



Department for
Science, Innovation
& Technology

Assuring a responsible future for **AI**

Accelerating the growth of the UK's
AI assurance market

Contents

3	Ministerial foreword
4	Executive summary
6	Introduction
7	The UK's AI assurance market
8	• UK government's vision
9	• The UK's AI assurance market
10	• Future potential
11	Opportunities for future market growth
12	Demand
14	Supply
15	Interoperability
17	Actions to accelerate the growth of the UK's AI assurance market
18	Driving demand for AI assurance
19	Action 1: AI Assurance Platform
20	Increasing supply of third-party AI assurance
21	Action 2: Roadmap to trusted third-party AI assurance
22	Action 3: Collaboration with AISI to drive AI assurance research, development and diffusion
23	Enabling the interoperability of AI assurance
24	Action 4: Terminology Tool for Responsible AI
25	Conclusion

Ministerial foreword



The Rt Hon Peter Kyle MP

**Secretary of State for Science, Innovation
and Technology**

Advances in Artificial Intelligence (AI) are increasingly impacting how we work, live, and engage with others.

Rapid improvements in AI capabilities have made the once unimaginable possible, from providing new ways to identify and treat disease to conserving our wildlife. The UK is home to many organisations at the forefront of driving innovation in AI, and we are committed to building an AI sector that can scale and win globally, ensuring the conditions are right for global AI companies to want to call the UK home.

AI is at the heart of the Government's plan to kickstart an era of economic growth, transform how we deliver public services, and boost living standards for working people across the country. My ambition is to drive adoption of AI, ensuring it is safely and responsibly developed and deployed across Britain, with the benefits shared widely.

AI assurance provides the tools and techniques required to measure, evaluate, and communicate the trustworthiness of AI systems, and is essential for creating clear expectations for AI companies – unlocking widespread adoption in both the private and public sectors. A flourishing AI assurance ecosystem is critical to give consumers, industry, and regulators the confidence that AI systems work and are used as intended. AI assurance is also an economic activity in its own right – the UK's mature assurance ecosystem for cyber security is worth nearly £4 billion to the UK economy.

This 'Assuring a responsible future for AI' report is an original publication that – for the first time – surveys the state of the AI assurance market in the UK, identifies opportunities for future growth, and sets out how the Department for Science, Innovation and Technology's (DSIT) work will seize these opportunities and drive the growth of this emerging industry. It demonstrates that the UK already has a growing AI assurance market, which could move beyond £6.53bn by 2035 if action is taken. Our proactive, targeted programme of work will help to realise this potential, accelerating innovation and investment in AI assurance and driving the adoption of safe and responsible AI across Britain.

Executive Summary

AI is transforming the way we live and work, and rapid developments in its capabilities provide exciting new opportunities to improve the lives of UK citizens.

AI has the potential to radically transform public services across the UK, driving a modern digital government which gives people a more satisfying experience and their time back. It also has an important role to play in realising this government's five missions to rebuild Britain, including building an NHS fit for the future and kickstarting economic growth. AI is already providing new ways to identify and treat disease and is transforming the speed and accuracy of diagnostic services. There are early indications that the UK AI market could grow to over \$1 trillion (USD) by 2035, demonstrating its enormous economic potential.

However, to fully realise its potential, AI must be developed and deployed in a safe and responsible way with its benefits more widely shared. Like all technological innovation, the use of AI poses risks; for example, bias, privacy, and other socio-economic impacts like job loss. Identifying and mitigating these risks will be key to ensuring the safe development and use of AI and driving future adoption.

AI assurance provides tools and techniques to measure, evaluate and communicate the risks posed by AI across complex supply chains. It can help to demonstrate the safety and trustworthiness of AI systems, and their compliance with existing – and future – standards and regulations. AI assurance is therefore a key driver of safe and responsible AI innovation.

The UK government is taking action to realise the benefits of AI and ensure it is developed and deployed safely, equitably and responsibly across Britain. DSIT's Responsible Technology Adoption Unit (RTA) – formerly the Centre for Data Ethics and Innovation – develops tools and techniques that enable responsible adoption of AI in the public and private sectors, and has been working to support the UK's emerging AI assurance industry. The UK government has been working to support the AI assurance ecosystem for some time but much has changed in the rapidly evolving AI landscape. Given the dramatic development of AI capabilities in recent years, alongside a range of governance frameworks developing in the UK and globally, there is need to take stock of the current state and future potential of the UK's AI assurance market.

This report combines analysis from a large-scale industry survey, deep-dive interviews with industry experts, focus groups with members of the public and – for the first time – an economic analysis of the market to reflect on the state of the AI assurance market in the UK. It surveys the UK's AI assurance market and its future potential, explores opportunities to further drive growth, and sets out targeted actions the government is taking to maximise the growth of the UK's AI assurance market to ensure the safe and equitable development and deployment of AI in Britain.

