


Research Briefing

20 August 2024

By Elizabeth Rough,
Nikki Sutherland

Artificial intelligence: A reading list



Summary

- 1 Overviews of artificial intelligence
- 2 UK Government policy and regulation
- 3 Use of AI in different sectors
- 4 Safety and ethics

Disclaimer

The Commons Library does not intend the information in our research publications and briefings to address the specific circumstances of any particular individual. We have published it to support the work of MPs. You should not rely upon it as legal or professional advice, or as a substitute for it. We do not accept any liability whatsoever for any errors, omissions or misstatements contained herein. You should consult a suitably qualified professional if you require specific advice or information. Read our briefing [‘Legal help: where to go and how to pay’](#) for further information about sources of legal advice and help. This information is provided subject to the conditions of the Open Parliament Licence.

Sources and subscriptions for MPs and staff

We try to use sources in our research that everyone can access, but sometimes only information that exists behind a paywall or via a subscription is available. We provide access to many online subscriptions to MPs and parliamentary staff, please contact hoclibraryonline@parliament.uk or visit commonslibrary.parliament.uk/resources for more information.

Feedback

Every effort is made to ensure that the information contained in these publicly available briefings is correct at the time of publication. Readers should be aware however that briefings are not necessarily updated to reflect subsequent changes.

If you have any comments on our briefings please email papers@parliament.uk. Please note that authors are not always able to engage in discussions with members of the public who express opinions about the content of our research, although we will carefully consider and correct any factual errors.

You can read our feedback and complaints policy and our editorial policy at commonslibrary.parliament.uk. If you have general questions about the work of the House of Commons email hcenquiries@parliament.uk.

Contents

Summary	5	
1	Overviews of artificial intelligence	11
1.1	Different types of AI	11
	Machine learning and deep learning	11
	Foundation models and generative AI	12
2	UK Government policy and regulation	13
2.1	UK Government publications	13
	Strategies and roadmaps	13
	White papers	13
	Risks and safety	13
	AI Safety Summit, November 2023	13
	AI Standards Hub	13
2.2	Government Office for Science (GO-Science)	14
2.3	Information Commissioner’s Office (ICO)	14
2.4	Digital Regulation Cooperation Forum	14
2.5	Commentary on UK government’s AI policy & regulation	15
2.6	UK Parliament select committee reports	16
	House of Commons	16
	House of Lords	16
	Government responses	16
3	Use of AI in different sectors	19
3.1	Business and industry	19
	Use of AI in supply chains	20
3.2	Creative industries	20
3.3	Defence and security	20
3.4	Education	21
3.5	Energy	21

3.6	Financial services	22
3.7	Healthcare	22
3.8	Retail	23
3.9	Travel and transport	23
4	Safety and ethics	24
4.1	General	24
4.2	The 'black box' problem	24
4.3	Bias in AI	25
4.4	Misuse and harms	25

Summary

What is artificial intelligence (AI)?

Artificial intelligence (AI) can take many different forms and there is no single, universally agreed definition. The term is frequently used as a shorthand to refer to technologies that perform the types of cognitive functions typically associated with humans, including reasoning, learning and solving problems.

To perform these types of functions, AI systems generally rely on vast amounts of data. This data may be 'structured' or 'unstructured'. Structured data is typically stored in a fixed format and can be more easily analysed and processed, such as financial transactions that have a date, time and amount. Unstructured data includes images, videos and text files; it is not organised according to a predefined structure, it is generally unformatted and is much harder to analyse.

Both types of data can be used to 'train' AI so that it can recognise patterns and correlations. This is achieved by the AI system applying rules (algorithms), based on the training dataset, to interpret new data and perform a specific task. In some instances, the AI is supervised and trained with data sets labelled by humans, as explained in this [example from IBM](#):

A data scientist training an image recognition model to recognize dogs and cats must label sample images as “dog” or “cat”, as well as key features—like size, shape or fur—that inform those primary labels. The model can then, during training, use these labels to infer the visual characteristics typical of “dog” and “cat.”

This is useful for AI systems designed to look for specific things, such as spam emails.

In other instances, the system is unsupervised and the data is left unlabelled. Under these conditions, the system autonomously identifies patterns in the data. This is useful where the AI is designed to find something that is not known in advance, such as online shopping recommendations.

Machine learning and deep learning

The process of an AI system developing and improving over time, without all its instructions being explicitly programmed, is called 'machine learning'.

