-5>
- 5 <https://cepr.org/voxeu/columns/rapid-adoption-generative-ai>
- 6 www.gov.uk/government/publications/public-attitudes-to-data-and-ai-tracker-survey-wave-4/public-attitudes-to-data-and-ai-tracker-survey-wave-4-report
- 7 <https://attitudestoai.uk/>
- 8 <https://link.springer.com/article/10.1007/s00146-024-02128-2>
- 9 <https://hai.stanford.edu/ai-index/2025-ai-index-report>
- 10 www.gov.uk/government/publications/public-attitudes-to-data-and-ai-tracker-survey-wave-4/public-attitudes-to-data-and-ai-tracker-survey-wave-4-report
- 11 <https://journals.sagepub.com/doi/full/10.1177/089444393241235175>
- 12 <https://pmc.ncbi.nlm.nih.gov/articles/PMC7553740/>
- 13 <https://www.pewresearch.org/science/2015/01/29/public-and-scientists-views-on-science-and-society/>
- 14 <https://en.wikipedia.org/wiki/AI%5Fwinter>
- 15 <https://incidentdatabase.ai/>
- 16 <https://uitspraken.rechtspraak.nl/details?id=ECLI:NL:RBDHA:2020:865>
- 17 <https://globaldatajustice.org/gdj/1876/>
- 18 <https://energy.sustainability-directory.com/question/how-does-public-perception-impact-nuclear-energy/>

- 19 <https://www.tandfonline.com/doi/full/10.1080/14786451.2024.2355642>
- 20 <https://institute.global/insights/climate-and-energy/a-new-nuclear-age>
- 21 <https://www.forbes.com/sites/cognitiveworld/2018/08/22/ai-knowledge-map-how-to-classify-ai-technologies/>
- 22 This can include tools that are powered by AI without users thinking about it, like web search engines.
- 23 Unless otherwise specified, results are based on all UK adults aged 16+ (n=3,727).
- 24 <https://openai.com/index/chatgpt/>
- 25 <https://www.gov.uk/government/publications/public-attitudes-to-data-and-ai-tracker-survey-wave-4/public-attitudes-to-data-and-ai-tracker-survey-wave-4-report>
- 26 <https://www.gov.uk/government/publications/public-attitudes-to-data-and-ai-tracker-survey-wave-4/public-attitudes-to-data-and-ai-tracker-survey-wave-4-report>
- 27 <https://www.turing.ac.uk/sites/default/files/2025-06/understanding%5Fthe%5Fimpacts%5Fof%5Fgenerative%5Fai%5Fuse%5Fon%5Fchildren%5F-%5Fwp1%5Freport.pdf>
- 28 <https://direct.mit.edu/coli/article/doi/10.1162/coli.a.16131631>
- 29 <https://dl.acm.org/doi/abs/10.1145/3703155>
- 30 <https://link.springer.com/article/10.1007/s10462-024-11024-6>
- 31 <https://ieeexplore.ieee.org/abstract/document/11023504>
- 32 <https://academic.oup.com/idpl/article/12/3/255/6545822>
- 33 www.gov.uk/government/consultations/copyright-and-artificial-intelligence/copyright-and-artificial-intelligence
- 34 www.ofcom.org.uk/media-use-and-attitudes/media-habits-adults/new-research-on-how-uk-adults-navigate-an-increasingly-online-world
- 35 <https://attitudestoai.uk/assets/documents/How-do-people-feel-about-AI-2025-Ada-Lovelace-Institute.pdf>; <https://www.turing.ac.uk/news/publications/how-do-people-feel-about-ai>
- 36 <https://attitudestoai.uk/assets/documents/How-do-people-feel-about-AI-2025-Ada-Lovelace-Institute.pdf>

- 37 <https://oms-www.files.svdcdn.com/production/downloads/academic/2023-FoW-Working-Paper-Generative-AI-and-the-Future-of-Work-A-Reappraisal-combined.pdf>
- 38 www.sciencedirect.com/science/article/pii/S0167268124004591
- 39 <https://assets.publishing.service.gov.uk/media/656856b8cc1ec500138eef49/Gov.UK%5FImpact%5Fof%5FAI%5Fon%5FUK%5FJobs%5Fand%5FTraining.pdf>
- 40 <https://www.turing.ac.uk/sites/default/files/2025-06/understanding%5Fthe%5Fimpacts%5Fof%5Fgenerative%5FAI%5Fuse%5Fon%5Fchildren%5F-%5Fwp1%5Freport.pdf>
- 41 <https://attitudestoai.uk/>
- 42 <https://attitudestoai.uk/assets/documents/How-do-people-feel-about-AI-2025-Ada-Lovelace-Institute.pdf>; <https://www.turing.ac.uk/research/research-projects/understanding-public-attitudes-ai>
- 43 <https://institute.global/insights/tech-and-digitalisation/tbi-globalism-study-tech-trust>
- 44 <https://institute.global/insights/tech-and-digitalisation/tbi-globalism-study-tech-trust>
- 45 <https://www.gov.uk/government/publications/public-attitudes-to-data-and-ai-tracker-survey-wave-4/public-attitudes-to-data-and-ai-tracker-survey-wave-4-report>
- 46 <https://www.gov.uk/government/publications/public-attitudes-to-data-and-ai-tracker-survey-wave-4/public-attitudes-to-data-and-ai-tracker-survey-wave-4-report>
- 47 <https://www.jstor.org/stable/10.7864/jj.10354693>
- 48 www.lawfaremedia.org/article/ai-regulation-s-champions-can-seize-common-ground-or-be-swept-aside
- 49 <https://hai.stanford.edu/ai-index/2025-ai-index-report/public-opinion>
- 50 <https://unityinsights.co.uk/our-insights/an-evaluation-of-rapid-health/>
- 51 <https://www.england.nhs.uk/london/our-work/gp-at-hand-fact-sheet/>
- 52 <https://deepmind.google/science/alphafold/>
- 53 The 2025 Technology Adoption Review by the Government Chief Scientific Adviser and the National Technology Adviser (<https://www.gov.uk/government/publications/technology-adoption-review>) identified the “lack of awareness of relevant use cases and the difficulties in putting forward a business case with incomplete information about the potential future benefits of a technology” as a key barrier to the adoption of technology by businesses.