Our research has found:

- There are currently an estimated 524 firms supplying AI assurance goods and services in the UK, including 84 specialised AI assurance companies.
- Altogether, these 524 companies are generating an estimated £1.01bn and employ an estimated 12,572 employees, making the UK's AI assurance market bigger relative to its economic activity than those in the US, Germany and France.
- Despite evidence that both demand and supply are currently below their potential, there are strong indications that the market is set to continue growing, with the potential to exceed £6.53bn by 2035 if opportunities to drive future growth are realised.

There are opportunities to drive safe, responsible, and trustworthy development and deployment of AI by addressing challenges related to demand, supply, and interoperability in the UK's AI assurance market. Lack of understanding among consumers of AI assurance about the risks posed by AI, relevant regulatory requirements, and the value of AI assurance is currently limiting demand for AI assurance tools and services. In addition, a lack of quality infrastructure and limited access to information about emerging AI models restricts the supply of third-party AI assurance tools and services. The use of differing frameworks and terminology between different sectors and jurisdictions has also fragmented the AI governance landscape, inhibiting the interoperability of AI assurance, restricting global trade in AI assurance and limiting adoption of responsible AI.


Executive Summary

UK government will look to mitigate the risks associated with AI and drive adoption of safe and responsible AI by maximising future growth of the UK's AI assurance market. Our proactive, targeted actions will increase demand and supply in the UK market and unlock its global potential.

- 1. We will drive demand for AI assurance tools and services by developing an AI Assurance Platform. This will provide a one-stop-shop for information on AI assurance and host new resources DSIT is developing to support startups and SMEs to engage with AI assurance, including an AI Essentials Toolkit. The platform will help to raise awareness of and drive demand for assurance tools and services among industry and those deploying AI systems.
- 2. We will increase the supply of third-party AI assurance by working with industry to develop a 'Roadmap to trusted third-party AI assurance' and collaborating with the AI Safety Institute to advance AI assurance research, development, and diffusion. This will spur collective action to increase the supply of independent, high-quality assurance tools and services, increasing confidence and trust in the UK's AI assurance market.
- 3. We will enable the interoperability of AI assurance by developing a Terminology Tool for Responsible AI. This will help industry and assurance service providers to navigate the international governance ecosystem, enabling the interoperability of AI assurance and supporting the international growth of the UK's AI assurance market.

Government actions to accelerate the growth of the UK's AI assurance market:

This infographic sets out three key needs for the UK's AI assurance market, opportunities to address these and drive market growth, and the targeted actions government will take to maximise the future growth of the UK's AI assurance market.



Market needs:
Throughout this document you will see these tags, which refer the reader to the market needs the UK government has identified, and the opportunities and government actions to address these needs.



Alan Warburton / © BBC / Better Images of AI / Quantified Human / CC-BY 4.0

Introduction

This report brings together findings from research commissioned by DSIT's Responsible Technology Adoption Unit (RTA) to reflect on the state of the AI assurance market in the UK.

The report identifies key challenges facing the market, future needs, and actions the UK government is taking to meet these needs.

Artificial intelligence (AI) offers transformative opportunities for society and the economy but also poses significant risks. To realise the full potential of AI, we must ensure these risks are mitigated and it is developed and deployed in a way that is safe, equitable, and trustworthy. AI assurance can enable us to measure, evaluate, and communicate the trustworthiness of AI systems and is a key driver of responsible innovation in AI. In the rapidly evolving AI landscape, innovation is happening at pace and there is a need to take stock of the current state and future potential of the UK's AI assurance ecosystem.

This report synthesises findings from:

- An economic analysis of the AI assurance market (delivered by Frontier Economics)
- A large-scale online survey of 1347 business leaders across seven industries (delivered by DG Cities)
- Qualitative interviews with 30 senior business leaders from across the Connected and Autonomous Vehicles (CAV), financial services, and HR and recruitment sectors (delivered by DG Cities)
- Focus groups with 35 members of the public (delivered by Thinks Strategy and Consulting)

These findings provide a snapshot of the AI assurance market at this moment in time, highlights its potential for future growth, and sets out how the UK government's ambitious programme of work on AI assurance will help to realise this potential and ensure AI is developed and deployed safely across Britain.


Assuring a responsible future for AI

The UK AI assurance market

UK government's vision

Assurance is critical to mitigate the risks associated with AI and government has an important role to play in catalysing the development of a world-leading AI assurance ecosystem.

The global assurance market is projected to become a multi-billion-dollar industry and fuel more widespread and impactful adoption of AI across the economy. The UK is well placed to capture this market and jumpstart adoption, with existing world-leading AI expertise.