'Deep learning' is a type of machine learning where the design of algorithms is informed by the structure and function of the human brain and the way it transmits information. Deep learning can be seen in numerous applications,

such as voice and image recognition, as well as in ‘foundation models’, of which ‘large language models (LLMs)’ like ChatGPT, are one example.

LLMs refers to those models that are typically (though not always) trained on very large, unlabelled data sets. They can be adapted to do many tasks, despite not having been trained explicitly to do those tasks. In other words, the model can take information it has learnt about in one situation and apply it to another, different situation.

Different types of AI

Distinctions have also been drawn between ‘narrow’ and ‘general’ AI. Narrow AI is designed to perform a specific task, such as speech recognition, using information from specific data sets. It cannot adapt to perform another task. These are often tools that aim to assist, rather than replace, the work of humans.

Artificial general intelligence (AGI) is an AI system that can reason, analyse and achieve a level of understanding that is on a par with humans; something that has yet to be achieved.

Generative AI has more sophisticated capabilities than narrow AI but has not reached the level of AGI. [Generative AI relies on deep learning models](#) that can use unstructured data to “‘learn’ to generate statistically probable outputs when prompted”, generating new content, such as images.

‘[Frontier AI](#)’, or the term ‘frontier model’, is also being increasingly used to describe AI systems that are “‘cutting edge” foundation models, though the [Ada Lovelace Institute](#), a data and AI research organisation, states that ‘frontier’ is a contested term and that “there is no agreed way of measuring whether a model is ‘frontier’ or not”.

UK Government policy and regulation

The UK does not have any AI-specific regulation or legislation covering the technology. Instead, AI is regulated in the context in which it is used, through existing legal frameworks, such as financial services legislation.

Some regulators, however, have oversight of the development, implementation and use of AI more broadly. For example, the Information Commissioner’s Office (the UK’s independent body established to uphold information rights) has guidance on its website covering [AI and data protection](#) and [explaining decisions made with AI](#).

The Johnson and Sunak governments started developing a more comprehensive regulatory framework for AI. This included publishing strategy documents and a white paper on AI.

The [UK's National AI Strategy](#) was published by the government in 2021. It is a 10-year plan to make the UK a “global AI superpower” by focusing on the following aims:

1. Invest and plan for the long-term needs of the AI ecosystem to continue our leadership as a science and AI superpower;
2. Support the transition to an AI-enabled economy, capturing the benefits of innovation in the UK, and ensuring AI benefits all sectors and regions;
3. Ensure the UK gets the national and international governance of AI technologies right to encourage innovation, investment, and protect the public and our fundamental values.

As part of its aim to “ensure the UK gets the national and international governance of AI technologies right” the government ran a [public consultation on regulating AI in 2022](#). A white paper – [A pro-innovation approach to AI regulation](#) – followed in March 2023.

In the paper, the government proposed that AI would continue to be overseen by existing regulators covering specific sectors, such as Ofcom (the UK’s communications regulator), Ofgem (the energy regulator in Great Britain), and the Financial Conduct Authority (the UK’s conduct regulator for financial services). This context-based approach to regulation was favoured by the government, rather than creating a single regulatory function, and uniform rules, to govern AI.

The government proposed that AI regulation would be informed by five, cross-sector principles which regulators will “interpret and apply within their remits in order to drive safe, responsible AI innovation”. The principles are:

- Safety, security and robustness
- Appropriate transparency and explainability
- Fairness
- Accountability and governance
- Contestability and redress

While the Sunak government decided against creating a single regulatory function to govern AI, it proposed that existing regulators would be aided by “central support functions”, established by the government, such as horizon scanning for emerging risks and trends, and monitoring the overall regulatory framework. The government also proposed that the five principles would not, at least initially, be subject to new regulation, but instead would be implemented by existing regulators.

[A further consultation accompanied the publication of the white paper](#). The [government's response](#) was published in February 2024 and confirmed that it remained committed to its cross-sector principles and “context-specific” approach to regulation. It said that it would seek to build on this in the future, only legislating when it was “confident that it is the right thing to do”.

In February 2024 the government also published guidance for regulators, [Implementing the UK's AI Regulatory Principles](#) (PDF). The guidance sets out the “considerations that regulators may wish to have when developing tools and guidance to implement the UK's approach to AI regulation”. Alongside the guidance, the [government asked regulators](#) to set out their approach to AI. Regulators' responses were published in May 2024 at [Regulators' strategic approaches to AI](#).

