Google

Executive Summary

Google welcomes the opportunity to respond to the UK Government's consultation on AI and copyright. Our response outlines how targeted, pro-growth and pro-innovation reforms to the copyright framework will facilitate the UK's ambition to become a leader for AI investment and innovation.

The UK is rightly positioning itself to be at the forefront of AI technologies to reap the societal and economic benefits of these developments, estimated at up to £400bn by 2030[1]. We welcome the UK's ambition to seize this opportunity and become a world leader on AI and the recent AI Opportunities Action Plan is a positive statement of intent. But for AI companies of all sizes to flourish, a truly competitive copyright framework will be crucial.

We are clear in our view that the training of AI models is a non-expressive use of open web content, which boosts economic growth, fosters scientific advancement, and enables the creation of valuable new works. AI platforms and services will supplement and support creativity, not replace it.

To ensure the UK delivers on its ambition to be an 'AI maker', we believe the Government should deliver an internationally competitive copyright regime that:

- Maintains the current non-commercial research exception
- Enables TDM for any purpose with rights reservation for rightsholders
- Provides the UK with a competitive advantage by introducing a commercial research exception.

While option 2 in the consultation is the most competitive option and would support the Government's ambitions to be an AI leader, we believe that rights holders should have choice and control, and acknowledge that option 3 would deliver on this aim. As we lay out in detail in our response, appropriate technical solutions exist (e.g., in the form of an existing and effective opt out mechanism) and are working to facilitate the detail of that proposal. However, we are concerned that excessive transparency requirements referenced in option 3 could hinder AI development and impact the UK's competitiveness in this space.

Our preferred approach is therefore an amended Option 3, which allows TDM for any purpose with rights reservation for rights holders, *and* secures a competitive advantage by introducing a commercial research exception.

There is a clear need for a balanced approach that encourages AI innovation while respecting the rights of content creators. We hope for reforms that provide legal clarity, support technological advancement, and foster collaboration between AI developers and rights holders, to the benefit of the UK.

---

[1] Google UK Economic Impact report

The UK can be an AI maker, not an AI taker

The AI Opportunities Action Plan sets the UK up for success

The UK has stated its ambition to be an attractive place for AI innovation and investment. Having a competitive copyright framework will facilitate AI companies of all sizes to thrive, and establish the foundations for realising the UK's vision.

We welcome the ambition of the UK AI Opportunity Action Plan to proactively shape AI's potential and to be an "AI maker, not an AI taker". As the Plan rightly notes, delivery of this vision will require bold and decisive action.

A copyright framework that supports innovation and creativity is one strong predictor of whether a country will be a leader on AI. It is no coincidence that AI innovation - and indeed tech innovation overall - has taken root in countries that have a long history of balanced copyright frameworks, including fair use or a specific TDM exception covering commercial uses, such as Japan and Singapore.

Balanced text and data mining exceptions within copyright law allow researchers, innovators and creators to use copyright-protected material under certain circumstances without permission from the copyright holder. Exceptions are necessary to facilitate new expression, follow-on creativity, and innovation, and specifically to understand the fundamental nature of ideas and creativity in order to develop new societally-beneficial tools and resources.

AI depends on workable TDM exceptions because responsible development of AI applications requires use of a broad and diverse range of data. Without clear legal protections for this type of use, the UK will struggle to achieve its stated ambitions in becoming an attractive place for AI investment and innovation, as training AI models in the UK will be more difficult for companies of all sizes.

We are pleased that the Government's response to the Action Plan notes that they will act to ensure that the UK has a competitive copyright regime that supports both the AI sector and creative industries. For this ambition to be realised, clear direction on TDM exceptions is vital.

The Prime Minister has stated that AI will be the centre piece of the Industrial Strategy, and it is also at the heart of the Government's 'Plan for Change' - to deliver a much-needed boost to growth and productivity in the UK. In addition to the concerns acknowledged in the consultation - that a lack of clarity could stunt AI innovation and adoption in the UK - an unclear approach will also have an impact on long term investment decisions.

Maximising AI's potential is critical to generating widespread opportunity

AI holds transformative economic and social potential. All industries - including the creative sectors - stand to benefit from advances in AI.

All industries stand to benefit from advances in AI. The AI copyright frameworks we adopt will affect and shape not only the creative industries - typically associated with copyright considerations - but virtually every sector of economic and scientific activity that relies on innovations in AI, including both generative and non-generative capabilities. The recognition in the consultation[2] that there are applications of AI that are unlikely to have as much impact on the creative industries, but are also impacted by the frameworks decided upon here, is important.

AI is already enabling major breakthroughs - from translating languages to mitigating climate change and understanding diseases - which can deliver significant benefits for people in the UK and around the world.

