

The logo for Wiggin, featuring the word "wiggin" in a white, lowercase, sans-serif font. The background of the entire page is a vibrant red. In the top left corner, there is a pattern of white dots that form a series of overlapping, wavy lines, creating a sense of depth and movement. On the right side of the page, there is a large, stylized profile of a human head, rendered in black and white vertical stripes, which appears to be looking towards the left.

Generative AI considerations for business

UK Q1 2024 update

Recent progress in transformer-based self-attention models has changed the AI game beyond recognition.

This new class of models, known informally as 'generative AI', has enabled the creation of completely new content and data never seen before. By learning patterns and relationships from large datasets, generative AI uses algorithms and models to generate outputs such as images, audio, text, video and even entire virtual environments.

The rapid advancement of AI, particularly generative AI, has opened new opportunities and challenges for businesses across various sectors. Industry announcements continue apace covering new model launches, investments, partnerships and innovations from AI vendors and their partners. Industries continue to grapple with how to most effectively capture value and defend against threats from the AI revolution - whether it be how to use their data to train AI models to produce better outputs, how to use AI outputs as more efficient inputs for their business, and how existing business functions (including the humans employed by them) can be expanded, optimised and/or replaced with AI systems.

As AI technologies and their use cases advance, so do the regulations and guidelines governing their use. Businesses must stay informed about the evolving legal landscape to ensure compliance and to avoid potential fines or legal disputes, as well as to maximise the opportunities presented by utilising AI. Effective regulation of AI by the UK Government and governments around the world collectively may be the greatest regulatory challenge of our generation.

In this document, we explore the legal complexities of regulating AI and key points that businesses operating in the UK should be aware of when developing, deploying and exploiting new AI technologies.

Authors



Marcus Bagnall
Partner



Mark Deem
Partner



Ben Towell
Regulatory Advisor



Cecilia Lovell
Regulatory Advisor



Recent regulatory approach in the UK

A pro-innovation approach to AI regulation

As we've previously reported, the UK Government in 2023 issued a [call for comments](#) on its AI White Paper 'A pro-innovation approach to AI regulation', proposing a non-statutory and cross-sectoral framework based on five core principles: (i) safety, security and robustness, (ii) transparency and explainability, (iii) fairness, (iv) accountability and governance, and (v) contestability and redress. The aim of the rules is to allow for a flexible and adaptive approach to respond to technological progress where regulators can implement these principles to the specific contexts in which AI is used.

On 6 February 2024, the Government released its long awaited [response](#) to its AI White Paper, acknowledging the strong support received from stakeholders and regulators who are already taking action in line with the Government's proposed approach to AI regulation, including: (1) the Competition and Markets Authority ("CMA") [initial review of foundation models](#), focused on the consumer principles that should best guide development in this area; (2) the Information Commissioner's Office ("ICO") updated [guidance](#) on data protection principles and how they apply to AI systems; and (3) the Office of Communications ("OFCOM") strategic approach to AI in its [Plan of Work for 2024/2025](#). The Government will continue to cooperate with regulators in these areas to ensure AI driven markets remain fair and competitive, and has asked regulators to publish their strategic approach by 30 April 2024.

Further reports on the regulators' proposed approach is available [here](#), [here](#) and [here](#).

The Communications and Digital Committee in the House of Lords recently released its [report](#) on focusing on large language models ("LLMs") and generative AI, which we reported on [here](#). The report outlined the potential advantages of LLMs in terms of economic and scientific advancements, but also underscored the importance of addressing associated risks, including threats to public safety, societal values, open market competition, and overall economic competitiveness. The report advocated for a balanced regulatory approach that not only prioritises AI safety but also fosters commercial opportunities. It expressed concerns about the concentration of market power and the potential for regulatory capture, emphasising the need for open competition and transparency. As discussed below, it also highlighted issues over copyright, echoing concerns regarding technology companies allegedly using copyrighted material without permission, and calling for decisive Government action. The Committee stipulated a two-month timeframe for the government to respond to the report.