Labour government

Both the Labour Party's 2024 election manifesto and the King's Speech indicate that the Labour government is looking to take a different approach to the previous government towards regulating AI. In its [2024 manifesto](#) (PDF) the Labour Party said it would “ensure the safe development and use of AI models by introducing binding regulation on the handful of companies developing the most powerful AI models”. Similarly, in the [King's Speech](#), the Labour Government said that it would “harness the power of artificial intelligence as we look to strengthen safety frameworks” and signalled that it would “place requirements on those working to develop the most powerful artificial intelligence models”.

Use of AI in different sectors

AI is currently being used across many different industries, from finance to healthcare. In 2022, [the UK Government reported](#) that “around 15% of all businesses (432,000 companies) had adopted at least one AI technology”. The [management consultancy firm McKinsey](#) identified the following examples of AI that can be applied across a range of industries:

- Predictive maintenance

[...] Rather than waiting until a piece of equipment breaks down, companies can use predictive maintenance to project when maintenance will be needed, thereby preventing downtime and reducing operating costs. Machine learning and deep learning have the capacity to analyze large amounts of multifaceted data, which can increase the precision of predictive maintenance. For example, AI practitioners can layer in data from new inputs, like audio and image data, which can add nuance to a neural network's analysis.

- Logistics optimization

Using AI to [optimize logistics](#) can reduce costs through real-time forecasts and behavioural coaching. For example, AI can optimize routing of delivery traffic, improving fuel efficiency and reducing delivery times.

- Customer service

[AI techniques in call centers](#) can enable a more seamless experience for customers and more efficient processing. The technology goes beyond understanding a caller's words: deep-learning analysis of audio can assess a customer's tone. If a caller is getting upset, the system can reroute to a human operator or manager.

[The business magazine Forbes](#), for example, highlighted the use of AI by the retailer Marks and Spencer which had deployed "autonomous event capture" to help identify unsafe behaviours in its warehouses. Since adopting the technology, it reported an 80% reduction of warehouse accidents.

AI is also used in the public sector. [The NHS AI Lab](#), for example, is focused on developing and deploying AI systems in health and care. In 2023, the [government provided funding of £21 million](#) to roll out AI imaging and decision support tools to help analyse chest X-rays, support stroke diagnosis and manage conditions at home.

[The Food Standards Agency](#) uses an AI tool to help local authorities prioritise inspections by predicting which establishments "might be at a higher risk of non-compliance with food hygiene regulations". The government also announced in January 2023 that it would be [using AI to help "find and prevent more fraud across the public sector"](#).

Safety and ethics

The UK Government interprets AI safety to mean the [prevention and mitigation of harms \(whether accidental or deliberate\) from AI](#). Harms may arise from the ethical challenges that complex AI can present. These include the 'black box problem', where the inputs to and outputs from an AI system are known but humans cannot decipher – and the AI cannot explain – the process it went through to reach a particular conclusion, decision, or output. The AI's decision-making process, in other words, is not transparent nor accountable to humans.

Such decisions may also be susceptible to biases, such as 'embedded bias'. Embedded biases arise from relying on training data that reflects social and historical inequalities. These inequalities are then perpetuated in the outputs from the system. As IBM explains, "using [flawed training data can result in algorithms that repeatedly produce errors](#), unfair outcomes, or even amplify the bias inherent in the flawed data". In addition, there are concerns about the privacy and security implications of AI, including the use of sensitive data to train AI systems, as well as the ability of those systems to infer personal information.

The Ada Lovelace Institute, an independent data and AI research organisation, [groups AI risks and harms into four areas](#):

- accidental harms from AI systems failing, or acting in unanticipated ways, such as self-driving car crashes, or discrimination when sifting job applications.
- misuse harms from AI systems being used in malicious ways, such as bad actors generating misinformation using ‘generative’ AI applications such as ChatGPT and Midjourney.
- supply chain harms from the processes and inputs used to develop AI, such as poor labour practices, environmental impacts, and the inappropriate use of personal data or protected intellectual property.
- structural harms from AI systems altering social, political, and economic systems, such as the creation of unequal power dynamics (for example through market concentration or inequitable access to AI systems), or the aggregate effect of misinformation on democratic institutions.

