Response ID ANON-2HEH-VQDT-6

Submitted to Consultation on Copyright and Artificial Intelligence
Submitted on 2025-02-25 22:20:59

Basic Information

1 What is your name?

Name:
Rachael Drury

2 What is your email address?

Email:
R.L.Drury2@liverpool.ac.uk

3 What is your organisation?

Organisation:
Rachael is a researcher and lecturer in Music Industry Studies, Copyright, and AI Ethics at the University of Liverpool, with twenty years' experience as a professional violinist and music practitioner. Her PhD, submitted in January 2025, explored the legal and economic challenges posed by generative AI in the music industry with a focus on the impact of AI on the economic framework of the music industry and the livelihoods of music creators. A recipient of the prestigious IASPM Andrew Goodwin Memorial Prize (2021), Rachael combines her extensive practical expertise with a passion for advancing ethical and sustainable approaches to innovation in the creative industries.

Copyright and Artificial Intelligence

4 Do you agree that option 3 - a data mining exception which allows right holders to reserve their rights, supported by transparency measures - is most likely to meet the objectives set out above?

No

Please give us your views:

No, I do not agree that Option 3 meets the government's stated objectives. As an academic researcher specialising in AI, copyright law, music industry studies, I find that Option 3 fails to balance the interests of creators, AI developers, and the broader public. It undermines copyright protections, weakens established licensing frameworks, and poses significant risks to the economic sustainability of the UK's creative industries.

Option 3 fails to meet the government's stated objectives for the following reasons:

1. The Economic Value of the UK Creative Industries

The UK's creative industries are a cornerstone of the national economy, contributing £124.6 billion to the UK economy in 2023, accounting for 5.7% of UK GVA. The sector supports 2.4 million jobs, with 311,000 in music, performing, and visual arts - a 15.1% increase since 2019. Between 2010 and 2022, the creative industries' GVA grew by 50.3%, far outpacing the overall UK economy's 21.5% growth.
(https://lordslibrary.parliament.uk/creative-industries-growth-jobs-and-productivity/)

This success is underpinned by robust copyright protections that incentivise creativity and allow creators to earn a living from their work. Copyright exceptions, such as s.29A CDPA (the non-commercial TDM exception), already provide AI developers with legal pathways to conduct public-interest research without undermining creators' rights. Those seeking to mine copyrighted works for commercial purposes can and should engage with the existing licensing frameworks that sustain the industry.

By introducing a broad commercial TDM exception with an opt-out system, Option 3 threatens to destabilise this delicate balance. It would disincentivise AI developers from seeking licences, undermining the core revenue streams that the music industry relies upon. In essence, this approach devalues creative works, treating them as free fuel for AI development while cutting creators out of the value chain.

2. Erosion of Copyright Enforcement and Creator Control

Even with an opt-out mechanism, Option 3 undermines copyright enforcement. AI developers would be able to mine copyrighted works without prior consent, making it difficult for rights holders to monitor, control, or challenge the use of their works. This undermines one of copyright's fundamental principles: that creators have the exclusive right to decide how their works are used.

AI-generated outputs that closely resemble human-created works could be shielded under the legal ambiguity created by Option 3. Developers could argue that their models were trained on "lawfully" accessed data under the TDM exception, even if the resulting outputs infringe on original works. This legal loophole severely limits creators' ability to protect their work.

3. Transparency as a Non-Negotiable Requirement

Transparency is essential for any copyright framework that interacts with AI development. Unfortunately, Option 3's transparency obligations are insufficient and would be difficult to enforce. The history of voluntary transparency initiatives, such as the EU Code of Practice on Disinformation, shows that tech companies are unlikely to comply without strict legal mandates.
(https://digital-strategy.ec.europa.eu/en/policies/code-practice-disinformation)

To protect creators, full transparency must be a legally enforced requirement, not a trade-off for weakened rights. AI developers should be required to disclose:
• The datasets used in model training.
• The copyrighted works included.
• The licensing status of these works.

Without such transparency, fair licensing agreements and enforcement will be impossible, leaving creators unable to monitor how their works are used or seek fair remuneration.

4. Disregard for Existing Licensing Frameworks

The music industry already operates under effective licensing systems that ensure creators are compensated when their works are used. AI developers argue that paying for training data is prohibitively expensive, yet they routinely invest billions in cloud computing, data storage, and engineering talent. Training data is a critical resource, and licensing fees should be treated as a standard business cost.

Licensing for AI is not only feasible—it is already happening. The APRA AMCOS report highlights voluntary licensing agreements between AI developers and rights holders (https://www.apraamcos.com.au/about/supporting-the-industry/research-papers/aiandmusic). Similarly, the GEMA/SACEM report emphasises the importance of robust licensing to ensure ethical AI development (https://www.gema.de/en/news/ai-study). These examples demonstrate that AI developers can and do license content when required, refuting arguments that licensing would stifle innovation.

The CISAC Report highlights the real-world impact of unlicensed AI training, predicting that the music industry will lose nearly 25% of its revenue to AI-generated content within four years
(https://www.cisac.org/Newsroom/news-releases/global-economic-study-shows-human-creators-future-risk-generative-ai). Without assurances that AI developers will adhere to licensing requirements, these figures will become a reality.

5. Exploiting the Non-Commercial TDM Exception

The non-commercial TDM exception (s.29A CDPA) is being exploited as a loophole by AI developers. Companies scrape vast amounts of copyrighted content under the guise of non-commercial research, only to later commercialise their models and outputs.

The LAION v. Robert Kneschke case in Germany exposed this issue
(https://www.euipo.europa.eu/en/law/recent-case-law/germany-hamburg-district-court-310-o-22723-laion-v-robert-kneschke). The Regional Court of Hamburg ruled that LAION's non-commercial dataset creation was lawful, even though downstream users—including commercial AI companies—relied on the dataset. This ruling highlights the legal grey area that Option 3 would perpetuate in the UK, allowing AI developers to sidestep licensing obligations while still profiting from creators' works.

6. The Streaming Problem: AI and Unfair Remuneration

The streaming economy already struggles with fair remuneration for artists. The pro-rata payment model has been widely criticised for disproportionately favouring high-streaming artists, leaving many musicians earning less than £20,000 annually from their music. The introduction of AI-generated music into streaming platforms exacerbates these issues.

AI-generated tracks can be produced at scale and uploaded en masse, flooding streaming services and competing directly with human-made music. Cases like Michael Smith's $10 million streaming fraud—where AI-generated songs were streamed by bots to siphon royalties—demonstrate how easily the system can be exploited (https://www.justice.gov/usao-sdny/pr/north-carolina-musician-charged-music-streaming-fraud-aided-artificial-intelligence).

Streaming platforms also have financial incentives to prioritise AI-generated content, which often incurs lower royalty obligations. Without strict regulation, AI-generated music could dominate streaming playlists, further marginalising human musicians and reducing their income opportunities.