While applying cross-sector principles in a context-based approach can be beneficial in the long-run, establishing regulatory frameworks is challenging as Government and regulatory bodies contend with the rapid AI technology advancement in balancing the needs of encouraging innovation while maintaining certainty and protecting against harms.



AI assurance and governance guidance

In February 2024, the Department for Science, Innovation and Technology (“**DSIT**”) published its [‘Introduction to AI assurance’](#) in the first of its series of guidance to help organisations and regulators upskill on topics around AI assurance and governance (the “**Guidance**”). This Guidance aims to establish AI assurance techniques and standards to support industry and regulators in understanding “how” to implement the UK Government’s five core principles to AI regulation as set out in its AI White Paper by providing agreed-upon processes, metrics and frameworks to build and deploy responsible AI systems.

The Guidance also sets out key immediate actions for organisations and business, including to raise AI assurance understanding levels by considering the [training workbooks](#) published by the [Alan Turing Institute](#) and the training platform provided by the [UK AI Standards Hub](#); and implementing effective internal AI governance processes.

DSIT plans to publish additional sector-specific guidance to provide more detail about AI assurance in particular contexts.

In a further effort by regulators to provide guidance while fostering AI innovation, the Digital Regulation Cooperation Forum (**DRCF**) comprising the CMA, Ofcom, ICO and Financial Conduct Authority (FCA), has recently launched the [AI and Digital Hub](#) for innovators to receive informal advice on regulatory requirements from DRCF member regulators. The AI and Digital Hub is intended to help innovators navigate overlapping regulatory landscapes to safely bringing new products to the UK market. Providers with an ‘innovative’ AI or digital product, service or business model intended for consumers in the UK may apply to the hub. Outcomes of queries will be shared as case studies to aid a broader range of innovators, with consideration for confidentiality.

Assessing business risk

As with any technology investment decision, business must carefully assess whether the AI tool being considered appropriately addresses user needs and understand what data will be required, used and stored in connection with its development and ongoing use. Businesses should also assess their readiness for implementation, considering factors such as the quality and scale of existing data, employee skills and the IT infrastructure in place. There may be little commercial sense investing in training an AI solution when their existing data is limited or of poor quality where a commercial off-the-shelf solution may be more appropriate.

Businesses looking to adopt AI solutions may choose to start with small-scale pilot projects to test feasibility, effectiveness, and appetite. For example, AT&T’s generative AI tool “[Ask AT&T](#)” was initially developed to help increase productivity levels for its coders and software developers. The tool’s uses have grown significantly since then, with the tool now being used to upskill customer care representatives, provide answers to employee HR questions, and help translate documents.

While AI solutions can automate complex tasks, they may also make errors, ‘hallucinate’ or produce unpredictable results. AI systems need to be accompanied by [appropriate processes for monitoring performance](#) and ensuring that their AI systems meet any relevant regulatory requirements, ethical guidelines, and organisational values. This is particularly relevant as businesses will be held accountable for transparency and explainability in how AI systems produce their outputs. For example, an AI-driven decision support software should be accompanied by documentation explaining how the AI system produced the relevant outputs to allow the user responsible for making the decision to independently verify the outputs from the AI system.

Intellectual Property

As we summarised in our [response](#) to the UK Intellectual Property Office's [call for views on AI](#), a widely recognised legal concern with training and using generative AI models is the potential for intellectual property rights infringement. As AI systems are developed, trained, generate content and make decisions, questions arise about the use of media outputs as inputs to train models, the ownership of trained models and of AI-generated outputs from those trained models (including as media inputs), and the implications for intellectual property rightsholders and AI vendors.

Generative AI models trained on large datasets will invariably include copyrighted material. There have already been a number of high-profile cases in key jurisdictions launched by rightsholders against AI developers for infringement in connection with the use of their data, often large volumes of data scraped or mined from data available over the internet, to train AI models and outputs generated by AI models.