As well as misuse risks, the UK Government (and others) has considered “[loss of control risks](#)” in which human oversight and control over a highly advanced, autonomous AI system is lost, leaving it free to take harmful actions. There is an ongoing, contentious debate about such ‘existential risks’ and their likelihood. The Ada Lovelace Institute has emphasised that we [should not lose sight of the harms](#) that can arise from existing (rather than futuristic) AI systems:

risks from AI do not simply grow in proportion with the technical capabilities of a given system, but depend intimately on the contexts in which those systems are deployed. Existing systems with capabilities far behind the ‘frontier’ can produce unexpected harms when deployed in new contexts, particularly without prior testing or consultation with users and other affected parties.

Elsewhere, thousands of AI experts [signed an open letter in March 2023](#), calling for a pause on the development and training of generative AI systems that are more powerful than ChatGPT-4. According to the signatories, a pause would have allowed time for the development and implementation of “a set of shared safety protocols for advanced AI design and development that are rigorously audited and overseen by independent outside experts”. No such pause occurred.

In November 2023, the UK Government hosted the [AI Safety Summit](#); a two-day global event at Bletchley Park, Buckinghamshire. Its aims were to examine the risks of AI, particularly in relation to ‘frontier AI’, and discuss how those risks could be mitigated by the international community. Frontier AI is [defined by the UK Government](#) (PDF) as “highly capable general-purpose AI models that can perform a wide variety of tasks and match or exceed the capabilities present in today’s most advanced models”.

Briefings published by the House of Commons and House of Lords Libraries, and the Parliamentary Office of Science and Technology on AI can all be found at: [Artificial intelligence - House of Commons Library \(parliament.uk\)](#)

1 Overviews of artificial intelligence

- Accenture, [What is Artificial Intelligence \(AI\)?](#), not dated
- Alan Turing Institute, [Data science and AI glossary](#), not dated
- Department for Science, Innovation and Technology, [A guide to using artificial intelligence in the public sector](#), June 2019
- Google Cloud, [What Is Artificial Intelligence \(AI\)?](#), not dated
- IBM, [Artificial Intelligence](#), December 2022
- McKinsey, [What is AI \(Artificial Intelligence\)?](#), April 2024
- Parliamentary Office of Science and Technology (POST) [Artificial Intelligence: An explainer](#), January 2024
- Parliamentary Office of Science and Technology (POST), [Artificial intelligence \(AI\) glossary](#), January 2024
- [Artificial intelligence \(AI\): recent trends and applications](#) / edited by S. Kanimozhi Suguna, M. Dhivya, and Sara Paiva. Publisher: Boca Raton, Florida ; London, England ; New York : CRC Press, 2021

1.1 Different types of AI

- Cloud Academy, [Types of AI explained](#), October 2023
- IBM, [Understanding the different types of artificial intelligence](#), October 2023

Machine learning and deep learning

- IBM, [What is a machine learning algorithm?](#), not dated
- NVIDIA, [The Difference Between AI, Machine Learning, and Deep Learning?](#), July 2016
- Ross Gruetzemacher, [The Power of Natural Language Processing](#), Harvard Business Review, April 2022
- Sara Brown, [Machine learning, explained](#), Massachusetts Institute of Technology, Sloan School of Management, April 2021

- University of California, Berkeley, [What Is Machine Learning \(ML\)?](#), June 2020

Foundation models and generative AI

- Elliot Jones, [Explainer: What is a foundation model?](#), [Ada Lovelace Institute](#), July 2023
- McKinsey, [What is ChatGPT, DALL-E, and generative AI?](#), April 2024
- Rishi Bommasani and Percy Liang, [Reflections on Foundation Models](#), Stanford University, October 2021
- Yoon Kim, Jacob Andreas and Dylan Hadfield-Menell, [Large Language Models](#) (PDF) Massachusetts Institute of Technology, Schwarzman College of Computing, November 2023

2 UK Government policy and regulation

2.1 UK Government publications

Strategies and roadmaps

- [AI Roadmap](#), January 2021
- [National AI Strategy](#), September 2021
- [National AI Strategy - AI Action Plan](#), July 2022
- [Regulators' strategic approaches to AI](#), May 2024

White papers

- [Establishing a pro-innovation approach to regulating AI](#) (consultation paper), July 2022
- [A pro-innovation approach to AI regulation \(white paper\)](#), March 2023 (updated August 2023)
- [A pro-innovation approach to AI regulation Government response to consultation](#) (PDF), February 2024

Risks and safety

- [Holistic AI: Risk Mitigation Roadmaps](#), September 2023
- [Capabilities and risks from frontier AI: discussion paper](#), October 2023
- [Introducing the AI Safety Institute](#), November 2023
- [Emerging processes for frontier AI safety](#), October 2023

AI Safety Summit, November 2023

Outputs from the summit, including the [Bletchley Declaration](#), can be found at: [AI Safety Summit 2023 - GOV.UK](#)

AI Standards Hub

Establishing an [AI Standards Hub](#), to develop technical AI standards, was one of the actions set out in the National AI Strategy. The hub is led by the Alan

Turing Institute and supported by the British Standards Institution and the National Physical Laboratory.