---

[2] "*But AI and data mining have many other applications, whose outputs are unlikely to affect the creative industries to the same degree. These include research applications – for example, the use of AI to identify candidates for drugs or treatment pathways. They also include non-generative commercial applications such as image and music recognition – technologies which may be employed by right holders to create commercial products or to prevent online piracy.*" Copyright and AI: Consultation

AI also has the potential to unlock significant economic benefits, driving prosperity through greater opportunity. Research conducted by Public First estimates that AI has the power to save over 700,000 hours a year in administrative work for GPs and teachers. This could help offset growing cost pressures in areas such as health and education, and free up over £8 billion in greater public sector productivity for other uses. This research also found that AI could help over 1 million people with disabilities at work, boosting the UK economy by over £30bn per year. [3] As the AI Opportunities Action Plan notes, AI could be the single biggest lever to deliver the Government's five missions, particularly economic growth[4].

In medicine, Med-Gemini, a family of Gemini models fine-tuned for multimodal medical domain applications, presents substantial potential. This family of models builds upon the Gemini large language model by fine-tuning on de-identified medical data, and achieves a 91.1% accuracy on benchmarking against the popular MedQA benchmark. These models can interpret complex 3D scans, answer clinical questions, and generate state-of-the-art radiology reports - and even calculate risk predictions. Med-Gemini demonstrates that powerful multimodal capabilities, driven by generative AI, have the potential to assist clinician, researcher and patient workflows. These are early findings but point to exciting capabilities on the horizon for healthcare applications from generative AI[5].

We recently introduced AI co-scientist, a multi-agent AI system built on Gemini 2.0 as a virtual scientific collaborator to help scientists generate novel hypotheses and research proposals, and to accelerate scientific and biomedical discoveries. Initial findings include multiple wet-lab validations of new insights in areas like drug repurposing for acute myeloid leukaemia, liver fibrosis and antimicrobial resistance.[6]

Beyond societal challenges and economic growth, AI has exciting and promising applications in the creative and media industries. AI has the potential to open up new opportunities for artists, journalists, and creators of all kinds - and we are already seeing creators exploring new approaches to the creative writing process, to music and visual art productions, to textile design development, and more. Newsrooms and journalists are also integrating AI into their work processes, with nearly three quarters of news organisations believing generative AI presents new opportunities for journalism[7]. Our research tool, Pinpoint, helps journalists and academics analyse and explore large collections of documents, and is already being put to use in award-winning investigative reporting[8].

We are committed to building tools that increase access to information and create new and expanded economic and creative opportunities for artists, small businesses, and creators of all kinds. To do this, we are working closely with the creative community to put these tools in the hands of creators and to tackle new challenges as they emerge. We see AI as a complement to, and not a substitute for, human creativity; YouTube creators have embraced AI to streamline and boost their creative processes, with more than 1.7 billion views of videos related to AI tools on YouTube in 2023 alone.[9]

From all of these examples, it is clear that the term "AI" describes a more wide-ranging and diverse set of technologies than the generative AI applications that have captured public imagination. It is important to note the sheer breadth of possible AI applications across sectors as diverse as healthcare, media and entertainment, retail, e-commerce, logistics, banking, finance and IT. All are already integrating AI-based solutions in a myriad different ways into their products and services.

---

[3] Public First, 'Google's Impact in the UK 2023', April 2024
[4] AI Opportunities Action Plan, January 2025
[5] Google Research, Med-Gemini, May 2024
[6] Google Research, Accelerating scientific breakthroughs with an AI co-scientist, February 2025
[7] LSE, 'Nearly three quarters of news organisations believe generative AI presents new opportunities for journalism', September 2023
[8] Google Pinpoint, 'New funding, trainings and tools to help journalists', February 2024
[9] YouTube, 'Our principles for partnering with the music industry on AI technology', August 2023

Working in partnership to develop and explore AI capabilities

YouTube & Google are creating new opportunities for the creative industries, lowering barriers to creative production, and using AI as a complement to human creativity - not as a substitute.

We support the emphasis on partnerships throughout the consultation and agree that AI can enhance creative productivity in a range of ways, just as creative content helps drive AI development. Google greatly values its relationships with its diverse range of partners; our partnerships stretch beyond the rightsholders and AI firms, to many different industries and sectors.

YouTube & Google are helping open a world of opportunities for the creative industries. Through increased investments in technology and a responsible approach in AI, we're enabling artists & creators to push the bounds of creative expression and lowering barriers to creative production. We view AI as a complement to and not a substitute for human creativity, and creators are also finding value in using AI tools to support their creative businesses. According to a Deloitte study, 62% of creators plan to use generative AI tools over the next year for a variety of purposes, from idea generation to editing to workflow management[10].

YouTube's long standing, deep partnership and collaboration with the music industry has been a catalyst for innovation, enabling us to create products, features, and experiences that inspire originality and connect fans worldwide. Central to this shared success is our commitment to protecting artists and the integrity of their work:

- We have developed a set of AI Music principles and a Music AI Incubator in partnership with the music industry, guiding our collective efforts towards responsible collaboration.
- Since then, we have been exploring the possibilities of how AI can empower creativity alongside artists, songwriters, producers and our partners while also identifying its challenges. Through innovative experiments such as DreamTrack in Shorts and Music AI Tools, we're committed to amplifying opportunities for artists and the industry at large, deepening fan engagement, and ultimately enriching the entire music ecosystem.
- Looking further into the future, we are actively developing new technology that will enable people from a variety of industries—from creators and actors to musicians and athletes—to detect and manage AI-generated content showing their faces on YouTube. [11]