7. Conflicts with Domestic Law: Competition and Data Protection

Option 3 conflicts with UK competition law. The Competition and Markets Authority (CMA) and the Digital Markets Competition and Consumers Act 2024 (https://www.legislation.gov.uk/ukpga/2024/13/contents) prohibit practices that distort market competition. Allowing AI companies to mine copyrighted works without licences gives them a significant market advantage, undermining fair competition and disadvantaging human creators.

Data protection is another significant concern. Under the UK GDPR, personal metadata embedded in music—such as artist names, ISRC codes, and licensing information—constitutes personal data. AI developers scraping music for training purposes often capture this data without consent, violating GDPR principles. The Information Commissioner's Office (ICO) has highlighted the risks of "invisible processing," where individuals are unaware their data has been used (https://www.skadden.com/insights/publications/2025/01/ico-publishes-outcomes-of-genai-consultation).

8. Violations of International Copyright Treaties

Option 3 violates the Berne Convention's three-step test (Article 9(2)), which requires that copyright exceptions:
• Be limited to special cases.

• Not conflict with the normal exploitation of the work.
• Not unreasonably prejudice the legitimate interests of rights holders.

A broad commercial TDM exception fails all three steps. It is not limited to special cases, conflicts with the normal exploitation of music through licensing, and severely prejudices creators' interests by depriving them of fair remuneration.

Option 3 also conflicts with the Council of Europe's Framework Convention on Artificial Intelligence, which mandates that AI respects human rights, including the economic rights of creators, fairness and individual autonomy (https://www.coe.int/en/web/artificial-intelligence/the-framework-convention-on-artificial-intelligence). By enabling AI developers to exploit copyrighted works without proper licensing or transparency, Option 3 undermines these principles.

Conclusion

As an academic focused on AI, copyright, and the music industry, I find that Option 3 not only fails to meet the government's objectives but also poses serious risks to the sustainability of the UK's creative industries. It weakens copyright protections, disrupts established licensing frameworks, and undermines creators' ability to earn a living from their work.

A more effective solution would be to enforce existing copyright laws, strengthen licensing frameworks, and mandate transparency obligations. These measures would ensure that AI development can continue while safeguarding the rights and livelihoods of human creators—an essential balance for the future of the UK's creative economy.

5 Which option do you prefer and why?

Option 2: A broad data mining exception

Please give us your views.:

I strongly support Option 1 which strengthens copyright by requiring licensing in all cases. In fact, Option 1 does not "strengthen" copyright as by law, licensing is already required for commercial purposes. However, both copyright and licensing frameworks can be strengthened with legally binding transparency obligations. This approach offers the most balanced and effective solution, safeguarding creators' rights while enabling AI developers to operate legally, ethically and sustainably.

Copyright law in the UK is already clear and fit for purpose: commercial generative AI training requires explicit licensing. However, widespread non-compliance by AI developers, who often bypass licensing obligations, has created an urgent need to strengthen existing frameworks. Option 1 would provide the clarity and legal certainty necessary to ensure that AI developers respect copyright law, while upholding the fundamental principle that creators must give permission before their works are used.

A strengthened licensing framework, paired with enforceable transparency measures, would benefit all stakeholders by ensuring that creators retain control over their works while allowing AI developers access to high-quality data through legal channels. Transparency obligations are critical, not as a trade-off, but as a foundational requirement in any future-proof copyright system.

Why Option 1 Meets the Government's Objectives:

• Maintains creators' control over how their works are used.
• Ensures fair remuneration through enforceable licensing agreements.
• Allows AI developers to access data legally and sustainably through clear licensing pathways.
• Promotes legal certainty for both rights holders and AI developers, reducing ambiguity.

Option 1 strikes the necessary balance between encouraging AI innovation and safeguarding the rights and livelihoods of human creators.

Why the Other Options Fall Short:

• Option 0: Maintain the Status Quo
While UK copyright law is clear, Option 0 fails to address the transparency gap, which is critical for ensuring that rights holders know when and how their works are being used. Without transparency, it is nearly impossible for creators to monitor unauthorised uses or enforce their rights effectively.

• Option 2: Broad TDM Exception
A broad text and data mining (TDM) exception would fundamentally undermine copyright protections by allowing AI developers to use vast amounts of copyrighted works without permission or fair remuneration. This would create an unfair marketplace, where AI-generated content—produced using unlicensed data—directly competes with human-made works, devaluing the creative economy.

• Option 3: TDM Exception with Opt-Out and Transparency Measures
While Option 3 superficially addresses transparency, it remains deeply flawed. The opt-out model shifts the burden onto creators, forcing them to police unauthorised uses of their work, which contradicts the foundational principle of copyright law—that creators should grant prior permission, not rely on retroactive enforcement.

It must be stated that transparency should be a universal legal requirement - not attached exclusively to Option 3 - and must apply to all licensing frameworks to ensure proper oversight.

The Illusion of Control in Opt-Out Models

Option 3 creates a false sense of control. While it allows rights holders to “opt out” of AI training, in practice, this system is impractical and ineffective. Most creators lack the resources to track where and how their works are being scraped and used. Moreover, AI models cannot "unlearn" data once it has been incorporated into their training sets, meaning that even if a creator opts out later, their works may continue to influence AI-generated outputs.

The burden should not fall on individual creators to protect their works. The responsibility must lie with AI developers, who should be legally required to seek licences before using copyrighted content for training.
Supporting Evidence: The Real-World Impact of AI Training Without Licensing

Multiple reports from leading industry organisations highlight the risks posed by AI training without proper licensing:

APRA AMCOS, GEMA/SACEM, and CISAC have shown that:

• AI systems often infringe upon input data, including copyrighted works.
• AI-generated outputs directly compete with the original works they are trained on, undermining the economic potential of human-made music.
• This dynamic leads to lost revenue for music creators and contributes to the devaluation of human-created content.

The CISAC Report projects that AI-generated music could replace up to 25% of human-created music revenues within four years, highlighting the urgent need for protective measures to ensure the sustainability of the creative industries.

Additionally, streaming fraud, exacerbated by AI-generated content, continues to siphon billions of dollars from legitimate creators. A case involving Michael Smith, who generated over $10 million in fraudulent royalties using AI-generated tracks combined with bot streaming, underscores the scale of potential abuse if licensing frameworks are not enforced.

Legal Clarity and International Compliance

The UK copyright system is already designed to protect creators and support licensing as the primary mechanism for content use. Strengthening this framework under Option 1 would:

• Enforce existing laws that require explicit permission for commercial use of copyrighted works.
• Align the UK with international copyright treaties, including the Berne Convention and TRIPS, which uphold creators’ rights to control and be fairly remunerated for the use of their works.
• Provide legal certainty for both AI developers and rights holders, reducing ambiguity and encouraging fair licensing practices.