- 54 www.england.nhs.uk/long-read/nhs-vaccination-strategy/
- 55 www.gov.uk/missions
- 56 <https://www.tandfonline.com/doi/full/10.1080/10447318.2025.2545467>
- 57 <https://www.nature.com/articles/d41586-025-02462-5>
- 58 <https://www.sciencedirect.com/science/article/pii/S1532046422002799>
- 59 Narayanan, A., & Kapoor, S. (2024). AI snake oil: What artificial intelligence can do, what it can't, and how to tell the difference. In AI Snake Oil. Princeton University Press.
- 60 <https://www.england.nhs.uk/long-read/planning-and-implementing-real-world-ai-evaluations-lessons-from-the-ai-in-health-and-care-award/>
- 61 <https://www.gov.uk/service-manual/user-research>
- 62 <https://onlinelibrary.wiley.com/doi/abs/10.1111/puar.12047>
- 63 <https://oecd.ai/en/dashboards/policy-initiatives/ai-register-2737>
- 64 www.oecd.org/content/dam/oecd/en/publications/reports/2025/02/towards-a-common-reporting-framework-for-ai-incidents_8c488fdb/f326d4ac-en.pdf
- 65 <https://link.springer.com/article/10.1007/s00146-024-02128-2>
- 66 <https://www.nature.com/articles/s41591-024-03456-y>
- 67 www.ipsos.com/en-uk/ipsos-ai-monitor-2024-changing-attitudes-and-feelings-about-ai-and-future-it-will-bring
- 68 <https://www.statista.com/statistics/886366/issues-facing-britain/>
- 69 <https://www.fca.org.uk/publication/corporate/ai-update.pdf>
- 70 <https://www.turing.ac.uk/news/publications/common-regulatory-capacity-ai>
- 71 <https://www.gov.uk/government/publications/trusted-third-party-ai-assurance-roadmap/trusted-third-party-ai-assurance-roadmap>
- 72 <https://academic.oup.com/ijlit/article/doi/10.1093/ijlit/eaee013/7701544?login=false>
- 73 <https://e-estonia.com/ai-leap-2025-estonia-sets-ai-standard-in-education/>

- 74 <https://demos.co.uk/research/citizens-white-paper/>
- 75 <https://kpmg.com/xx/en/our-insights/ai-and-technology/trust-attitudes-and-use-of-ai.html>
- 76 <https://arxiv.org/abs/2505.22073v2>
- 77 <https://www.normaltech.ai//p/chatgpt-is-a-bullshit-generator-but>
- 78 <https://digital-transformation.hee.nhs.uk/building-a-digital-workforce/dart-ed/horizon-scanning/developing-healthcare-workers-confidence-in-ai>
- 79 <https://www.gov.uk/government/publications/technology-adoption-review>
- 80 <https://www.gov.uk/government/publications/ai-opportunities-action-plan/ai-opportunities-action-plan>
- 81 <https://www.cs.ox.ac.uk/news/2281-full.html>
- 82 <https://www.youtube.com/playlist?list=PLqYmG7hTraZBwZQwCxzllsyFxC3WKH%5Fli>
- 83 <https://dcmslibraries.blog.gov.uk/2025/02/26/how-public-libraries-break-down-barriers-to-digital-inclusion/>
- 84 <http://www.librariesconnected.org.uk/facts-and-figures>
- 85 <https://www.politico.eu/article/labour-uk-starmer-ai-policy-growth-reform-election-trust-development-open-source-tech-welfare-populist/>

Follow us

facebook.com/instituteglobal

x.com/instituteGC

instagram.com/institutegc

General enquiries

info@institute.global

Copyright © September 2025 by the Tony Blair Institute for Global Change

All rights reserved. Citation, reproduction and or translation of this publication, in whole or in part, for educational or other non-commercial purposes is authorised provided the source is fully acknowledged Tony Blair Institute, trading as Tony Blair Institute for Global Change, is a company limited by guarantee registered in England and Wales (registered company number: 10505963) whose registered office is One Bartholomew Close, London, EC1A 7BL.