We are committed to supporting a healthy ecosystem that drives value for consumers, creators, and news publishers. Open access to information is core to our mission, and we will be thoughtful about how we can continue to support a healthy open web and the creators of the content:

- We have continued to build ways for publishers and content creators to monetize on the web by providing advertising tools: News publishers keep between 70-95% of the revenues generated through these tools[12].
- As we develop new Search and News products using generative AI, we'll continue to work in collaboration with news publishers and prioritize approaches that will allow us to send valuable traffic to them.
- Through the Google News Initiative, we are supporting training programs for journalists — so they can use AI in their work — and research into how AI can support the news ecosystem.[13]
- We have launched tools like Pinpoint, a research tool that helps journalists and academics analyse and explore large collections of documents. This tool is already being put to use in award-winning investigative reporting. A recent report by JournalismAI, the journalism think tank at the London School of Economics, reveals that survey respondents, including journalists and managers at news organizations, expect AI to

[10] Deloitte, Gen AI and the Creator Economy, 2023
[11] YouTube, New tools to protect creators and artists, September 2024
[12] Google, A look at how news publishers make money with Ad Manager, June 2020
[13] *See* The London School of Economics and Political Science, *JournalismAI Home Page*, Dept. of Media and Communications, https://www.lse.ac.uk/media-and-communications/polis/JournalismAI (last visited Oct. 27, 2023).

free up their capacity for more creative work by helping with time-intensive tasks, such as interview transcriptions[14].

These findings are echoed in the 2023-24 WAN-IFRA World Press Trends Outlook report. It found that:

- AI was the biggest priority technology/product for investment, with 87% saying it was a key area for investment over the next 12 months, followed by data analytics and intelligence (86%), video (79%) and audio/podcasts (74%).
- Just over a third (34%) were very optimistic that generative AI presented opportunities for their business, while 58% were somewhat optimistic and 8% were not optimistic at all.
- At the time of the survey 67% felt their business was poorly prepared to take advantage of those opportunities, with 16% each saying they were either not prepared at all or well prepared.[15]

Google Arts & Culture has partnered and supported a range of projects that allow artists and audiences to push the boundaries of creative experimentation with, such as:

- Collaborating on an "AI artist residency" at Somerset House to support emerging artists.
- Co-hosting a roundtable on AI and creativity to foster cross sector dialogue with Serpentine Galleries.
- An experimental project called *National Gallery Mixtape*, focussed on engaging audiences with the museum's collection in new ways, by using Generative AI to mix a personalised soundtrack inspired by the paintings in National Gallery's collection.

All these projects focus on AI as a collaborator, supporting artists interested in new modes of expression, or to find more routes to empower and enable a participatory experience for audiences.

Reforming the UK copyright framework to provide legal clarity

The right outcome will balance copyright with other policy objectives, including securing a future with AI to the benefit of society. Existing international frameworks make a provision for such a balancing of rights.

The copyright framework looks to address two public policy goals - incentivising the production and dissemination of public goods to the benefit of society, and protecting a private interest. This is why more legal clarity is needed on TDM for commercial purposes.

Existing international frameworks clearly make provision for the balancing of copyright with other fundamental rights and policy objectives. Article 7 of the World Trade Organisation's TRIPS Agreement recognises the balance between the importance of limitations and exceptions to secure the promise of knowledge goods to improve the welfare of society as a whole[16], and the preamble to the WIPO Copyright Treaty explicitly recognizes the need to "maintain a balance between the rights of authors and the larger public interest, particularly education, research and access to information."[17]

The role of limitations and exceptions in promoting public welfare is a matter of importance not only for users of works but for (follow-on) creators as well. Without the appropriate balance between protection and access, copyright systems risk limiting the key societal and economic advances offered by AI, but may also undermine creativity over the longer-term.

We believe that the most important factors for AI communities and web publishers are choice and control and we are working to improve machine readable approaches to achieve these goals, as outlined in the section below on opt-outs.

---

[14] London School of Economics, 'Nearly three quarters of news organisations believe generative AI presents new opportunities for journalism', September 2023.
[15] WAN-IFRA, World Press Trends Outlook 2023-24, 2024
[16] World Trade Organization, 'TRIPS Agreement - General Provisions and Basic Principles, Article 7', January 1995.
[17] World Intellectual Property Organization, WIPO Copyright Treaty, December 1996.

Creating the right environment for the UK tech and AI sector to grow

Taking a balanced, pro-innovation approach towards AI innovation

Foundational language models learn by training on diverse information and data from the open web - learning in much the same way as people do.

AI innovation fundamentally depends on the ability to learn from the widest possible variety of publicly available material. We believe that this learning is a non-expressive use of open web content that encourages innovation, science, creativity, and useful new works.

For AI development to succeed, we need copyright systems that take a balanced, pro-innovation approach. This means continuing to allow use of the data needed to train AI models, while still ensuring that rights holders can protect their creative works and the goal of copyright systems to foster creativity is honoured.