Why Option 1 Is the Only Sustainable Path Forward

Option 1 is the only choice that upholds the rights of creators, ensures fair remuneration, and maintains a healthy, competitive marketplace for both human and AI-generated content. It strikes the necessary balance between fostering innovation and protecting the creative economy, offering a clear and enforceable framework for AI development that respects existing copyright laws.

By reinforcing licensing obligations and establishing legally binding transparency measures, Option 1 creates a system where:

• Creators retain control over their works.
• AI developers access data through legal, transparent means.
• The integrity of the UK’s creative industries is preserved.

In contrast, Options 2 and 3 would erode the core principles of copyright law, undermining the livelihoods of music creators and threatening the long-term sustainability of the UK’s creative economy.

The UK has long been a global leader in both creative and technological innovation. Adopting Option 1 allows the government to maintain this leadership while protecting the rights of those who contribute so significantly to the cultural and economic fabric of the nation.

Our proposed approach: Exception with rights reservation

6 Do you support the introduction of an exception along the lines outlined in section C of the consultation?

No

Please give us further comments.:

No. The introduction of a copyright exception with a rights reservation (opt-out) system is neither fair nor workable for music creators.

I find the proposal for an opt-out-based TDM exception fundamentally flawed. It undermines the core principles of copyright law, fails to provide meaningful protections for creators, and risks severe economic and cultural consequences for the UK’s creative industries.

In addition to the concerns raised in my response to Question 4, this answer outlines five key reasons why I do not support the introduction of an opt-out exception as outlined in Section C of the consultation:

1. Opt-Out Undermines Copyright's Core Principles

The UK copyright system is built on a permission-based framework, not assumption. Copyright law grants creators the exclusive right to copy, distribute, adapt, and publicly communicate their works (CDPA, s.16(1)(a-e)), ensuring that permission is required before any of these acts occur. Musicians, like all creators, expect automatic legal protection the moment their work is fixed in a tangible form - a foundational principle of copyright law.

Option 3 undermines this principle by reversing the burden of enforcement. It allows AI developers to use copyrighted works by default, forcing creators to actively opt out to protect their rights. This fundamentally contradicts CDPA s.16(1)(a), which defines copying as a restricted act, and s.17(2), which extends this definition to electronic reproduction. Training generative AI models, which involves copying vast amounts of copyrighted material for datasets, clearly falls within this definition and should require explicit licensing.

Additionally, the opt-out system is practically ineffective. Once a work is scraped and used to train an AI model, it becomes impossible to fully retract. Even if a creator successfully opts out after the fact, the AI model cannot "unlearn" the data it has already absorbed. This means that creators' works continue to influence AI-generated outputs indefinitely, without permission, credit, or remuneration.

Copyright operates on an opt-in, permission-based system. AI developers should not be exempt from this core principle simply because of the scale at which they operate. To uphold creators' rights and ensure fair licensing, AI developers must be legally required to seek explicit permission before using copyrighted works in training datasets.

2. The EU Opt-Out Model Has Already Failed

The UK's proposed opt-out system mirrors the EU's approach under Article 4 of the DSM Directive (2019/790) (https://eur-lex.europa.eu/eli/dir/2019/790/oj/eng), which has already proven ineffective in protecting rights holders. Several critical failures in the EU model highlight why replicating it in the UK would be detrimental:

• Rights holders remain unaware of the need to opt out, resulting in unauthorised AI training on their works. Many creators, particularly independent musicians, lack the legal or technical knowledge to implement opt-out mechanisms.
• Technical barriers prevent many independent artists from successfully using machine-readable opt-outs. Smaller rights holders often lack the resources to implement complex technological solutions to protect their works.
• AI developers routinely ignore opt-outs: GEMA's lawsuit against OpenAI and Suno alleges that copyrighted works were used without consent, despite opt-outs being in place (https://www.gema.de/en/news/ai-and-music/ai-lawsuit). In The 'Intercept Media v. OpenAI', the plaintiff claims that OpenAI stripped copyright management information from articles before using them for model training (https://www.courtlistener.com/docket/68290804/the-intercept-media-inc-v-openai-inc/). This highlighting the lengths to which some AI companies go to evade responsibility.
• There is no evidence that the EU opt-out model has led to meaningful licensing opportunities for creators. Instead, AI developers exploit legal loopholes to avoid licensing altogether, undermining the economic framework that supports the creative industries.

If the UK adopts a similar opt-out system, it will only replicate these failures, leaving music creators vulnerable and without meaningful control over their works.

3. Opt-Out Fails the Government's Objectives

The government's stated objectives for AI and copyright policy include ensuring creator control, promoting fair licensing and remuneration, and maintaining transparency. Option 3 fails on all counts.

• Control
Opt-out systems strip creators of control by allowing AI developers to use their works without prior consent. The burden falls on creators to monitor the use of their works and opt out—a near-impossible task given the scale of data scraping. This creates an unfair power imbalance where tech companies dictate how copyrighted works are used.

• Licensing:
By permitting default access to copyrighted content, Option 3 undermines the licensing framework that supports the creative industries. AI developers can bypass licensing entirely, reducing opportunities for rights holders to negotiate fair terms and fees.

• Fair Remuneration:
Without mandatory licensing obligations, AI developers exploit copyrighted music without compensating the artists who created it. This undermines the economic foundation of the music industry, leading to lost income for musicians.
• Transparency:
Voluntary transparency requirements are inadequate. The Cisco AI Readiness Index (UK) shows that 66% of AI firms lack basic data tracking, making it nearly impossible to verify how copyrighted works are used. Without enforceable transparency obligations, creators cannot monitor AI training practices or enforce their rights.

4. AI Developers Can and Should Pay for Licensed Content

The AI industry operates under an unfair expectation: that high-quality copyrighted content should be accessible for free, despite the vast sums invested in computing power, data storage, and engineering talent. This business model exploits creators while allowing AI developers to profit from their works.

Reports from APRA AMCOS (https://www.apraamcos.com.au/about/supporting-the-industry/research-papers/aiandmusic), GEMA/SACEM (https://www.gema.de/en/news/ai-study), and CISAC

(https://www.cisac.org/Newsroom/news-releases/global-economic-study-shows-human-creators-future-risk-generative-ai) confirm that:

• Licensing is feasible. Many AI developers already have licensing agreements with music publishers, record labels, and collecting societies. The argument that licensing is impractical is unfounded.
• AI-generated music directly competes with human-made music in streaming, advertising, and sync licensing markets, leading to significant income losses for musicians.
• The CISAC report predicts that unlicensed AI training could result in the music industry losing up to 25% of its revenue within four years. Without clear licensing frameworks, the economic sustainability of the creative industries is at risk.