The UK AI assurance market

Today, there are an estimated 524 firms supplying AI assurance goods and services in the UK, who are generating an estimated £1.01bn Gross Value Added (GVA).

Most significantly, 84 of these firms are UK-based specialised AI assurance companies providing assurance services and products as part of their core offering. This compares to just 17 specialised AI assurance companies identified in HMG's 2023 [AI Sector Study](#), indicating that the market for specialised AI assurance companies has grown significantly in the space of two years. Specialised AI assurance companies are estimated to contribute £0.36bn GVA to the total market and are predominantly microbusinesses and SMEs.

Alongside these 84 specialised AI assurance companies, there are an estimated 225 AI developers, 182 diversified firms, and 33 in-house adopters contributing to the UK ecosystem. Third sector organisations – including not-for-profit organisations, academic institutes and research organisations – also play an important role in the market but are excluded from these figures as they do not provide commercial services.

Our research found that AI assurance suppliers provide two main groups of products and services. The first includes consulting, advisory, training and/or procedural services and tools to help developers and deployers think through their AI assurance strategy and set up the right processes to ensure effective AI assurance. The second is technical tools that are used to assess AI systems. AI accreditation services may become a third group of products provided in this market in the future, though there are currently no fully active certification schemes for AI assurance in the UK. Findings from a [recent DRCF study](#) suggest that the type of services offered tend to vary on the basis of the size of the assurance provider, with larger providers tending to offer a full lifecycle audit with bespoke recommendations and smaller providers tending to offer more hands-off governance audits.


AI assurance suppliers vary in size, with specialised firms and AI developers accounting for the largest share of the AI assurance workforce in the UK (96%), and 60% of suppliers classified as SMEs. Much like related markets, the majority of AI assurance suppliers are concentrated in London (47-69%), with smaller hubs in the South East (5-12%), Scotland (5-11%), and the North West (2-12%).

When compared with other countries, these figures indicate that the UK's AI assurance market is bigger relative to its

economic activity than those of the US, Germany, and France.¹ The UK also presents other characteristics which indicate it can be a leader in the AI assurance market in the future, including high levels of venture capital investment and a competitive advantage in the financial, business, and professional services sectors.

Demand for AI assurance tools and services appears to be most developed in sectors with high sensitivity to risk, and where assurance is already part of their operations and organisational culture. For specialised firms, key sectors of custom include financial services, life sciences, and pharmaceuticals.

Findings from the industry survey and sector deep-dive interviews echoed this, indicating that familiarity with AI assurance and assurance techniques – and similar concepts like AI risk management – is higher in industries in which similar processes and techniques are already embedded into their organisational practices. For example, some business leaders from the connected and autonomous vehicles (CAV) sector were familiar with AI assurance due to pre-existing safety assurance processes in vehicle testing. In financial services, others recognised the value of augmenting existing processes or incorporating AI into pre-existing risk management processes.



¹Economic activity is measured in total GDP in 2023.

The UK AI assurance market

524

Firms supplying AI assurance goods and services in the UK, generating an estimated £1.01bn GVA.

84


UK-based specialised AI assurance companies providing assurance services and products.

£360m

Estimated contribution of AI assurance companies in GVA to total market.

96%

Specialised firms & AI developers accounting for the largest share of the AI assurance workforce (UK).



The majority of AI assurance suppliers are concentrated in London (47-69%), with smaller hubs in the South East (5-12%), Scotland (5-11%), and the North West (2-12%).



The UK's AI assurance market is bigger relative to its economic activity than those of the US, Germany, and France.

Future potential

Our economic analysis suggests that both supply and demand of AI assurance is currently below its potential. However, there are strong indications that the UK market for AI assurance will continue growing in the future.

In total, there is an estimated gap of £1.08bn to £5.05bn between current estimated market turnover and possible market demand expenditure, demonstrating the potential for growth in demand for AI assurance products and services. Alongside demand, there are also signs that supply in the AI assurance market is set to increase. Our research suggests that the market for specialised AI assurance companies has grown significantly over last couple of years, supported by the fact that 80% of UK-based specialised firms in the AI assurance market show growth signals.

AI assurance is also a derivative product, meaning increasing use of AI is likely to result in greater uptake of AI assurance. The growth of the AI market itself is therefore another indicator that demand for AI assurance will continue to rise. The US International Trade Organisation has predicted that the UK AI market will grow to over \$1 trillion (USD) by 2035. The global AI industry is also growing rapidly, with some studies suggesting it may reach \$184bn in 2024.

This booming international AI economy presents a valuable opportunity for the UK to capitalise on its position as a first mover in AI assurance and realise the potential of its AI assurance market.