Google's foundational language models are trained on publicly available data from the internet, drawn from sources like blog posts, media transcripts, and public conversation forums. The responsible development of AI requires use of a diverse range of training data to learn from; training a model requires billions of data points to provide information, culturally relevant content, varied perspectives, and different forms and use of language to help guard against bias and take into account societal norms.

Training an AI model is a non-expressive use of data

There are no copies of content in the model itself, and the use that models make of data is non-expressive.

Generative AI models use what they have learned to create new content, such as text, images, music, and computer code. For example, a “large language model” (LLM) is a generative AI model that finds patterns in human language, making it suitable for a range of writing tasks, including predicting the next words to complete a sentence or suggesting grammatical edits that preserve what you mean to say.

During training, a model evaluates the proximity, order, frequency, and other attributes of portions of words, called tokens, in its training data. In fact, the model itself selects which attributes to use. In this way, training is the discovery of probabilities of relationships between the tokens — ultimately not in any individual text, but in all of the text on which the model is trained. The trained model then comprises a large network of weights that represent these learned relationships and can then respond to a prompt and generate new content with a probability of addressing the prompt as determined by its training. As such, no single piece of content is necessary to the training of AI models; the value comes from the total collection. It is the entirety of the dataset that makes it useful.

Generative AI models are not databases or information retrieval systems- there are no copies of content in the model itself. When, for instance, an LLM is prompted with a query, the model is generating responses based on a statistical estimation of what a satisfactory response should look like. Put simply, it produces an average group of words, pixels, or sounds related to a prompt. As such there are no specific content pathways in the model.

AI models that currently exist on the market differ as to their architecture, outputs and training processes. However, what is common is the fact that they make non-expressive use of content (i.e. they may involve temporary copying in a technical sense, but not in a normative sense, in that such uses do not communicate the expressive elements of the work to be read or otherwise enjoyed).

The problem is that the legal status of non-expressive uses is not articulated clearly in current UK copyright law. Therefore, while non-expressive uses should not be covered by the exclusive rights of the rights holder, we see a need for adding further legal clarity by having a commercial exception for TDM, which would undoubtedly enable the training of AI models in the UK. Such legal certainty is particularly important given the scale of investment required for AI development.

It is also important to note that the vast majority of generative AI models allow for substantial non-infringing uses and are a welcome tool used by many creators which should not be undermined. Moreover, many leading providers of GAI tools have safeguards to prohibit infringing uses. To support creators to prohibit the use of their works for training AI models, there are a number of tools that already exist such as robots.txt. These tools also provide transparency, as will be explained later in our response.

Training on the open web must be free; rights holders are still able to exercise their exclusive rights

Given the volume of data that models need to train on, any particular work in and of itself is not necessary for that training. Instead, it is the total collection of works that is needed to train an AI model.

We support web publishers in exercising their rights under UK law to prohibit or authorise use of their content for training AI models through a myriad of tools available. However, these rights are exclusive and do not translate to remuneration rights.

To develop and improve our AI models, we train on publicly available data from the web. We believe the use of such content is non-expressive and that it enables innovation. As such, we believe training on the open web must be free.

We understand some publishers are calling for payment for training on material already on the web. However, we reject the framing of rights reservation as a licensing mechanism. Web publishers have rights to authorise or prohibit use (which they already can do under existing law), and Google-Extended and other similar mechanisms allow them to exercise that right by opting out. However, this does not extend to a right to be paid. Reserving rights will not automatically translate to a license, not least because AI developers will have different strategies for data acquisition. A key reason for this is that no individual type - or piece - of content - be it news, blogs, recipes or reviews - has unique value in the training of an AI model. Given the volume of data that models need to train on, any particular work in and of itself is not necessary for that training. Instead, it is the total *collection* of works that is needed to train an AI model.

To enable the responsible development of AI - that can guard against bias and ensure cultural relevance and varied perspectives - this training will run into billions of data points. At which point, payment for training becomes prohibitively complex - particularly large foundational models that are often the key to unlocking new breakthroughs that require enormous compute power and data to find patterns in the dataset.

A study from think tank Bruegel[18] argues that data inputs in an AI model do not adhere to the Euler Theorem - defined by the thesis that the "remuneration of inputs should be in accordance with their marginal contribution to the value of outputs". The study notes that "the Euler Theorem comes into problems when applied to generative AI models" because "unlike physical goods, data inputs and media outputs are non-rival products, as they can be re-used without limits. [...] Non-rivalry undermines the Euler Theorem and remuneration according to marginal productivity. Doubling the volume of data inputs will not double generative AI model outputs or productivity."

Furthermore, large input data sets are needed to improve model performance, but if licensing is mandated, there is a risk that AI firms, particularly start ups and scale ups, will struggle to determine an economically meaningful and efficient licensing price for media inputs. No single piece of content has value, and as such, pricing becomes a pure bargaining issue. As such, we urge the Government to avoid any suggestion that the rights reservation should automatically lead to licensing.