If AI developers can pay for cloud infrastructure and technical talent, they can and should pay for the data that gives their models value.

5. Transparency Alone Is Not Enough

While transparency is essential, it cannot serve as a substitute for proper licensing. Transparency alone does not prevent infringement—it only reveals it after the fact. Without clear legal obligations requiring AI developers to obtain licences before using copyrighted content, transparency measures will be ineffective.

The failure of voluntary self-regulation is well-documented. The EU Code of Practice on Disinformation demonstrated that tech companies routinely evade transparency obligations, often citing trade secrets to avoid disclosure. Without binding legal requirements, there is little incentive for AI developers to be forthcoming about their data practices.

Transparency measures must be legally enforced and paired with mandatory licensing requirements. Only then can creators effectively monitor how their works are used and ensure they receive fair remuneration.

Conclusion: Why the Opt-Out Model Must Be Rejected

Option 3 fails to meet the government's stated objectives. It undermines the fundamental principles of copyright law, strips creators of control, and exposes the UK's creative industries to economic harm.

• Control: Opt-out denies creators meaningful control over how their works are used.
• Licensing: It allows AI developers to bypass licensing frameworks, undermining the music industry's economic structure.
• Fair Remuneration: It enables AI developers to use copyrighted works without paying for them.
• Transparency: Voluntary measures are ineffective and will not hold AI developers accountable.

The only way to safeguard the rights of creators and ensure the long-term sustainability of the UK's creative industries is to adopt a permissions-based system, where AI developers are legally required to seek explicit consent before using copyrighted works in training datasets.

The UK has a thriving creative economy, contributing £124.6 billion to GVA and supporting 2.4 million jobs. Weakening copyright protections through an opt-out exception would threaten this success and undermine the rights of the very creators who fuel the industry.

The UK must reject Option 3 and instead strengthen copyright protections, ensuring that music creators are fairly compensated, their rights are respected, and the integrity of the creative economy is preserved in the age of AI.

7 If so, what aspects do you consider to be the most important?

Please give us your views.:

8 If not, what other approach do you propose and how would that achieve the intended balance of objectives?

Please give us further comments.:

The UK's existing copyright law and licensing frameworks are already fit for purpose. Rather than introducing new exceptions, the focus should be on strengthening and enforcing current mechanisms to ensure that AI developers respect creators' rights.

I propose a twofold approach that both reinforces existing licensing frameworks and utilises Section 9(3) of the Copyright, Designs and Patents Act 1988 (CDPA), the protection of computer-generated works, to create a balanced, sustainable system that protects human creators while accommodating AI innovation.

1. Strengthening Existing Licensing Frameworks

At its core, UK copyright law already provides the necessary tools to regulate the use of copyrighted music in AI training. AI developers should be required to obtain explicit consent before using copyrighted works, just as they do when licensing music for streaming, radio, advertising, and public performance. Licensing frameworks, managed by Collective Management Organisations (CMOs) like PRS for Music and PPL, ensure that rights holders are compensated fairly when their works are used. These systems are robust, functional, and have been successfully adapted to new technologies in the past.

Training data is a core resource for AI developers - equally as essential as computing power and engineering expertise, both of which they invest heavily in. There is no valid justification for allowing AI companies to exploit copyrighted works for free, especially when generative AI systems directly compete with human creators by producing music that can displace human-made content in key revenue streams, such as streaming, sync licensing, and advertising.

To ensure fairness, AI developers should:

• Obtain licenses for copyrighted content used in training datasets.
• Pay upfront licensing fees through CMOs, similar to the PPL Broadcast Licence or TheMusicLicence for public performances.
• Disclose datasets to ensure transparency and allow rights holders to monitor how their works are used.

This licensing-first approach would guarantee that music creators are fairly remunerated while enabling AI developers to access the data they need legally and sustainably.

2. Utilising Section 9(3) CDPA: A Framework for Fair Remuneration

A complementary approach would leverage Section 9(3) of the CDPA, which uniquely positions the UK to address the complexities of AI-generated works. Section 9(3) assigns authorship of computer-generated works (CGWs) to the person who makes the "arrangements necessary" for their creation. This offers a legal basis for attributing ownership and responsibility in cases where no human author directly creates the output.

By interpreting Section 9(3) as a mechanism to establish related rights for AI-generated works, the UK could implement a licensing and royalty-sharing framework that ensures fair remuneration for all stakeholders:

• AI developers would be assigned a portion of the rights to AI-generated outputs under Section 9(3) but would also be responsible for licensing and compensating the human-created works used to train their models.
• A two-tier licensing system, as proposed in my PhD thesis, could be established:
• Collective Licence (for general data usage): An annual licence fee (e.g., TheAILicence) would be paid by AI companies to access large datasets, with royalties distributed to rights holders through CMOs.
• Custom Licence (for curated datasets): Bespoke datasets that focus on specific artists or genres would incur higher licensing fees, directly benefiting identified rights holders.

• Royalties generated by AI-generated works could be distributed using a joint ownership model:
• AI developers (for their investment and technical expertise) (e.g., 33%).
• Human creators (whose works were used in training datasets) (e.g., 33%).
• AI users (who contribute to the creative process through prompts and refinements) (e.g., 34%)

This system aligns with how sound recordings are treated under related rights and would ensure that all contributors—human and non-human—are fairly compensated.

3. Addressing Transparency and Accountability

Transparency obligations are essential for any framework to function effectively. AI developers must be legally required to:

• Disclose training datasets to CMOs, enabling proper licensing and royalty distribution.
• Label AI-generated outputs to distinguish them from human-created content, ensuring that consumers, rights holders, and industry stakeholders can make informed decisions.
• Implement traceability measures to allow creators to track when and how their works have been used.

Transparency would also facilitate retrospective licensing, an issue identified in my thesis as a significant challenge. AI developers who have already trained models on unlicensed data should be required to backdate licensing agreements, ensuring that rights holders are compensated for past use.

4. Why This Approach Achieves Balance

This proposed approach achieves the government's objectives of balancing innovation with the protection of creators' rights in several ways:

• Control: Creators retain control over how their works are used, as AI developers must obtain explicit permission through licensing.
• Fair Remuneration: Rights holders receive fair compensation through upfront licensing fees and ongoing royalty payments, ensuring they benefit economically from the use of their works in AI training.
• Licensing Framework: Utilising existing CMOs and adapting Section 9(3) offers a practical pathway for managing AI-generated works within the current legal system.
• Transparency: Legally binding obligations for dataset disclosure and labelling ensure accountability and traceability throughout the AI development process.

This framework avoids the pitfalls of opt-out models - which undermine copyright's core principles - and broad TDM exceptions, which unfairly privilege AI developers at the expense of creators. Instead, it creates a balanced ecosystem where AI innovation can flourish without exploiting human creativity.
Conclusion

Rather than introducing new exceptions that undermine copyright's core principles, the UK should focus on strengthening its existing legal framework, while leveraging the unique opportunities offered by Section 9(3). This approach fosters innovation while ensuring that music creators are fairly compensated, preserving the integrity and economic sustainability of the UK's creative industries.