If the market continues to grow at a similar rate to the last five years, it has the potential to reach £6.53bn by 2035, even if many current market barriers continue to exist. It is therefore possible that the market could grow beyond this. However, further growth will require actions to address remaining challenges. The rest of this report will unpack opportunities for future growth and set out actions UK government will take to capitalise on these opportunities and accelerate the growth of the UK's AI assurance market.

Assuring a responsible future for AI

Opportunities for future market growth

Opportunities for future market growth



Wes Cockx & Google DeepMind / Better Images of AI / AI large language models / CC-BY 4.0

Demand

As the use of AI becomes increasingly widespread and its capabilities continue to advance, it is more important than ever that organisations engage with AI assurance.

AI assurance is a crucial component of wider organisational risk management frameworks. It is also a key pillar for realising the UK's AI governance framework, as it can help organisations to demonstrate compliance with existing and future regulation.

In the King's Speech, the UK government confirmed that it intends to introduce binding requirements on the handful of companies developing the most powerful AI systems. These proposals will build on the voluntary commitments secured at the Bletchley and Seoul AI Safety Summits and will strengthen the role of the AI Safety Institute. They will complement the existing focus on proportionate, sector-specific regulation, and build on the government's ongoing commitment to ensure that the UK's regulators have the expertise and resources to effectively regulate AI in each of their respective domains.

In our research, some companies pointed to limited understanding of existing regulations as a barrier to uptake of AI assurance. Our economic analysis highlighted uncertainty around how rules and regulations will be applied in the UK in the future and how these will interact with other regimes being developed internationally, notably the EU AI Act. This correlates with findings from a recent DRCF study, which suggest an increase in UK assurance users looking to

international legislation — like the EU AI Act and New York Bias Audit Law — when seeking out AI assurance products and services.

Although the UK's regulatory approach for AI is still developing, AI assurance can help to address risks in the here and now. Regardless of the regulatory framework around them, assurance mechanisms will be required to demonstrate the trustworthiness and compliance of AI systems. Demand for AI assurance, however, still relies on awareness and understanding among those developing, deploying and using AI systems of the potential risks associated with AI and how assurance can help to mitigate them.

Our research found that awareness of the risks posed by AI is currently limited among firms deploying AI systems and members of the public. Findings from our economic analysis suggest that most assurance users currently seek AI assurance to mitigate reputational risks, and may not fully appreciate the benefits it can provide to them in mitigating the full spectrum of risks associated with AI. This correlates with insights from the survey and deep-dive interviews, in which industry leaders focused predominantly on the risk of AI causing physical, economic, and psychological harm, which can lead to reputational damage.

There is some evidence to suggest that awareness of risks is improving among the UK public. However, we found that much like industry, members of the public are also primarily concerned by a limited range of risks. These include risks of psychological or physical harm — particularly in high-risk settings — and systemic risks associated with the widespread use of AI, such as job loss and loss of human connection. As a result, buying decisions by end consumers do not provide a clear signal to businesses deploying AI as to how important different risks are to mitigate.

Opportunities for future market growth

“The physical and psychological harm. That is worrying... for example if you're using a GP surgery or some sort of medical system where it makes a diagnosis based on the information you're putting in. It's quite alarming, being given the wrong diagnosis for example.”

Male, distrusting data sceptic

Our research findings indicate that firms deploying AI systems and members of the public also have limited understanding of AI assurance and how it can help to mitigate the risks associated with AI. Only 44% of survey respondents agreed that they felt comfortable demonstrating that the AI systems their organisation develops or deploys comply with existing regulations in the UK. Limited understanding of AI assurance, both at a high level and what it means in practice, was cited as one of the biggest barriers to organisations feeling comfortable demonstrating that the AI systems their organisation develops or deploys comply with existing regulations in the UK.

Knowledge of AI assurance also differs significantly across sectors and between different assurance mechanisms. AI assurance was generally familiar to participants working in the connected and autonomous vehicles sector, whereas participants from other sectors – notably financial services – significantly favoured concepts like risk management and governance. Our research also found that familiarity of AI assurance mechanisms like risk assessments (98%), performance testing (94%), and compliance audits (94%) was much higher among UK industry than mechanisms like bias audits (62%) and conformity assessments (61%). Familiarity with more nascent mechanisms like model cards (37%) and red teaming (35%) was even lower.

“I would say ‘performance testing’ is a term we use a lot. Specifically, ‘safety performance testing’... safety performance indicators are the first things that people look at.”

CAV industry, developer, private sector


Although findings from our focus groups indicate that the public have a high-level understanding of assurance and associated concepts, they do not always know how these apply to AI specifically. Levels of understanding of assurance were higher in areas like safety and security as compared to others like fairness, as participants were able to draw on a greater number of references from other sectors.