The consultation also references Collective Management Organisations (CMOs). However, licensing a CMO's repertoires for the purpose of generative AI training would be extremely challenging and disproportionate due to the variety of works and economic rights potentially involved. This would result in a fragmentation of rights, not least as regards crawled content where the copyright status is in most cases unknown. The latter adds additional layers of complexity as regards 'unknown works', which CMOs would be unable to provide to the respective

[18] Bruegel, 'Economic Arguments In Favour Of Reducing Copyright Protection For Generative AI Inputs And Outputs', April 2024

licensors.

However, having a balanced copyright framework, such as TDM exceptions, does not prevent AI model providers and right holders from finding new commercial solutions for access to content: negotiating agreements and partnership deals for a variety of situations, including programmatic access to custom APIs, access to data, digitisation, etc. At Google we have concluded such agreements, and we see many of our competitors in the AI space agreeing to similar deals. Given the early stages of this technology, of commercial decisions, and the complexity of license deals themselves, it would be premature to introduce new measures to mandate licensing. Early signs show that a market is developing and as such a voluntary licensing market should be given time to develop and grow.

Supporting greater control and choice for web publishers

Opt-outs are technically feasible

Over half of news publishers, among others, are already using robots.txt to block web crawlers. These mechanisms are well-understood, and used on websites of all kinds and sizes.

The consultation itself states that over half of news publishers block the main generative AI web-crawlers using the robots.txt standard and AI developers generally respect this standard and offer various implementations of it.[19] A study from the Reuters Institute at Oxford University likewise found that nearly a quarter of news publishers are already blocking Google's AI crawler, and half of sites are already blocking OpenAI's crawler[20]. Nor is the use of robots.txt limited to use by news publishers. Companies like YouTube use robots.txt to control how their content is accessed by web crawlers.

In September 2023, Google launched Google-Extended, a new control that web publishers can use to manage whether their sites help improve Gemini and Vertex AI generative APIs, including future generations of models that power those products. Making simple, scalable and well-recognized controls, like Google-Extended, available through robots.txt is an important step in providing transparency and control.

Any web publisher can implement Google-Extended to opt out of model training. It uses the well-established robots.txt technology with which publishers are already familiar, and it does not impact a site's inclusion or ranking in Google Search. Publishers such as The Guardian have implemented Google-Extended and you can still see them in Search. It is easy for anyone to check that Google Extended has been applied to a web publisher site by adding /robots.txt to the web address.

Overwhelmingly, the feedback we have received from site owners is that they understand robots.txt and its mechanisms, and are able to use it appropriately to achieve their goals. Independent studies such as the open-source Web Almanac for 2024 have shown that the vast majority of the web uses these mechanisms, and they also use them specifically for the newer AI crawlers. These mechanisms are well-understood, and used on websites of all kinds and sizes.

Furthermore, discussions around opt-outs and what "*machine readable*" means are ongoing across industry, government institutions, and civil society. Rightholders across different creative sectors may have different preferences for different types of machine-readable opt-outs, while AI developers, in turn, require clarity, market adoption, and technical feasibility to make sure this can all work in practice. Building on the long-established robots.txt protocol, it is clear that an opt-out mechanism is technically practical and achievable as part of a compromise to enable AI development and rights holder control. While governments may be interested in

---

[19] *"82: There is a growing use of rights reservation protocols and standards by both AI developers and right holders. Over half of news publishers block the main generative AI web-crawlers using the robots.txt standard. AI developers generally respect this standard and offer various implementations of it. For example, Google, Microsoft Bing and OpenAI each adopt their own approach."* Copyright and AI: Consultation

[20] University of Oxford Reuters Institute, 'How many news websites block AI crawlers?', February 2024.

promoting a one-size-fits-all-solution/common standard, we believe it is important that, at this time, industry has flexibility and room for the evolution of standards and best practices.

If however the Government plans to standardise opt-outs, we urge this to be aligned with robots.txt type controls that are already widely accepted and used by right holders and model providers alike. Robots.txt provides the granularity needed, beyond what is referenced in the consultation (paragraph 86)[21], as it allows users to specify individual URLs to opt out. Robots.txt can be used on a site-level, for URL patterns, and for individual URLs. A URL can be a page or a file (for example, a web-page can contain images, with each image potentially being a separate file with a separate URL). Our view is that any work on standards should build on these effective and long-established protocols.

While standardisation as referenced in the consultation is an acceptable goal, this is fundamentally different from imposing a standard that might not work for all providers. As outlined, we believe that this consultation should consider all types of AI developers, and have an element of flexibility to allow a standard to develop organically. In addition, by applying one standard to AI developers, the UK risks being outdated quickly in an evolving market.

Transparency requirements should not undermine AI activities

Excessive transparency measures can be unworkable for AI companies of all sizes, including start-ups. They risk requiring companies to share trade secrets, to police the internet, and some proposals are impossible due to the lack of available information.

As outlined above, an opt-out is an effective mechanism to deliver choice and transparency for rightsholders, whereas certain types of transparency measures can be unworkable for AI firms of all sizes, including start-ups.

The absence of a requirement under international and UK copyright law to register works or transfers of rights, inconsistent quality of rights ownership metadata and a generally low threshold for copyright protection make it impossible for AI developers to identify and document the copyright status of third party works used for training and make such information available to rightholders.