By integrating licensing frameworks, transparency obligations, and the related rights model under Section 9(3), the UK has the opportunity to lead globally in developing fair, effective copyright protections that adapt to the challenges of the AI era.

9 What influence, positive or negative, would the introduction of an exception along these lines have on you or your organisation? Please provide quantitative information where possible.

Please give us your views.:

The introduction of a TDM exception with an opt-out mechanism would have a profoundly negative impact on the music industry, particularly on the livelihoods of music creators. As both an academic specialising in AI, copyright law, and the creative economy, and as a performing musician, I am deeply concerned about how such an exception would undermine existing copyright protections, disrupt established licensing frameworks, and create an unfair competitive landscape where human-made music is devalued by AI-generated content.

The implications of this policy shift would be felt across the industry, leading to revenue losses, weakened copyright enforcement, and a substantial erosion of musicians' ability to sustain viable careers. Below, I outline the core issues, supported by quantitative data and real-world examples.

1. Unfair Market Competition and Revenue Losses

AI-generated music poses a direct economic threat to human musicians. Generative AI systems can produce vast amounts of music quickly, cheaply, and at scale, all while relying on existing human-created works as their training data. This dynamic creates an unfair market where AI-generated music competes directly with the very works it was trained on, without offering fair remuneration to the original creators.

Allowing AI developers to train on copyrighted music without explicit permission or compensation risks decimating the creative industries. Human musicians, who already operate in an economically precarious environment, will face increasing challenges as AI-generated content floods the market, further reducing their share of revenue.

Quantitative Evidence:

The CISAC report (https://www.cisac.org/Newsroom/news-releases/global-economic-study-shows-human-creators-future-risk-generative-ai) projects that AI-generated music could replace up to 25% of human-created music revenues within four years, representing a significant income loss for creators.

• APRA AMCOS (https://www.apraamcos.com.au/about/supporting-the-industry/research-papers/aiandmusic) and GEMA/SACEM (https://www.gema.de/en/news/ai-study) reports corroborate these findings, highlighting that AI-generated content is already displacing human-made music in key revenue streams such as streaming, advertising, and sync licensing.
• Beatdapp, a music tracking firm, estimates that streaming fraud diverts between $2 billion and $3 billion annually from the global music industry (https://news.sky.com/story/fraud-gangs-stealing-billions-from-music-industry-via-fake-streams-13163016). Much of this fraud involves AI-generated tracks combined with bots to inflate streaming numbers.
• The Michael Smith case (https://variety.com/2024/music/news/ai-songs-spotify-apple-michael-smith-fraud-charged-artificial-intelligence-1236130680/) exemplifies the scale of potential abuse, where AI-generated tracks combined with bot streaming fraudulently generated over $10 million in royalties, highlighting systemic vulnerabilities in the streaming ecosystem.

These examples underscore how AI-generated music not only competes with but actively undermines the economic sustainability of human creators.

2. Undermining Licensing Frameworks

Licensing is the foundation of the music industry. It ensures that creators are compensated when their works are used and allows for sustainable revenue generation across various platforms, including streaming services, broadcasting, and live performances.

A TDM exception that allows AI developers to bypass licensing requirements fundamentally threatens this system. It devalues music as a commodity by enabling companies to exploit it as free training data, without compensating the original creators. This approach would have severe consequences, not only for individual musicians but also for the broader economic framework that supports the music industry.

AI-generated music, trained on unlicensed datasets, will directly compete with human-created works, reducing licensing opportunities and eroding revenue streams for rights holders. This undercuts the very purpose of copyright law, which exists to protect creators and ensure they can earn a living from their work.

Quantitative Evidence:

• In 2023, the UK music industry contributed £6.7 billion to the economy (https://www.ukmusic.org/wp-content/uploads/2024/06/This-Is-Music-2023-Economic-Report.pdf) and supported over 311,000 jobs (https://lordslibrary.parliament.uk/creative-industries-growth-jobs-and-productivity/).

Undermining licensing frameworks would jeopardise this contribution, leading to job losses and reduced investment in new music and talent. The music industry's economic success depends on a strong licensing system. Allowing AI developers to bypass these frameworks could significantly disrupt this balance, with long-term consequences for creators and the wider industry.

3. The Opt-Out Model is Impractical and Ineffective

While the idea of an opt-out system may appear to offer creators some level of control, in practice, it is ineffective and unworkable, especially in the context of music. Most musicians distribute their work widely across multiple platforms—streaming services, social media, websites, and more. Expecting individual creators to track the use of their works across the entire internet and manually opt out on each platform is unrealistic, particularly for independent musicians who lack the legal and technical resources of larger rights holders.

Practical Concerns:

• Opt-out models fail to address the irreversible nature of AI training. Once an AI model has been trained on a dataset, it cannot "unlearn" the data, even if the creator opts out later.
• This means that AI-generated outputs will continue to be influenced by the original works indefinitely, without permission or compensation.
• The burden of enforcement falls entirely on the creator, who must not only opt out but also attempt to police how their works are used—an impossible task at the scale at which AI systems operate.

In effect, an opt-out system offers creators the illusion of control while doing little to protect their rights in practice.

4. Misleading "Lawful Access" Arguments

AI companies often claim that they will only use "legally accessible" material under a TDM exception for commercial purposes. However, this argument is deeply misleading.

The fact that music is publicly available online—on platforms like Spotify, YouTube, SoundCloud, or even personal websites—does not equate to permission for it to be used in AI training. There is a critical distinction between public availability and the public domain. Just because a work can be accessed online does not mean it is free to use for any purpose, especially when it comes to commercial exploitation.

Key Concerns:
• Once music has been scraped and used to train an AI model, there are no safeguards to ensure that the influence of that work is removed from the system. Even if the use of the data is later deemed unauthorised, the AI's outputs will still carry traces of the original music.
• The Information Commissioner's Office (ICO) has also raised concerns regarding GDPR violations tied to AI training (https://www.arnoldporter.com/en/perspectives/advisories/2024/02/ico-launches-consultation-series). "Invisible processing" occurs when personal metadata attached to musical works—such as artist names, licensing information, and other identifiers—is scraped without consent. This process may breach data protection laws, particularly when creators are unaware that their personal data has been used in this way.

These issues highlight the legal and ethical risks of relying on the "lawful access" argument to justify the scraping of copyrighted works.

Conclusion: The Risk to Music Creators and the Broader Industry

The introduction of a TDM exception with an opt-out system would have far-reaching negative consequences for music creators, the broader music industry, and the cultural sector as a whole.