The interests of rightsholders and AI vendors diverge significantly. Rightsholders want protection from at-scale copyright infringement and better transparency for how their data manifests in trained AI models and their outputs, while AI vendors need large, quality datasets to produce more effective AI models and want room to innovate. At the same time, such training and the production of AI outputs from those trained models is considered to undermine the commercial basis for rightsholders producing their content in the first place. There is also increasing need for certainty for rightsholders who use AI model outputs as inputs to producing content.

On 11 January 2024, the Government [responded](#) to the Culture Media and Sport Committee's 30 August 2023 report on [Connected Tech: AI and creative technology](#), supporting the Committee's position that "reproduction of copyright-protected works by AI will infringe copyright unless permitted under licence or an exception" and confirming that it will not proceed with proposals to broaden copyright exceptions allowing text and data mining for AI model training. At the same time, the Government has confirmed in its response to the AI White Paper that the Intellectual Property Office has been unsuccessful in brokering a voluntary code of practice between AI developers and rightsholders regarding using copyright materials to train AI models. While the road ahead is unclear, DSIT has expressed its intention to continue engaging with stakeholders to agree an approach that will allow both the AI and creative sectors to grow together in partnership. Read further [here](#) for our more detailed report.

In the meantime, data sharing and data licensing models are evolving to expressly account for AI training and use, such as licences granted by rightsholders to AI developers or licensee authorisations for specific purposes that cover making licensed data available to train AI models.¹ AI model training also presents a data commercialisation opportunity for rightsholders and data owners that can create new revenue streams for rightsholders who can shape the further beneficial development of AI models.² These arrangements offer the potential for AI vendors to differentiate based on training data, which if these data sharing agreements are exclusive could make them difficult to compete with.

¹ See, for example, Microsoft's proposed AI data sharing agreement template to use data to train an AI model to be made available on an open source basis: <https://query.prod.cms.rt.microsoft.com/cms/api/am/binary/RE4Rfg>.

² See, for example, Reddit entering into an AI content licensing deal with Google - <https://www.reuters.com/technology/reddit-ai-content-licensing-deal-with-google-sources-say-2024-02-22/>.

Data Protection

The use of AI can involve collecting, processing, and storing large amounts of data, which raises concerns about data privacy and security. Ensuring that AI systems comply with relevant regulations and maintain the confidentiality and protection of personal data and sensitive information is crucial for mitigating these risks.

The ICO has launched a series of [public consultations](#) on generative AI, focusing on how aspects of data protection law should apply to the development and use of AI technologies, as well as how AI developers could subsequently establish a lawful basis for AI development as required under the UK GDPR and DPA 2018.

The [first consultation](#) considers the lawful basis for web scraping data used to train generative AI models. It proposes that the *legitimate interest* lawful basis for processing data under the GDPR may be available for AI developers where they can satisfy the requisite three-part test: (i) that the purpose of the processing is *legitimate*, (ii) that the processing is *necessary* for that purpose, and (iii) that individuals' interests do not override the interest being pursued. It is possible this test could be satisfied in the case of, for example, web-scraped publicly available data being used to train a defined purpose AI model where there is effectively no alternative large dataset source and steps have been taken to mitigate infringement of data subjects' rights. Further information on this consultation is available in our previous article [here](#).

The [second consultation](#) considers purpose limitation in the generative AI lifecycle. The purpose limitation principle requires organisations to be clear and open about why they are processing personal data and that such data is used for its intended purpose. With regards to training generative AI models, organisations may process different types of personal data throughout its lifecycle, for example: the purpose

of training a core model will require training data and test data, while the purpose of adapting the core model may require fine-tuning data from a third-party developing its own application. In this consultation, the ICO proposes how purpose limitation should apply in each stage of the training process of generative AI models and how organisations should process data in compliance with data protection principles. The expectations of the ICO at this stage are that defined purposes will be necessary in order to understand how the training and deployment of an AI solution will comply with data protection laws. Further information on this consultation is available in our previous article [here](#).