Moreover, there are major risks associated with the UK requiring the disclosure of training data as outlined in the consultation:

- This could equip bad actors with the knowledge to attack the model, causing security concerns.
- The selection and preparation of training data may be protected by trade secrets. Certain transparency requirements raise the risk of the disclosure of sensitive information and thus jeopardises investment, including in and by start-ups, in the development and deployment of AI models here in the UK.

Under current legal frameworks, disclosure is already ensured in specific enforcement proceedings. This disclosure is purposeful, meaning it concerns a specific subject matter and a specific rights holder, and is subject to the normal checks and balances of the judicial system, including proportionality, protection of trade secrets and confidentiality of proceedings. This approach is welcome as it provides safeguards for all parties involved, helping to build trust. We would urge caution against unnecessary divergence from current legal frameworks due to the risks and practical hurdles outlined above.

Additionally, if evidence is required on the compliance with rights reservation mechanisms - such as the opt-out - this could be met through publishing information on web crawlers. However, the scope of information required again risks the sharing of commercially sensitive information that is tantamount to a trade secret.

---

[21] *"Additionally, the most widely-adopted standard – the robots.txt standard –cannot provide the granular control over the use of works that many right holders seek. It allows works to be blocked from web crawling at the site level but does not recognise reservations associated with individual works. It also does not enable right holders to distinguish between uses of works. For example, they may be content for web crawlers to use their works for search indexing or language training, but not for generative AI. Robots.txt does not currently allow for this degree of control."* Copyright and AI: Consultation

Finally, there are practical hurdles associated with any requirement to publish detailed transparency reports. This is because this reporting is effectively asking developers to disclose and summarize all the content on a web that is constantly changing and dynamic. This makes it impossible to produce and maintain an up-to-date summary.

Our view on the consultation options

Our strong recommendation is an amended option 3 that enables TDM for any purpose with a rights reservation for rightsholders. We also encourage the UK to secure a competitive advantage by introducing a commercial research exception.

Below we outline our views on the consultation options and how we see these working in practice to support the UK's ambitions.

Option 0: Do nothing: Copyright and related laws remain as they are

Option 0 does not offer sufficient clarity under the current UK framework and does not meet the Government's stated aims as set out in the consultation. As the consultation outlines, the current uncertainty regarding the application of UK copyright law to the training of AI models is hindering innovation and undermining the UK's broader ambitions for AI, as well as the growth of the creative industries. This was also noted by Lord Patrick Vallance in his past review[22], which recommended that the previous government announce a clear policy position on the relationship between intellectual property law and generative AI to create an environment in which TDM is enabled in the UK, providing confidence to innovators and investors.

Option 1: Strengthen copyright by requiring licensing in all cases

Option 1 would make the UK significantly less competitive compared to other jurisdictions and may not increase the level of licensing, as noted in the consultation proposals.

It is important to address the idea that strengthening copyright is equal to mandating licensing. Copyright's fundamental purpose is to balance rewarding and incentivizing creators with enabling beneficial uses. Requiring licensing in all scenarios disrupts this balance, hindering the evolution of creative models and the work of follow-on creators, especially those reliant on AI, and ultimately stifling future innovation.

Further, restricting models trained in other jurisdictions, which do not meet UK standards, would create an environment of regulatory uncertainty that would hinder confidence in AI development. This may mean, as also noted in the consultation itself, that providers would not be able to release new models onto the UK market - irrespective of where the training happened - thereby, denying UK users from assessing these tools as well as businesses of all sizes across all sectors of the economy.

Given the rapid rate of model improvement, with major breakthroughs occurring on 6-12 month cycles, even short delays in access to models in the UK could negatively impact the possible productivity gains of AI.

Finally, in the context of global AI competition, it is important to recognise the potential consequences of restrictive policies. Licensing and transparency mandates may impede technological progress in some areas, but they will not stop the world from advancing AI. This approach risks leaving the UK at a distinct disadvantage.

Option 2: A broad data mining exception

Option 2 is the most competitive option for the UK to be a leader in AI. It would bring a similar approach to the UK that aligns with Japan and Singapore - and would fully support the Government's ambitions laid out in the AI Opportunities Action Plan.

Option 3: A data mining exception which allows right holders to reserve their rights, underpinned by supporting measures on transparency

[22] UK Government, 'Pro-innovation Regulation of Technologies Review: Digital Technologies', March 2023.

Option 3 would be a minor improvement from what is in place today but, as set out, is unlikely to deliver on the Government's ambitions outlined in the consultation.

We agree that rights holders should have choice and control, and acknowledge that rights reservation in option 3 would deliver on this aim. As outlined above, appropriate technical solutions exist and are working to facilitate the reservation of rights. Google Extended is one example of a tool that enables choice and transparency over whether content is used to inform the improvement of AI models.

However, as drafted, proposal 3 risks being undermined by transparency requirements, particularly suggestions in the consultation that AI firms and developers should have some level of control over the expression and duplication of work.[23] It is not the role of AI developers to police the internet and for good reasons AI developers cannot reliably know if a piece of content is using the expression of work as there is no registration of copyright works, and we have no way of knowing whether that use is legitimate or not because it could be covered by a copyright exception or a licence.