We also found that the public do not always desire in-depth understanding of AI assurance itself; instead, when assessing the trustworthiness of technology, participants looked for heuristics, such as brand familiarity, word of mouth, and certification or kitemarks.

“With bias, it comes down to the morals that someone has. How do you put that in AI and how do you measure that?”

Male, distrusting data sceptic

Participants wanted more information about assurance for high-risk use cases, whereas this was deemed unnecessary for lower-risk applications. This suggests that providing detailed information about assurance to end users may not always be necessary but providing ‘shortcuts’ can be an effective way to build trust in many cases.



Alexa Steinbrück / Better Images of AI / Explainable AI / CC-BY 4.0

Opportunities for future market growth



Catherine Breslin & Team and Adobe Firefly / Better Images of AI / Chipping Silicon / CC-BY 4.0

Supply

Even where there is demand for AI assurance, the quality of available tools and services is not clear.

Though a variety of different products and services are available on the UK market, there is currently no quality infrastructure in place to assess and assure that they can effectively identify and mitigate AI harms. This is concerning as one [World Privacy Forum report](#) found that than 38% of AI governance tools they assessed either mention, recommend, or incorporate problematic metrics that could result in harm. Additionally, in a [recent DRCF study](#), UK assurance providers emphasised the importance of technical standards and the ability to develop accreditations against them. However, these standards are still in development, making it is unclear exactly which measures assurance tools and auditing services should adhere to.

In addition, supply of effective AI assurance tools and services relies on access to the production process of AI models; without this, suppliers do not have the right inputs for the development of relevant AI assurance tools. Given their proximity to and visibility of the AI development and testing process, developers have access to information about models that makes them well-informed about the ways they've been developed and the risks they pose. This may help to explain why AI developers are not only key consumers of assurance but also the biggest group of assurance service providers in the UK market.

Despite this, developers alone may not possess the diversity of expertise – including ethics, legal, sociological, and more – that is required to produce the most effective assurance tools. In-house suppliers of AI assurance may also have differing risk appetites to end users of AI systems, leading to misjudged levels of investment in assurance tools and services. Finally, our research found that for the public, (who are often the end users of AI,) the organisation assuring the AI product or service is as important as the process. Participants felt that, if assurance were delivered only by an AI developer that is profit motivated, it would not be trustworthy. These findings highlight the need for third-party AI assurance that is high-quality, independent, and deemed trustworthy by industry and the end users of AI systems.

However, the need for access to AI models poses a challenge for third-party assurance providers. Without access to AI models and collaboration with AI developers, third-party AI assurance providers will not be able to produce appropriate assurance tools, leading to a lack of quality third-party supply in the market. Given the value end users of AI place on independent assurance, lack of such supply may lead to a loss of confidence in the quality and trustworthiness of AI assurance tools and services.

“I'm troubled by the lack of transparency in the whole framework. Who is accountable when it goes wrong? We need a body we can trust, not the makers of this stuff.”

Female, less confident digital user

Opportunities for future market growth

Interoperability

For assurance to be an effective way of measuring, evaluating, and communicating the trustworthiness of AI systems, actors across the ecosystem require a common understanding of what AI assurance is and how AI systems can be assured in practice.

Fostering shared understanding among actors in the UK enables clear and effective communication around the trustworthiness of AI systems, supporting uptake of AI assurance and enabling the development of future standards and accreditation. At the global level, promoting a common understanding across different jurisdictions can make different AI governance regimes interoperable, helping to ensure that AI systems assured in one location are fit for purpose and legally compliant in another.

AI assurance is an emerging market, and industry-wide norms and standards are still developing. Our research found a lack of shared understanding and common terminology being used among consumers of AI assurance. This correlates with findings from a recent DRCF study, in which UK assurance providers emphasised the need for assurance taxonomies and for definitions to be clarified.

Deep-dive interviews with industry leaders revealed differing preferences across sectors for high-level terms to describe the goal of AI assurance. For example, while the terms ‘trustworthy AI’ and ‘responsible AI’ were more commonly used in CAV sector, the term ‘ethical AI’ was more familiar to industry leaders in HR and recruitment, perhaps due to its association with concerns around bias and discrimination, which are frequently cited risks in this sector. Similarly, industry leaders expressed greater familiarity with and use of terms like ‘AI risk management’ and ‘AI governance’ than ‘AI assurance’, with these terms favoured particularly in financial services.

“Is the language of AI assurance clear? I don't know whether it's the language per se; I think there's probably a lack of vocabulary. For example, when you're talking about safe autonomous vehicles... what does ‘safe’ actually mean? ... There's not an identifiable term to capture all of that.”