It would:

• Undermine copyright protections, allowing AI developers to exploit creative works without permission or fair compensation.
• Undermine licensing frameworks, leading to revenue losses and reduced economic contributions from the music industry.
• Create unfair competition between AI-generated content and human-created works, further marginalising musicians.
• Expose creators to legal and data protection risks, particularly around metadata scraping and GDPR compliance.

From both an academic and practitioner's perspective, this policy would destabilise the creative industries and threaten the livelihoods of thousands of music creators. It would shift the power dynamics further in favour of large tech companies, at the expense of the artists who generate the content that fuels both AI systems and the broader cultural economy.

The UK government must prioritise the protection of music creators by rejecting opt-out-based TDM exceptions and instead reinforcing existing licensing frameworks. This will ensure that AI innovation can continue while respecting the rights and contributions of human creators, preserving the cultural and economic value that the music industry brings to society.

10 What action should a developer take when a reservation has been applied to a copy of a work?

Please give us your views.:

The proposed opt-out system is an unworkable solution. However, should a reservation be applied to a copy of a work, AI developers should legally be required to remove that work from training datasets and retrain their models from scratch without those works. This is the only way AI companies can ensure that works with reservations applied are not used in the development of their models.

11 What should be the legal consequences if a reservation is ignored?

Please give us your views.:

The proposed opt-out system is unworkable and does not provide meaningful protection for music creators. However, if a rights reservation is applied to a work and ignored by an AI company, there must be serious legal consequences to ensure compliance. Penalties should include significant fines for non-compliance, a legal obligation to retrain AI models to remove the influence of copyrighted works that were used without permission, take-down notices requiring the removal of AI-generated content that unlawfully incorporates reserved works. These penalties should be extended to restriction of services in the UK and finally legal action if there continues to be non-compliance.

12 Do you agree that rights should be reserved in machine-readable formats? Where possible, please indicate what you anticipate the cost of introducing and/or complying with a rights reservation in machine-readable format would be.

Please give us your views.:

No. Requiring musicians to reserve their rights in machine-readable formats imposes an unrealistic, costly, and ineffective burden on individual creators, particularly independent musicians who lack the resources and technical expertise to comply.

As both an academic specialising in AI and copyright law and a performing musician, I see significant flaws in the proposal to rely on machine-readable rights reservations as a mechanism for protecting creators' works in the context of AI training. While the idea may appear to offer an efficient technical solution, in practice, it shifts the responsibility onto creators, who are often ill-equipped to implement or enforce such protections.

1. Unreasonable Burden on Musicians

The majority of musicians - especially independent artists - do not have the technical skills or legal expertise necessary to implement machine-readable rights reservations. Most creators focus on composing, performing, and distributing music, not on navigating complex digital rights management systems.

Implementing machine-readable opt-outs would require:

• Specialized software tools or services to tag works correctly across platforms.
• Ongoing legal or technical assistance to ensure that rights reservations are consistently applied and remain up to date.
• Manual tracking of where and how music is distributed across various platforms (e.g., streaming services, social media, personal websites), adding to the workload and cost burden for creators.

For many musicians, particularly those without major label backing or significant financial resources, these demands would be prohibitively expensive and time-consuming.

Anticipated Costs:

• Legal support for drafting and implementing machine-readable rights reservations could cost several hundred to thousands of pounds, depending on the complexity.
• Technical tools or services to automate the tagging process would likely involve subscription fees or licensing costs, adding ongoing financial pressure.
• Administrative time spent monitoring compliance across multiple platforms would divert time away from creative work, resulting in indirect economic losses.

For an industry where 62% of musicians earn less than £20,000 annually from their music activities (https://www.gov.uk/government/publications/music-creators-earnings-in-the-digital-era/executive-summary), these added costs are simply unsustainable.

2. Ineffectiveness of Machine-Readable Opt-Outs

Even if musicians overcome these challenges and successfully implement machine-readable rights reservations, there is no guarantee that AI developers will respect them. Historical examples show that machine-readable opt-outs are frequently ignored, bypassed, or circumvented by data scrapers.

The EU's DSM Directive (Article 4) (https://eur-lex.europa.eu/eli/dir/2019/790/oj/eng), which introduced a similar opt-out model, provides a cautionary tale:

• Rights holders remain unaware of the need to implement machine-readable opt-outs, leading to unauthorized scraping and AI training.
• Independent creators struggle with the technical complexity of machine-readable rights reservations, resulting in incomplete or inconsistent coverage.
• AI developers routinely ignore opt-outs, as evidenced by the GEMA lawsuit against OpenAI and Suno (https://www.gema.de/en/news/ai-and-music/ai-lawsuit), where copyrighted works were allegedly used without consent, despite opt-outs being in place.
• In 'The Intercept Media v. OpenAI', the plaintiff alleges that OpenAI stripped copyright management information from articles before using them in model training (https://www.courtlistener.com/docket/68290804/the-intercept-media-inc-v-openai-inc/). This highlights the risks of bad faith actors bypassing rights reservations.

These examples illustrate that machine-readable rights reservations, without robust legal obligations and strong enforcement mechanisms, offer little practical protection for creators.

3. Shifting the Burden onto Creators

At its core, the reliance on machine-readable rights reservations shifts the burden of enforcement from AI developers - who are profiting from the use of creative works - to the creators themselves. This approach contradicts the foundational principle of copyright law, which is based on prior permission, not retroactive enforcement.

Rather than requiring musicians to opt out of AI training by embedding machine-readable reservations, the onus should be on AI developers to seek explicit permission through licensing frameworks before using copyrighted content.

A permissions-based system:

• Ensures that creators retain control over how their works are used.
• Eliminates the need for individual musicians to navigate complex technical processes.

• Places responsibility on AI developers to respect copyright law and secure proper licenses.

This approach aligns with existing licensing practices in the music industry, where Collective Management Organisations (CMOs) like PRS for Music and PPL facilitate the licensing of works for streaming, broadcasting, and other uses.

4. Transparency and Enforcement Are Essential

If machine-readable rights reservations are to play any role, they must be supported by strong legal obligations and robust enforcement mechanisms:

• AI developers must be legally required to respect machine-readable opt-outs and face significant penalties for violations.
• Transparency measures must be in place, requiring AI developers to disclose their training datasets, enabling rights holders to verify whether their works have been used.
• Independent oversight bodies should be established to monitor compliance and handle disputes, ensuring that creators have access to effective remedies if their rights are infringed.

Without these safeguards, machine-readable rights reservations will remain an ineffective and burdensome solution that offers creators little meaningful protection.

Conclusion

While machine-readable rights reservations may seem like a technologically elegant solution, they are impractical, ineffective, and unfair in the context of AI training. They place an unreasonable burden on musicians - many of whom lack the technical expertise or financial resources to implement such systems - while offering little guarantee that AI developers will respect them.