Preferred option: Amended option 3

To ensure the UK enhances its competitive position in AI development, and to support the government's wider objectives of being home to future AI growth and investment, we believe the Government needs to catch up with other jurisdictions and gain a competitive edge by implementing a copyright regime that:

- Maintains the current non-commercial research exception;
- Enables TDM for any purpose with rights reservation for rightsholders;
- Provides the UK with a competitive advantage by introducing a commercial research exception.

Therefore, our strong recommendation is an amended option 3.

AI Outputs

Infringement and liability relating to AI-generated content:

Generative AI is a technology engineered to create new works, not to copy or facilitate the copying of existing works. Like previous technology transitions in the past, including digital photography, smartphones, synthesizers, and image editing software changed the artistic field, these tools augment and create new opportunities for artists and open up opportunities to more people.

While occurrences are rare - and are becoming even more so as the performance of models increases - AI systems can output content that seems similar to individual pieces of content on which they were trained. According to the flagship research paper in the field, this is an exceedingly rare occurrence, even under adversarial prompting.[24] The possibility that AI models can occasionally, despite the best efforts of their developers, output content that replicates existing expression is a bug not a feature, and developers are taking a range of measures to limit that occurrence even further, including deduplication of training data before training.

The possibility that a generative AI system can, through "prompt engineering," be made to replicate content from its training data does, however, raise questions around the proper boundary between direct and secondary

[23] *"81. Such an approach will prevent the use of copies of a work to which a machine-readable reservation has been applied. But other copies of the same work may exist to which it has not been applied. In that situation, it may be desirable that the effects of a rights reservation apply more broadly than the individual copy to which it applies. This would require developers to make efforts to ensure that the use of any expression of the same work is avoided."* Copyright and AI: Consultation

[24] The researchers reported that 94 near duplicate images were created, out of 175,000,000 attempts. These attempts were focused on the 350,000 images with the largest number of duplicates in the training set (500 attempts per image). The researchers needed to have not only information about which of the 160,000,000 images in the training set to duplicate, but also the caption data used to identify those most-duplicated images. They also needed access to immense compute power, which is something that the majority of users do not access to. *See* Nicholas Carlini et al., *Extracting Training Data from Diffusion Models* at 4-7, January 2023

infringement. Such activities go against our terms of service and the way we envisage the use of such tools. When an AI system is deliberately prompted by a bad actor to produce an infringing output, any resulting liability should attach to said actor. A rule that would hold AI developers directly (and strictly) liable for any infringing outputs created independently of the developers intention would impose unreasonable burdens on AI developers, even if they have undertaken significant measures to prevent infringing activity. Had that standard applied in the past, we would not have legal access to photocopiers, personal audio and video recording devices, or personal computers — all of which are capable of being used for infringement as well as for substantial beneficial purposes.

AI output labelling

Being able to determine whether content is AI-generated is critical to empowering people with knowledge of when they're interacting with generated media, and for helping prevent the spread of misinformation. As we invest in more capable models, we are also deeply investing in AI responsibility. That includes developing tools to ensure people know when an image or video has been altered or generated by AI, while preserving privacy and protecting artistic expression.

We believe all developers should be encouraged to develop voluntary provenance mitigations, such as metadata inclusion, watermarking, or similar techniques for audio-visual content, which may be helpful to determine if a particular piece of content was created with their system. For example, Google's SynthID toolkit watermarks and identifies AI-generated content. SynthID embeds digital watermarks directly into AI-generated images, audio, text or video. The toolkit is now being integrated into a growing range of Google products, helping empower people and organizations to responsibly work with AI-generated content. And while SynthID isn't built to directly stop motivated adversaries like cyberattackers or hackers from causing harm, it can make it harder to use AI-generated content for malicious purposes[25]. We believe information literacy is a multifaceted endeavor that must encompass the wider technology ecosystem in order to be effective. To this extent, last year, we open sourced SynthID for text, enabling developers to use this technology to help detect whether text outputs have come from their LLMs.

In addition, on YouTube, we have rolled out a policy that requires creators to disclose altered or synthetic content that is realistic, meaning that a viewer could easily mistake what's being shown with a real person, place, or event[26]. Labels will then appear within the video description information, and if content is related to sensitive topics like health, news, elections, or finance, a label will also be displayed on the video itself. In some cases, YouTube may add a label even when a creator hasn't disclosed it, especially if the altered or synthetic content has the potential to confuse or mislead people. And on our Ads products, we were also the first tech company to require election advertisers to prominently disclose when their ads include realistic synthetic content that's been digitally altered or generated, including by AI tools. In February 2024, Google became a steering member of the C2PA, a cross-industry effort to help provide more transparency and context for people when it comes to AI-generated content. C2PA has the benefit of being tamper evident and highly interoperable, making it an excellent vehicle for cross-ecosystem technical collaborations to signal the provenance of content at scale. We have also recently joined the International Press Telecommunications Council (IPTC) as a Voting Member.