CAV industry, developer, private sector

Beyond the UK, different governance frameworks have emerged internationally, which use different terminology to describe similar principles, processes and mechanisms to ensure AI is trustworthy and responsible. While UK government refers to ‘AI assurance’, other governance frameworks – such as that of the National Institute of Science and Technology (NIST) in the US – focus on ‘AI risk management’. The use of different terminology has fragmented the international landscape for AI governance, making it challenging for suppliers and consumers of AI assurance to navigate. This was reflected in findings from the industry survey, which indicated that 43% of firms felt that they had limited understanding of both the UK's regulatory approach and those in other countries.

Many different barriers were cited to understanding and using the terminology of AI assurance, including lack of international (31%) and UK standardisation of terminology (25%), lack of clear (26%) and consistent (25%) use of definitions in the UK ecosystem, and limited organisational understanding of AI assurance (26%). There are examples of emerging standards in some domains, with AI safety in particular specified in deep-dive interviews with industry leaders. However, the standardisation process moves slowly, and a lack of consensus on key concepts was echoed throughout.

These findings highlight how the domestic and international governance landscape for trustworthy AI is fragmented, with differing frameworks and terms being adopted in different countries and sectors. This lack of shared understanding is problematic given the vital role that assurance plays in communicating trustworthiness, as it risks limiting the ability of actors in the AI assurance market to effectively communicate the trustworthiness and responsible use of AI with trading partners and consumers both in the UK and internationally.

Assuring a responsible future for AI

Actions

To accelerate the growth of
the UK's AI assurance market


Actions to accelerate the growth of the UK's AI assurance market

For the AI assurance market to flourish both in the UK and globally, we need clear interventions to drive demand for AI assurance, increase the supply of quality third-party assurance tools and services, and support the interoperability of AI assurance within and outside of the UK's borders.

While the UK government continues to develop the expertise, institutions, and understanding to ensure AI regulation is adaptable, proportionate, and pro-innovation, the supporting structures of the AI assurance ecosystem – particularly government – have a vital role to play in advancing these interventions and realising the future potential of the UK market.

Government actions to accelerate the growth of the UK's AI assurance market:

This infographic sets out three key needs for the UK's AI assurance market, opportunities to address these and drive market growth, and the targeted actions government will take to maximise the future growth of the UK's AI assurance market.




Actions

Driving demand for AI assurance

Our research suggests that AI deployers and end users have limited understanding of AI risks and how these can be addressed through AI assurance.

This is confounded by the rapid pace of technological change and the fact that regulatory approaches to AI are still developing.

The UK government has already taken steps to improve understanding of AI assurance. The Responsible Technology Adoption Unit (RTA) published a refreshed [introduction to AI assurance](#) earlier this year, as well as a living [portfolio of AI assurance techniques](#) to support industry to understand how assurance techniques can be used in, and benefit, business. The DSIT-sponsored UK AI Standards Hub also has a [training platform](#), with e-learning modules on AI Assurance. However, further action is required to address knowledge gaps among assurance users and to clarify the UK's AI governance landscape and associated assurance requirements.



Actions

Action 1: AI Assurance Platform

To address risks and ensure that AI is governed responsibly, an increasing number of AI assurance tools and services, AI governance frameworks and relevant standards are emerging on the market.

However, as our research has identified, many organisations have difficulty navigating this increasingly complex landscape. Organisations struggle to understand or identify which frameworks, standards, and regulatory requirements are relevant to a particular context of use. This landscape is particularly challenging for startups and SMEs to navigate, given their relative size and experience operating in the market. Action is required to increase understanding, as well as aid navigation of the UK's AI assurance and governance landscape.

DSIT is seeking to develop an AI Assurance Platform to help AI developers and deployers to navigate this complex landscape. The platform will act as a one-stop-shop for AI assurance, bringing together existing assurance tools, services, frameworks and practices together in one place. By raising the profile of assurance and signposting users to best practice, this platform will drive demand for assurance, supporting the growth of the UK's AI assurance market.

The AI Assurance Platform will house existing DSIT guidance and tools, such as the Introduction to AI Assurance and the Portfolio of AI Assurance Techniques. DSIT will also develop new resources for the platform, including an AI Essentials toolkit. Like Cyber Essentials, AI Essentials will distil key tenants of relevant governance frameworks and standards to make these comprehensible for industry. Over time, we will create a set of accessible tools to enable baseline good practice for the responsible development and deployment of AI. This suite of products will help support organisations to begin engaging with AI assurance, and further establish the building blocks for a more robust ecosystem.