The [third consultation](#) considers accuracy in relation to the outputs of generative AI models, and the impact that the accuracy of training data has on the output. The call for evidence emphasises accuracy as a principle of data protection law requiring organisations to ensure that the personal data they process is accurate and up to date. Users relying on generative AI models with factually wrong information can cause reputational damage, financial harms and spread of misinformation. The ICO recognises that the specific purpose for which a generative AI model will be used is what determines whether the outputs need to be accurate. The key is for there to be clear communication between the developers and end-users to ensure that the application is properly used and inform the end-user of the level of accuracy, such as providing clear information about the statistical accuracy of the application, labelling the outputs as '*generated by AI*' or '*not factually accurate*', amongst others. Further information on this consultation is available in our previous article [here](#).

Further consultations will be published in the coming months on purpose limitation, accuracy and the rights of data subjects.

Bias and Discrimination

AI models learn from the data they are trained on.

What if the training data used contains inherent biases? Bias arises from the presence in data sets, or the absence in data sets, and through the manner in which data sets are collated. Bias can also be introduced through algorithms and the manner in which it latches onto a correlation and amplifies it, resulting in inherent bias which we cannot see. AI models trained on data containing inherent biases may therefore unfortunately contribute to perpetuating and reinforcing these biases in their outputs, which if left unchecked could lead to further unfairness and discrimination. This is particularly the case where AI model outputs are used in processes that directly or replace human decision-making processes and discretions, such as in job application processes, enforcing rules or selection procedures. The use of AI models in these situations [could hold legal implications](#) for AI developers or those that adopt AI solutions where harm or discrimination results from decisions due to any inherent biases.

It is arguable that biased AI outputs resulting from biased training data simply reflect systemic bias within the existing human process it looks to replace, and therefore the AI is not “at fault” per se nor is its development the “cause” of any biased AI outputs. However, there is arguably a policy distinction between a biased AI model that can only produce outputs dictated by its (biased) training data compared to a “human” process that we can (and must) expect to adjust following the identification of any bias to mitigate the risk.

On the other hand, generative AI can be interrogated, which may allow biases to be more quickly identified than it can be from within human discretion. Such probing may actually be able to reveal human biases within a data set that previously had gone unnoticed.

In terms of counteracting inherent biases, the UK government implemented the [Artificial Intelligence Impact Assessment](#) (“AIIA”) as a valuable tool to assess the potential impacts of AI systems. The AIIA, amongst other things, puts in place controls to assess potential bias and fairness impacts in relation to an AI system throughout all stages of the system lifecycle. Controls such as testing are key to identifying any adverse impacts before AI models are released. Audit trails then need to be used to show that it is performing in the manner anticipated.

The DSIT, in collaboration with various organisations, released [guidance](#) on responsible AI use in HR and recruitment. The document addresses risks such as bias and discrimination, providing examples like discriminatory job review software and biased chatbots, and emphasising the need for AIIA. We will most likely see competent authorities releasing guidance on the use of AI models as these are increasingly used across various industries.

Further information on the guidance can be found [here](#).



AI regulatory developments in the European Union

The EU AI Act

On 13 March 2024, the European Parliament plenary session formally adopted at first reading the [EU AI Act](#).³ The text reveals further details of the extensive changes the co-legislators have made to the Commission's original proposal published in 2021, including measures to protect fundamental rights, democracy, the rule of law, and environmental sustainability from high-risk AI.

The text establishes bans on specific AI applications that threaten citizens' rights, including biometric categorisation systems, untargeted scraping of facial images for recognition databases, emotion recognition in workplaces and schools, social scoring, predictive policing based solely on profiling, and AI that manipulates human behaviour or exploits vulnerabilities. Law enforcement exemptions are strictly regulated, permitting the use of biometric identification systems only in limited and defined situations, with strict safeguards and judicial or administrative authorisation.