This said, we'd caution against giving the impression that any regulatory or technical solution is a silver bullet to the challenge. Artificial intelligence innovation raises complex questions that neither Google, nor any other single company, can answer alone. Getting it right will require continued collaboration among companies, academic researchers, civil society, governments, and other stakeholders. This is still a relatively new area of research and product development, and the range of technical solutions available are rapidly evolving.

As such, legislation in this space should ensure that requirements are technically feasible, improve end user experience, do not require disclosure of trade secrets, and do not compromise product safety nor human rights. Specifically, requiring that a creator's identity be linked to a piece of content, for example via C2PA metadata,

[25] See John Kirchenbauer et al, On the Reliability of Watermarks for Large Language Models, June 2023
[26] YouTube, How we're helping creators disclose altered or synthetic content, March 2024

would create significant risks in terms of privacy and free expression. According to WITNESS, a human rights and technology group, "We should be wary of how [watermarking and metadata] could be used by governments to capture personally identifiable information to supercharge surveillance and stifle freedom of expression, or facilitate abuse and misuse by other individuals." A wide range of people — for example, human rights defenders and survivors of domestic violence — have legitimate reasons to maintain anonymity, and therefore should be able to make their own choice as to whether to attach information about their identity to a piece of content.

Finally, it is important to bear in mind that not all AI-generated content is deceptive and not all deceptive content is AI-generated. In fact, we expect that over time a large amount of high-value content will be AI-generated or augmented, and we want to make sure users can easily discover and benefit from that content as needed. Labeling some content as "synthetic" has the potential to increase trust in unlabeled content, even if it may be deceptive for other reasons, like being misrepresented or taken out of context. This phenomenon, which is well documented in academic research[27], creates a danger of distorting users' understanding in general, where they will believe all unlabeled content is "real" and may lose trust in the sources of labeled content. We outlined other limitations of user-facing labels in our public paper last year[28].

Digital replicas and other issues

Copyright law has always been about protecting the way or form in which ideas are expressed, and does not extend to protecting any subject matter that exists independent of copyright, e.g., a mathematical equation, a flower or, in this case, an individual's likeness. For example, where a photographer takes a photo of a landscape with individuals in the photo, the photographer owns the copyright in the photo, and copyright law has never afforded the individuals the right to prevent any subsequent exploitation of the photo. This, however, does not preclude the application of other laws, e.g., those dealing with passing off, breach of confidence, defamation, online harm, etc. Expanding copyright law to cover generative AI model outputs of individuals' likenesses would not only impermissibly expand the scope of copyright to facts, doing so would not guarantee that the alleged benefits will reach the individuals depicted, flowing instead to the owners of the copyright interest in the copyright-protectable work containing the depiction. Care must be taken, as new technologies like AI impact aspects of "personality rights" that are only tangential to copyright, not to expand copyright in ways that restrict creative activity and free expression.

Where countries are considering new measures to guard against harmful synthetic digital imitations ("digital replicas"), we believe the right approach is to focus on addressing the actual harms that might occur. In general, with respect to digital replicas, such harm is caused by the publication of such replicas and not by mere private creation. Attaching liability to the creation of synthetic digital imitations alone will disincentivize developers from designing tools that benefit creative output for fear of being held liable as a co-creator of violative content. It is the user who will have prompted the tool—often working to circumvent existing safeguards designed into the tool itself—to produce a harmful output and as such should be held responsible for this action. To be liable, the person making the unauthorized publication must have had actual knowledge of both the synthetic nature of the material and the harmfulness of the publication.

It is also critical to recognize that certain use cases should be protected because they are key to creative expression; examples of such may include content which is not explicitly misleading, communicated for the purpose of parody, satire, news reporting, education, and research, as well as an incidental use.

Finally, it is important that the existing conditional safe harbours for intermediary services continue to apply. Thus, we believe changes to address these rights should ensure a focus on liability solely for the individuals who seek to disseminate harmful digital replicas, while ensuring protections for valuable expressive content.

---

[27] See Pennycook et al, The implied truth effect: Attaching warnings to a subset of fake news headlines increases perceived accuracy of headlines without warnings, 2020

[28] Google, Determining trustworthiness through context and provenance, December 2024

Other emerging issues: Synthetic Data

Synthetic data is artificially created data that mimics real-world data, but doesn't contain actual real-world observations. It's created using computational methods and simulations, such as generative AI algorithms, to mimic the statistical properties of real-world data. Synthetic data is useful when there isn't enough real-world training data to improve ML accuracy, such as when testing or creating new models. For example, in a space where the number of real-world examples is small, such as due to lack of data collection or an obscure topic, a specialized data generating model can be trained on a set of the existing real-world examples, then that data generating model can be used to generate a number of purely synthetic additional examples. An example of this is in the field of computer code generation or mathematics, where a model could be used to generate both new questions and solutions to questions, which could then be used in subsequent training. This can broaden the scope of available relevant and useful data for training and testing of a more general purpose model. By broadening the scope of training data, synthetic data helps ensure that the model learns from patterns in the data, and not from specific examples. Therefore, the very limited, if any, copyright relevance of synthetic data is neutral-to-beneficial to the copyright ecosystem.