The focus should instead be on:

• Strengthening licensing frameworks to ensure that AI developers obtain explicit permission before using copyrighted works.
• Implementing transparency and enforcement mechanisms to hold AI developers accountable.
• Shifting responsibility away from individual creators and onto AI companies, who profit from the use of copyrighted content.

Only by reinforcing the existing permission-based model of copyright law can the UK protect the rights and livelihoods of music creators while fostering innovation in the AI sector.

Technical Standards

13 Is there a need for greater standardisation of rights reservation protocols?

Please give us your views:

If an opt-out system were to be introduced, standardisation would be essential, but it would not solve the underlying issues. The core problem with opt-out systems is not the lack of standardisation but the flawed principle that places the burden on creators to protect their works, rather than on AI developers to seek permission.

I consider the call for standardisation an attempt to make an inherently flawed system appear workable. While standardised protocols could help ensure some level of consistency in how rights reservations are applied and interpreted, they would do little to protect creators in practice.

Key Issues with Standardisation in Opt-Out Systems:

1. Burden Still Falls on Creators:
Even with standardised protocols, the responsibility would still lie with musicians and rights holders to proactively implement rights reservations. This is particularly burdensome for independent creators, who often lack the technical expertise or financial resources to manage complex digital rights management (DRM) systems.

2. Ineffectiveness Against Large-Scale Scraping:
Standardisation does not address the core issue—that opt-out systems cannot effectively prevent copyrighted works from being scraped and used in AI training. Given the scale and complexity of online content distribution, even standardised opt-outs will be circumvented or ignored by AI developers, particularly in jurisdictions with weaker enforcement.

3. No Control Over Downstream Use:
Once a work is scraped and incorporated into an AI training dataset, even standardised opt-out protocols cannot prevent its downstream use. AI models cannot "unlearn" data, meaning that even opted-out works will continue to influence AI-generated outputs.

The Alternative: A Permissions-Based System

The only fair and effective approach is to maintain a permissions-based system, where AI developers are legally required to obtain explicit opt-in consent before using copyrighted works.

This approach:

• Respects the core principles of copyright law by ensuring that creators retain control over their works.
• Eliminates the burden on musicians to implement complex rights reservation systems.
• Ensures fair remuneration through upfront licensing agreements.

Standardisation should only be considered within a permissions-based system, where AI developers are responsible for respecting rights, not within an opt-out framework that places the burden on creators.

14 How can compliance with standards be encouraged?

Please give us your views.:

Compliance with rights reservation standards can only be encouraged through strong legal obligations, robust enforcement mechanisms, and clear accountability for AI developers. Without these, any standardised system - whether opt-in or opt-out - will fail to protect creators effectively.

1. Enforce Existing Copyright Laws First
The UK copyright framework already requires AI developers to seek permission before using copyrighted works. Instead of introducing new, complex standards, the government should focus on enforcing existing laws and ensuring that AI developers comply with established licensing frameworks.

2. Responsibility Must Lie with AI Developers
Musicians should not be responsible for ensuring compliance with rights reservation standards. The burden must rest entirely on AI developers, who should be required to:

• Detect and respect rights reservations at every stage of data collection, model training, and deployment.
• Maintain accurate records of the data used, including documentation of licenses and permissions.
• Implement proactive systems that prevent non-compliant data from being ingested into training datasets.

3. Legal Obligations and Enforcement
Legally binding obligations must be introduced to compel AI developers to respect rights reservations. Strong penalties for non-compliance - for example, significant financial sanctions and potential legal action - must be established to discourage misuse. Enforcement mechanisms should include:

• Regular audits of AI developers' practices.
• An independent regulatory body to oversee compliance and handle disputes.
• Transparency obligations requiring AI developers to disclose their training datasets and licensing practices.

4. Transparency as a Core Principle
Transparency is crucial for ensuring compliance. AI developers should be required to:

• Disclose data sources used in model training.
• Provide creators with access to clear information about how their works have been used.
• Implement labelling systems for AI-generated content, allowing consumers to distinguish between human-made and AI-generated works.

15 Should the government have a role in ensuring this and, if so, what should that be?

Please give us your views.:

Yes, the government must play a central and active role in ensuring AI developers comply with copyright protections and that rights reservation standards are upheld. Without robust legal oversight and meaningful enforcement mechanisms, there is little incentive for AI developers to respect the rights of creators or engage in fair licensing practices.

The existing UK copyright framework already requires AI developers to seek permission before using copyrighted works, but enforcement has been lacking. The government should focus on upholding these laws rather than introducing new exceptions that would further undermine creators' rights. A permissions-based system, where AI developers must obtain explicit consent before using copyrighted works, remains the most effective approach.

To ensure compliance, the government needs to establish clear legal obligations for AI developers. This includes mandatory licensing requirements, enforceable transparency obligations, and strict penalties for non-compliance. Without legal mandates, AI developers will continue to exploit loopholes and avoid licensing obligations, leading to further erosion of creators' rights and income.

An independent regulatory body should be created to oversee AI developers, monitor compliance, and conduct regular audits. This body would be responsible for investigating violations, imposing penalties, and ensuring transparency in AI training practices. AI developers should be legally required to disclose their training datasets and demonstrate that they have secured appropriate licenses for all copyrighted content used.

Beyond enforcement, the government should also facilitate industry-wide licensing solutions. This could involve creating centralised licensing platforms or collective licensing schemes that simplify the process for AI developers to access creative works legally while ensuring fair remuneration for creators. Such frameworks would reduce friction between rights holders and AI companies, promoting a more balanced and sustainable ecosystem.

Ultimately, the government has a duty to safeguard the creative industries, ensuring that AI innovation does not come at the expense of human creators. By prioritising strong legal protections, effective enforcement, and transparent licensing frameworks, the government can help maintain the integrity of the UK's copyright system while supporting responsible AI development.

Contracts and licensing

16 Does current practice relating to the licensing of copyright works for AI training meet the needs of creators and performers?

Please give us your views:

Yes, current licensing frameworks are fundamentally fit for purpose, but they require stronger enforcement, particularly regarding transparency. Existing UK copyright law, alongside the well-established music licensing infrastructure, already provides the necessary tools to ensure that creators and performers are fairly compensated when their works are used, including in AI training.

The success of the UK's licensing system is evident. In 2023, the UK music industry contributed £7.6 billion to the economy (https://www.prsformusic.com/m-magazine/news/uk-music-industry-contributes-7-billion-to-uk-economy-this-is-music-2024) and supported over 311,000 jobs (https://lordslibrary.parliament.uk/creative-industries-growth-jobs-and-productivity/), a testament to the robustness of licensing frameworks across streaming, broadcasting, and sync licensing. Collecting societies like PRS for Music and PPL effectively manage vast licensing operations, ensuring that royalties flow to the appropriate rights holders.