- An explanation of its work on AI standards is available at: [What are standards? - AI Standards Hub](#)
- The AI Standards database can be searched at: [AI Standards Search](#)
- Separately, the Department for Science, Innovation and Technology has published an [Introduction to AI assurance](#), February 2024

2.2 Government Office for Science (GO-Science)

- [Rapid Technology Assessment: Artificial Intelligence](#), March 2023
- [Future Risks of Frontier AI](#) (PDF), October 2023

2.3 Information Commissioner's Office (ICO)

- [Explaining decisions made with AI](#), October 2022
- [Guidance on AI and data protection](#), March 2023
- [Generative AI: eight questions that developers and users need to ask](#), April 2023

2.4 Digital Regulation Cooperation Forum

The [Digital Regulation Cooperation Forum](#) (DRCF) brings together the major UK regulators tasked with regulating digital services:

- the Competition and Markets Authority (CMA),
- the Financial Conduct Authority (FCA),
- the Information Commissioners Office (ICO), and
- the Office of Communications (Ofcom).
- [Digital Regulation Cooperation Forum: Plan of work for 2022 to 2023](#), April 2022
- [The benefits and harms of algorithms: a shared perspective from the four digital regulators](#), September 2022

- [Digital Regulation Cooperation Forum: Plan of work for 2023 to 2024](#) (PDF), April 2023
- Competition and Markets Authority, [AI Foundation Models: Initial report](#), September 2023
- Competition and Markets Authority, [AI Foundation Models: Update paper](#), 11 April 2024
- Financial Conduct Authority, [Our emerging regulatory approach to Big Tech and Artificial Intelligence](#), July 2023
- Ofcom, [What generative AI means for the communications sector](#), June 2023

2.5

Commentary on UK government's AI policy & regulation

- Ayesha Bhatti, [The UK's Agile, Sector-Specific Approach to AI Regulation Is Promising](#), Center for Data Innovation, February 2024
- Elliot Jones, [Keeping an eye on AI](#), Ada Lovelace Institute, July 2023
- [Expert comment: Oxford AI experts comment on the outcomes of the UK AI Safety Summit](#), University of Oxford, November 2023
- Information Commissioner's Office, [The Information Commissioner's response to the Government's AI White Paper](#) (PDF), 2023
- Kir Nuthi, [An Overview of the UK's New Approach to AI](#), Center for Data Innovation, April 2023
- Matt Davies, Michael Birtwistle, [Policy briefing: Regulating AI in the UK](#), Ada Lovelace Institute, July 2023
- Paul Shepley, Matthew Gill, [Artificial intelligence: how is the government approaching regulation?](#), Institute for Government, October 2023
- Harry Booth, [What We Know About the New U.K. Government's Approach to AI](#), Time Magazine, 15 July 2024

2.6

UK Parliament select committee reports

House of Commons

- Science and Technology Committee, [Algorithms in decision-making](#), Fourth Report of Session 2017–19, May 2018
- Science, Innovation and Technology Committee, [The governance of artificial intelligence: interim report](#), Ninth Report of Session 2022–23, 31 August 2023
- Culture, Media and Sport Committee, [Connected tech: AI and creative technology](#), Eleventh Report of Session 2022–23, 30 August 2023
- Science, Innovation and Technology Committee, [Governance of artificial intelligence \(AI\)](#), Third Report of Session 2023–24, 28 May 2024

House of Lords

- Select Committee on Artificial Intelligence, [AI in the UK: ready, willing and able?](#) (PDF), HL Paper 100, April 2018
- Liaison Committee, [AI in the UK: No Room for Complacency](#) (PDF), HL Paper 196, December 2020
- Select Committee on AI in Weapon Systems, [Proceed with Caution: Artificial Intelligence in Weapon Systems](#), HL Paper 16, December 2023
- Communications and Digital Select Committee, [Large language models and generative AI](#), HL Paper 54, February 2024