The first tool we are developing is AI Management Essentials. Drawing on key principles from existing AI-related standards and frameworks – including ISO/IEC 42001 (Artificial Intelligence - Management System), the EU AI Act, and the NIST AI Risk Management Framework – AI Management Essentials will provide a simple, free baseline of organisational good practice, supporting private sector organisations to engage in the development of ethical, robust and responsible AI. The self-assessment tool will be accessible for a broad range of organisations, including SMEs. In the medium term we are looking to embed this in government procurement policy and frameworks to drive the adoption of assurance techniques and standards in the private sector. The insights gathered from this self-assessment tool will additionally help public sector buyers to make better and more informed procurement decisions when it comes to AI.

Spotlight on supporting initiatives

AI Standards Hub

Global technical standards can help operationalise the cross-sectoral principles by providing common, agreed-upon metrics and benchmarks that reflect industry best practice and that support industry and regulators to achieve these goals. Global technical standards allow others to trust the evidence and conclusions presented by assurance providers.

The UK government has undertaken work to support the development of AI technical standards, including conceptualising and funding the AI Standards Hub. The AI Standards Hub supports UK stakeholders to participate in Standards Development Organisations (SDOs) through knowledge sharing, community and capacity building, and strategic research.

Introduction to AI assurance

It is more important than ever for organisations to start engaging with AI assurance and leveraging its critical role in building and maintaining trust in AI technologies. However, the AI assurance landscape can be complex and difficult to navigate, particularly for SMEs.


To help organisations upskill on topics around AI assurance and governance, the UK government recently published its 'Introduction to AI Assurance'. This guidance aims to provide an accessible introduction to both assurance mechanisms and global technical standards to help industry better understand how to build and deploy responsible AI systems. In doing so, it will help to build common understanding of AI assurance among stakeholders across different sectors of the UK economy.

Actions

Increasing supply of third-party AI assurance

Supply of AI assurance is currently concentrated in AI developers, with a lack of quality infrastructure and information asymmetries restricting the ability of third-party service providers to develop and deploy independent assurance tools and services.

This is problematic given the value that end users place on independence as a mark of trustworthiness. DSIT has already undertaken work to support the growth of the market for third-party AI assurance, including identifying lessons learned from mature certification schemes. However, additional interventions are required to drive the supply of third-party assurance tools and services to increase confidence and trust in the UK's AI assurance market.



Actions

Action 2: Roadmap to trusted third-party AI assurance

To increase the supply of independent, high-quality assurance, DSIT will work with industry to develop a ‘Roadmap to trusted third-party AI assurance’ by the end of the year.

This roadmap will set out our vision for a market of high-quality, trusted AI assurance service providers and the actions required to realise this vision. To develop this roadmap, we will work closely with industry stakeholders and the UK’s quality assurance infrastructure to understand their needs and requirements, ensuring that actions to support third-party supply of AI assurance adds value without being overly burdensome or restricting the growth of the market.

The roadmap will provide an opportunity for us to explore all possible avenues to realise a trusted third-party AI assurance market, including professionalisation. Other sectors – such as cybersecurity, accountancy, and medicine – are underpinned by robust professional regimes, with professional bodies that provide specific training and uphold minimum professional standards. Our research has demonstrated the importance of independence in driving trust, and the valuable role that kitemarks and similar heuristics can play in effectively communicating the trustworthiness of technology with the end users of AI systems. By developing a ‘Roadmap to trusted third-party AI assurance’, we intend to spur collective action among stakeholders across the ecosystem to help drive supply of third-party assurance and increase confidence in its quality and trustworthiness.



Actions

Action 3: Collaboration with the AI Safety Institute to advance assurance research, development and adoption

As the capabilities of AI continue to advance, new techniques for evaluating and assuring AI systems are required to ensure these systems are developed and deployed safely and responsibly.

To scale up supply and drive adoption of new safety and assurance techniques, the Responsible Technology Adoption Unit and the AI Safety Institute will work together to advanced assurance research, development and adoption.

Government will allocate additional funding for the AI Safety Institute's Systemic Safety Grant programme, as well as extra funding to expand work to stimulate the AI assurance ecosystem. The Department for Science, Innovation and Technology will also explore other options for growing the domestic AI safety market, including potentially capital investment or additional grant mechanisms, and provide a public update on this by Spring 2025.

Spotlight on supporting initiatives

External assurance of public-facing AI models

For some AI use cases — particularly public-facing models and services like social media platforms and publicly accessible large language models (LLMs) — [external scrutiny](#) is critical to ensure policymakers, civil society, and the users of these models can influence decisions related to their training, deployment and use. Here, independent assurance has an important role to play in ensuring meaningful public accountability.

One key enabler of external scrutiny is researchers' access to data from these services, which can help to understand the capabilities and controllability of models. In many cases, this data contains sensitive or commercially confidential information. Privacy Enhancing Technologies (PETs) offer opportunities to enable this data sharing while minimising risks to privacy or commercial confidentiality, helping to address information asymmetries between AI developers and independent assurance providers.