Clear obligations have been agreed upon for other high-risk AI systems, covering critical areas such as infrastructure, education, employment, essential services, law enforcement, migration and border management, justice, and democratic processes. Citizens are granted the right to launch complaints about AI systems affecting their rights, with transparency requirements for general-purpose AI systems, including compliance with EU copyright law during training. To support innovation and provide opportunities for SMEs and start-ups, regulatory sandboxes and real-world testing will be established at the national level.

Accompanying the AI act, the European Commission is expected to provide guidelines to assist with practical implementation with illustrative examples, a template for surveillance and data collection, an annual reporting structure and specific technical documentation guidelines to assist with compliance.

Once the Act comes into force, the provisions in respect of prohibited AI systems will apply within six months, the provisions on GPAI will apply within 12 months and the remaining provisions will apply within 24 months, save in respect of high-risk AI under Annex II which will apply within 36 months. GPAI models already on the market when the Act comes into force will be given a two-year grace period.

Our further analysis of the text is available [here](#). The Parliament is expected to approve the compromise text in April 2024, and would then need to be formally endorsed by the Council of the European Union.

³ Artificial Intelligence Act: MEPs adopt landmark law - [Artificial Intelligence Act: MEPs adopt landmark law | News | European Parliament \(european-council.europa.eu\)](#)

Generative AI and competition considerations


Competition regulators have expressed concerns regarding the growing influence of a handful of large technology and infrastructure firms across the AI value chain on competition and consumer protection outcomes. It is [expected](#) that major firms who already hold significant market power in key digital markets (hardware/chips, software and ecosystems) will exploit strong positions in both the development of AI models, including the supply of critical inputs like data, and in the deployment of AI models through key access points and routes to market.

In the UK, the CMA has published an updated report as part of its review into AI foundation models, which we have reported on [here](#). In this report, the CMA outlines a set of principles to promote competition and consumer outcomes in the use of AI models to ensure access, diversity, choice, fair dealing, transparency and accountability. These are designed to complement DSIT's principles set out in its [AI White Paper](#) and the ICO's [guidance](#) on AI and data protection. The CMA is vigilant against the possibility that incumbent firms may try to use partnerships and investments to quash competition. To this end, the CMA signposts actions it intends to take to ensure fair, open and effective competition in the use of AI, including using its market investigatory and merger control powers regarding those digital activities that are critical access points or routes to market for AI model deployment, such as mobile ecosystems, search and productivity software.

In the EU, the European Commission has initiated [calls for contributions](#) on this subject and has specifically requested information from major digital market players, which we previously [reported](#) on. The EC's objective is to gather insights on the competitive landscape within these domains and explore how competition laws can preserve market competitiveness. The EC is specifically examining agreements between digital market players and generative AI developers, investigating their impact on market dynamics. It acknowledges the rapid growth and significant impact of generative AI technologies on businesses, prompting the need for a forward-looking analysis of potential regulatory issues and any consequential adaptations to EU legal concepts.

The EU approach has traditionally emphasised the importance of effective enforcement of competition rules to sustain competition in the Single Market. Generative AI systems that produce synthetic content in response to user prompts are recognised as disruptive technologies with substantial potential. The EU has proactively taken steps to address challenges posed by generative AI with initiatives such as the EU AI Act noted above, alongside other initiatives such as the Digital Markets Act, we have reported on [here](#).

These actions from competition regulators will have a direct impact on how AI solutions are offered to the market, what sorts of partnerships AI vendors and incumbent digital ecosystem players may enter into, and how AI models may be developed and launched in the future.



wiggin

Find out more

[wiggin.co.uk/expertise/ai-data-driven-technology/](https://www.wiggin.co.uk/expertise/ai-data-driven-technology/)