Government responses

- Science and Technology Committee, [Algorithms in decision-making: Government Response to the Committee's Fourth Report](#), July 2018
- Science, Innovation and Technology Committee, [The governance of artificial intelligence: interim report: Government response to the Committee's Ninth report](#) (PDF), First Special Report of Session 2023–24, HC 248, November 2023
- Culture, Media and Sport Committee, [Connected tech: AI and creative technology: Government Response to the Committee's Eleventh Report of Session 2022–23](#), Third Special Report of Session 2023–24, January 2024
- Department for Business, Energy & Industrial Strategy, [AI in the UK: ready, willing and able? - government response to the select committee report](#), June 2018

- Department for Digital, Culture, Media and Sport, [Government Response to the House of Lords Select Committee on Artificial Intelligence](#) (PDF), CP 390, February 2021
- Ministry of Defence, [The Government Response to the Report by the House of Lords AI in Weapon Systems Committee: ‘Proceed with Caution: Artificial Intelligence in Weapon Systems’](#) (PDF) (Session 2023–24 HL Paper 16), CP 1023, February 2024

1 European AI governance

European Union AI Act

The EU AI Act entered into force on 1 August 2024 with many of the provisions starting to apply two years later, on 2 August 2026. It is described by the [European Parliament](#) as “the world’s first comprehensive AI law”.

It is a ‘horizontal’ EU legislative instrument: the act applies to all AI systems placed on the market or used in the EU, regardless of where the provider, or the AI system, is based. It also covers AI systems outside of the EU (such as those in the UK) where the output produced by the AI system is used in the EU.

The act takes a risk-based approach to regulating AI, so that the requirements placed on providers, deployers and distributors vary depending on the type of AI and the purposes for which it will be used. Certain AI systems / practices are deemed harmful in the act and are banned, including untargeted scraping of facial images from the internet or CCTV footage to create facial recognition databases. Further information can be found in:

- [Regulation \(EU\) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence](#)
- European Parliamentary Research Service, [Artificial intelligence act](#), March 2024
- European Parliament, [EU AI Act: first regulation on artificial intelligence](#), June 2024

Council of Europe: first international treaty on artificial intelligence

In May 2024, the [Council of Europe adopted](#) the “first-ever international legally binding treaty aimed at ensuring the respect of human rights, the rule of law and democracy legal standards in the use of artificial intelligence (AI) systems”. It covers the use of AI by both the public and private sector and is technology-neutral: it does not seek to regulate the technology but rather sets out eight “fundamental principles” that the lifecycle of an AI system must

comply with, including “human dignity and individual autonomy” and “transparency and oversight”. It will be opened for signature on 5 September 2024. The UK is a member state of the Council of Europe.

Further information from the Council of Europe on the treaty is available at:

- [Council of Europe Framework Convention on Artificial Intelligence and Human Rights, Democracy and the Rule of Law](#)
- [The Framework Convention on Artificial Intelligence](#)
- [Artificial Intelligence: Ensuring respect for democracy, human rights and the rule of law](#)

3 Use of AI in different sectors

3.1 Business and industry

- All-Party Parliamentary Group on the Future of Work, [The New Frontier: Artificial intelligence at work](#), November 2021 (PDF)
- Andrew Evans and Anja Heimann, [AI activity In UK Businesses. An assessment of the scale of AI activity in UK businesses and scenarios for growth over the next twenty years](#) (PDF), A report by Capital Economics for the Department for Digital, Culture, Media, and Sport (DCMS) January 2022
- BSA, [Artificial intelligence in every sector](#), 2019 (PDF)
- Department for Digital, Culture, Media & Sport; Department for Science, Innovation & Technology, [AI activity in UK businesses: Executive Summary](#), January 2022
- Department for Science, Innovation & Technology; Department for Business, Energy & Industrial Strategy, [The potential impact of AI on UK employment and the demand for skills](#), October 2021
- European Commission, [Industrial applications of artificial intelligence and big data](#), 2020
- Eurostat, [Use of artificial intelligence in enterprises - Statistics Explained](#), December 2023
- McKinsey, [AI in government: Capturing the potential value](#), July 2022
- McKinsey, [Unlocking the promise of AI in industrials](#), December 2022
- Perspective Economics, [Artificial Intelligence Sector Study Research report for the Department for Science, Innovation & Technology \(DSIT\)](#), March 2023 (PDF)
- Xiang Hui, Oren Reshef, Luofeng Zhou, [The Short-Term Effects of Generative Artificial Intelligence on Employment: Evidence from an Online Labor Market](#), CESifo Working Papers, August 2023 (PDF)
- EY, [Data foundations and AI adoption in the UK private and third sectors](#) (PDF), July 2021