Several ongoing initiatives are exploring the opportunities for PETs to enable assurance of AI models, for example, work within the [Christchurch Call](#) on terrorist and extremist content online, and [DSIT's AI Safety Institute's partnership](#) with OpenMined.


The UK government is working to support wider adoption of PETs. As part of this, it has published an interactive PETs adoption guide to help organisations assess whether PETs could be useful in their context. DSIT also co-ran PETs prize challenges with the US government to incentivise novel innovation in PETs and the winners of these challenges were announced at the Summit for Democracy in 2023.

Actions

Enabling the interoperability of AI assurance

The domestic and international governance landscape for trustworthy AI is currently fragmented, with differing frameworks and terms being adopted in different countries and sectors.

Lack of shared understanding across the ecosystem restricts the ability of actors in the AI assurance market to effectively communicate about the trustworthiness of AI with trading partners and consumers in the UK and internationally, inhibiting the vital role assurance plays in communicating trustworthiness. The UK government has taken action to promote a shared understanding of AI assurance among actors within the UK ecosystem, including publishing an accessible [introduction to AI assurance](#) and [sector-specific guidance](#) for industry. However, further government action is required to enable the interoperability of AI assurance tools and services and support cross-border trade in trustworthy AI.



Actions

Action 4: Terminology Tool for Responsible AI

AI assurance is a nascent industry and standards and norms are still emerging. Our research found there are currently notable differences in the way AI assurance is understood across different sectors within the UK and different jurisdictions internationally.

Though we are unlikely to see all actors conform to the use of standardised language in the near future, we require common understanding to enable effective communication between different actors in the ecosystem in the interim.

To aid this common understanding, DSIT is developing a Terminology Tool for Responsible AI, which will define key terminology used in the UK and other jurisdictions and the relationships between them. The Tool will help industry and assurance service providers to navigate key concepts and terms in different AI governance frameworks to communicate effectively with consumers and trading partners within the UK and other jurisdictions, supporting the growth of the UK's AI assurance market.

We have already started working with the US National Institute for Standards and Technology (NIST) and the UK's National Physical Laboratory (NPL) to develop this tool to aid interoperability across UK-US governance frameworks. The US has the largest AI market in the world and is a world leader in responsible AI, already committing to work with UK government to advance global AI safety. By promoting common understanding between our AI governance regimes, the tool provides an exciting opportunity to help unlock the commercial potential of the US market and further strengthen our governments' shared ambitions to drive the responsible and trustworthy use of AI.


Longer term, we envisage the tool acting as a global point of reference on terminology for responsible AI. We are considering ways to expand its scope to incorporate terminology from other countries and make it accessible to industry and assurance service providers all over the world, including ensuring the tool is aligned with the OECD's AI Principles and other interoperable frameworks.

Spotlight on supporting initiatives

Sector-specific guidance

To aid common understanding of AI assurance within different sectors of the UK economy, DSIT is developing sector-specific guidance focusing on assurance good practice. Recognising that assurance is understood and used differently in different sectors, this guidance will set out how tools for responsible AI can manage the risks associated with AI and help to build trust, translating the UK's AI assurance framework into sector-specific recommendations to drive adoption of AI assurance across UK industry. The guidance will be aimed at a non-technical audience, including organisations without a comprehensive AI strategy.

We published the first piece of guidance for companies procuring and deploying AI for recruitment in March 2024 and will publish subsequent guidance for other sectors – including financial services – in the near future.



Nidia Dias & Google DeepMind / Better Images of AI / AI for Biodiversity / CC-BY 4.0

Conclusion

This report explored opportunities to further grow the UK's AI assurance market and outlined the targeted actions UK government will take to drive demand, increase third-party supply, and ensure the interoperability of AI assurance. The UK government is committed to continuing to support the development of an effective AI assurance ecosystem.

However, we cannot deliver the actions outlined in this report alone and collective action is required among actors across the AI assurance ecosystem. We welcome organisations with ideas or opportunities for future collaboration, insights or resources to share.

To get in touch, email us at ai-assurance@dsit.gov.uk.

Together, we can drive the growth of this emerging industry and help to ensure the safe, equitable and responsible development and deployment of AI.



Department for
Science, Innovation
& Technology

Acknowledgements

We would like to thank the organisations who conducted the research that informed this report, including DG Cities, Frontier Economics, and Thinks Strategy and Insight.

The images used in this publication have been sourced from Better Images of AI, a non-profit exploring, creating and commissioning new ways to visually represent AI. The images can be downloaded free of charge under a Creative Commons license from betterimagesofai.org