Use of AI in supply chains

- IBM, [AI is reshaping the supply chain](#), June 2017
- Maxime C. Cohen & Christopher S. Tang, [The Role of AI in Developing Resilient Supply Chains](#), Georgetown Journal of International Affairs, February 2024
- Remko Van Hoek and Mary Lacity, [How Global Companies Use AI to Prevent Supply Chain Disruptions](#), Harvard Business Review, November 2023

3.2

Creative industries

- Creative Industries Policy and Evidence Centre, [AI and the Creative Industries: The art in the artificial - Creative Industries Policy and Evidence Centre](#), December 2023
- David De Cremer, Nicola Morini Bianzino, and Ben Falk, [How Generative AI Could Disrupt Creative Work](#), Harvard Business Review, April 2023
- European Commission, [Study on Opportunities and Challenges of Artificial Intelligence \(AI\) Technologies for the Cultural and Creative Sectors](#), March 2022
- House of Commons Culture, Media and Sport Committee, [Connected tech: AI and creative technology](#), Eleventh Report of Session 2022–23, August 2023, HC 1643
- World Economic Forum, [How is AI impacting and shaping the creative industries?](#), February 2024

3.3

Defence and security

- Alice Saltini, [UK thinking on AI integration and interaction with nuclear command and control, force structure, and decision-making](#) (PDF), European Leadership Network, November 2023
- House of Commons Library, [Emerging and disruptive defence technologies, November 2023](#)
- Ministry of Defence, [Defence AI Playbook](#) (PDF), January 2024
- Mariarosaria Taddeo, Marta Ziosi, Andreas Tsamados, Luca Gilli and Shalini Kurapati, [Artificial Intelligence for National Security: The Predictability Problem](#), CETaS Research Reports, Centre for Emerging Technology and Security, Alan Turing Institute, September 2022

3.4

Education

- Cambridge Assessment, [Is education ready for artificial intelligence? Machine learning and EdTech](#), not dated
- Department for Education, [Generative artificial intelligence \(AI\) in education - GOV.UK](#), October 2023
- Department for Education, [The impact of AI on UK jobs and training - GOV.UK](#), November 2023
- House of Lords Library, [Educational technology: Digital innovation and AI in schools](#), November 2023
- Jisc, [Generative AI - a primer](#) (considers the impact of generative AI on Education), February 2024
- Jisc, [Artificial intelligence \(AI\) in tertiary education](#), September 2023
- Oxford University Press, [AI in education: where we are and what happens next](#), October 2023
- Parliamentary Office of Science and Technology, [Use of artificial intelligence in education delivery and assessment](#), January 2024
- UNESCO, [Artificial intelligence in education](#), 2023

3.5

Energy

- David G. Victor, [How artificial intelligence will affect the future of energy and climate](#), January 2019
- Energy Systems Catapult, [AI in Energy](#), not dated
- E.ON, [How Artificial Intelligence is accelerating the Energy Transition](#), May 2019
- International Energy Association, [Why AI and energy are the new power couple – Analysis](#), November 2023
- McKinsey, [An AI power play: Fueling the next wave of innovation in the energy sector](#), May 2022
- World Economic Forum, [Harnessing Artificial Intelligence to Accelerate the Energy Transition](#), September 2021

3.6

Financial services

- Alan Turing Institute, [AI in financial services](#), 2021
- Bank of England, [The AI Public-Private Forum: Final report](#), February 2022
- Chartered Insurance Institute, [AI in Claims: the good, the unknown, and the opportunity](#), November 2023
- International Monetary Fund, [Generative Artificial Intelligence in Finance: Risk Considerations](#), August 2023
- Oliver Wyman, [The AI Revolution In Banking](#), 2022
- OECD, [OECD Business and Finance Outlook 2021 : AI in Business and Finance](#), September 2021
- UK Finance, [How Artificial Intelligence is Transforming the Financial Services Industry](#), November 2023
- UK Finance, [The impact of AI in financial services: opportunities, risks and policy considerations](#), November 2023

3.7

Healthcare

- IBM, [The benefits of AI in healthcare](#), July 2023
- Imperial College, London, [AI for healthcare](#), not dated
- National Institute for Health and Care Research, [AI in healthcare - 10 promising interventions](#), July 2023
- Parliamentary Office of Science and Technology, [AI and healthcare](#), January 2021
- NHS, [The NHS AI Lab - NHS Transformation Directorate](#), not dated
- Thomson Reuters Institute, [Understanding the advantages and risks of AI usage in healthcare](#), September 2023
- The Health Foundation, [AI in health care: hope or hype?](#), March 2023
- World Health Organization, [WHO issues first global report on Artificial Intelligence \(AI\) in health and six guiding principles for its design and use](#), June 2021

3.8

Retail

- Retail Insight Network, [How is AI used in retail in 2024?](#), August 2023
- US National Retail Federation, [How artificial intelligence will change retail](#), June 2023
- World Economic Forum, [Here's how artificial intelligence can benefit the retail sector](#), January 2023
- Women's Wear Daily, [Harnessing Generative AI: The New Frontier for Fashion Retailers](#), February 2024

3.9

Travel and transport

- European Parliamentary Research Service, [Artificial intelligence in transport. Current and future developments, opportunities and challenges](#), March 2019 (PDF)
- Department for Science, Innovation & Technology, [AI Barometer Part 4 - Transport and logistics](#), December 2021
- Innovate UK, [The road to innovation: 4 key opportunities for AI in transport](#), March 2024
- techUK, [Artificial Intelligence as a vehicle for Transport innovation and economic growth](#), April 2023

4 Safety and ethics

4.1 General

- Centre for Data Ethics and Innovation, [Enabling responsible access to demographic data to make AI systems fairer](#), June 2023
- Council of Europe, [Common ethical challenges in AI](#), in The Impact of Artificial Intelligence on the doctor-patient relationship, December 2021
- David Leslie, [Understanding artificial intelligence ethics and safety: A guide for the responsible design and implementation of AI systems in the public sector](#), The Alan Turing Institute, 2019
- Department for Science, Innovation and Technology, [Understanding artificial intelligence ethics and safety](#), June 2019
- Department for Science, Innovation and Technology, [Introducing the AI Safety Institute](#), 17 January 2024
- GCHQ, [Pioneering a New National Security, The Ethics of Artificial Intelligence](#), not dated
- HM Government, [Generative AI framework for HM Government](#) (PDF), not dated
- National Cyber Security Centre, [Guidelines for secure AI system development](#), November 2023
- Jerome Beranger, [Societal Responsibility of Artificial Intelligence: Towards an Ethical and Eco-responsible AI](#), Publisher: Hoboken, New Jersey : John Wiley & Sons, Inc. June 2021

4.2 The ‘black box’ problem

- [Examining the Black Box](#), Ada Lovelace Institute, April 2020
- Neil Savage, [Breaking into the black box of artificial intelligence](#), Nature, 29 March 2022
- Stephen Bush, [Beware the rise of the black box algorithm](#), Financial Times, 20 September 2022

- Tom Cassauwers, [Opening the 'black box' of artificial intelligence](#), Horizon, the EU research and innovation magazine, 1 December 2020

4.3

Bias in AI

- Centre for Data Ethics and Innovation, [Review into bias in algorithmic decision-making](#) (PDF), November 2020
- IBM, [Shedding light on AI bias with real world examples](#), October 2023
- Jonny Freeman and Carolyn Ashurst, [Data protection, AI, and fairness](#), [The Alan Turing Institute](#), not dated
- McKinsey, [Tackling bias in artificial intelligence \(and in humans\)](#), June 2019
- Sray Agarwal and Shashin Mishra, [Responsible AI - Implementing Ethical and Unbiased Algorithms - Microsoft Industry Blogs](#), November 2021

4.4

Misuse and harms

- Ardi Janjeva and colleagues, [The Rapid Rise of Generative AI: Assessing risks to safety and security](#), Centre for Emerging Technology and Security, Alan Turing Institute, December 2023
- Heather Frase and Owen Daniels, [Understanding AI Harms: An Overview | Center for Security and Emerging Technology](#), Center for Security and Emerging Technology, Georgetown University, August 2023
- Julia Smakman, Matt Davies and Michael Birtwistle, [Mission critical: Lessons from relevant sectors for AI safety](#), Ada Lovelace Institute, October 2023

The House of Commons Library is a research and information service based in the UK Parliament. Our impartial analysis, statistical research and resources help MPs and their staff scrutinise legislation, develop policy, and support constituents.

Our published material is available to everyone on commonslibrary.parliament.uk.

Get our latest research delivered straight to your inbox. Subscribe at commonslibrary.parliament.uk/subscribe or scan the code below:



 commonslibrary.parliament.uk

 [@commonslibrary](https://twitter.com/commonslibrary)