In 2023, PRS for Music reported record-breaking revenues (https://www.prsformusic.com/about-us/track-record/2023-financial-results), collecting £944 million, a 12.8% increase from the previous year. This growth was observed across multiple income streams:

- Public performance: revenue increased to £188.2 million, recovering from the pandemic's impact.
- Broadcast: collections remained stable at £110.7 million.
- Online: revenue surged to £360.3 million, reflecting the growing consumption of digital music.
- Synchronisation: generated £40 million (https://www.ukmusic.org/wp-content/uploads/2024/11/TIM-Report-2024-reduced.pdf)
- Secondary use: £1.5 to £2 million (https://musiciansunion.org.uk/news/2023-end-of-year-round-up)

These figures evidence that large-scale licensing frameworks can function smoothly, even in complex markets.

Importantly, a functioning market for AI training data already exists. Some AI developers have voluntarily entered into licensing agreements with music publishers, record labels, and collecting societies, showing that licensing AI training data is not only feasible but practical. Reports from APRA AMCOS (https://www.apraamcos.com.au/about/supporting-the-industry/research-papers/aiandmusic) confirm that such agreements are in place, proving that AI developers can operate within existing systems when required. This evidence undermines the argument that a new TDM exception is necessary.

The music industry has long demonstrated flexibility in licensing models, ranging from Creative Commons licenses (https://creativecommons.org/) to royalty-free samples and bespoke plugin licenses (https://www.lucidsamples.com/blog/what-is-sample-pack#Copyright_Concerns). AI-generated content can follow similar paths. Royalty-free licenses, for example, could allow AI-generated music to be used without ongoing payments but still require upfront fees. Alternatively, custom licenses could restrict AI-generated outputs to non-commercial use or impose royalties for commercial applications. Even levy-based models (https://www.tandfonline.com/doi/full/10.1080/17510694.2020.1839702), akin to those previously used for home taping and private copying, could be adapted. AI developers could contribute to a trust fund, which would distribute royalties to creators whose works were used in AI training, with technologies like blockchain (https://www.elgaronline.com/view/journals/qmjip/10-4/qmjip.2020.04.05.xml) ensuring transparency in royalty distribution.

Licensing underpins a sector worth billions and supports hundreds of thousands of jobs. There is no justification for introducing exceptions that would allow AI developers to bypass these frameworks. Instead, the government should focus on enforcing existing laws and ensuring that licensing obligations are respected. An opt-in, permissions-based system, supported by transparency obligations, will allow AI innovation and creator rights to thrive together.

17 Where possible, please indicate the revenue/cost that you or your organisation receives/pays per year for this licensing under current practice.

Please provide further evidence.:

18 Should measures be introduced to support licensing good practice?

Please give us your views:

Yes. The most effective way to support good licensing practices is through the robust enforcement of existing copyright law. The UK's current licensing framework already provides the mechanisms needed to ensure creators are fairly compensated, but stronger oversight is necessary to guarantee compliance. By reinforcing these systems and promoting transparency in AI training practices, the government can safeguard the integrity of the licensing ecosystem and ensure that creators' rights are protected.

19 Should the government have a role in encouraging collective licensing and/or data aggregation services?

Yes

20 If so, what role should it play?

Please provide further comments:

The government must focus on enforcing existing copyright law and supporting the established licensing frameworks within the music industry. By promoting collective licensing and data aggregation services, the government can simplify the process of obtaining licenses for AI training and reduce administrative burdens for both AI developers and rights holders. Furthermore, once an AI model has trained on a dataset, the music contained within

that dataset continue to contribute to the outputs of that model until it is retrained from scratch. Collective licences can be paid annually which will cover ongoing use.

In addition to facilitating licensing, the government should introduce strong transparency obligations. AI developers must be required to disclose the copyrighted works used in their training datasets. This level of transparency will make it significantly easier for rights holders and collecting societies to monitor the use of their works and ensure that appropriate licenses are in place.

Moreover, the government should help develop centralised licensing platforms that streamline the process for AI developers to secure permissions while ensuring that creators receive fair remuneration. These initiatives would foster a more balanced and sustainable ecosystem where AI innovation can thrive without undermining the rights of creators.

21 Are you aware of any individuals or bodies with specific licensing needs that should be taken into account?

Please give us your views.:

Yes. Copyright law is designed to protect all creators and give them control over how their work is used. Those human creators have specific licensing needs that must be taken into account. AI training should not be treated differently from any other commercial use of music—licensing must remain the standard.

A wide range of stakeholders in the music industry rely on licensing for their livelihoods, including music creators, performers, producers, publishers, record labels, and collecting societies. Each of these groups has specific licensing needs that must be considered. Independent musicians, smaller rights holders, and emerging artists are particularly vulnerable to market shifts and would be disproportionately harmed if AI developers were allowed to use music without permission or payment.

Maintaining a strong, transparent licensing framework is essential to ensure that all rights holders - from global music companies to individual creators - can continue to be fairly compensated when their works are used, including for AI training. By supporting these diverse stakeholders, the government can help sustain the UK's creative industries while fostering responsible AI innovation.

Transparency

22 Do you agree that AI developers should disclose the sources of their training material?

Please give us your views.:

Yes, transparency is non-negotiable and must be a legally binding and enforceable obligation for AI developers, irrespective of which copyright framework the government ultimately adopts. If AI developers are sourcing their training data ethically and legally, there should be no reason for them to conceal the contents of their datasets.

Without full disclosure, creators cannot know whether their work has been used without permission, making it impossible to enforce their rights or seek fair remuneration. Transparency is crucial not only for enabling fair licensing but also for holding AI developers accountable and building trust between the tech and creative sectors. Past experience has shown that voluntary transparency simply doesn't work—legally binding and enforceable requirements are essential to uphold the rights of creators.

23 If so, what level of granularity is sufficient and necessary for AI firms when providing transparency over the inputs to generative models?

Please provide further comments:

AI developers must provide comprehensive, publicly available records of all datasets and sources used in training their models. This should include detailed information on individual copyrighted works—such as artist names and track titles—where possible, so creators can verify whether their material has been used. Developers should also disclose the sources and owners of any datasets employed, as well as timestamps indicating when data was scraped or collected.

Equally important is clear documentation on how data was acquired, specifying whether licensing agreements or permissions were obtained. These transparency obligations must be ongoing; whenever AI models are updated or retrained with new data, the lists of training inputs should be updated accordingly and made publicly available. This ensures transparency is not a one-time event but an enduring commitment.

24 What transparency should be required in relation to web crawlers?

Please give us your views.:

Operators of web crawlers must be held to stringent transparency standards. They should clearly identify themselves and state their purpose when scraping content. It is critical that they respect copyright protections and any opt-out mechanisms that rights holders have implemented. Furthermore, web crawler operators should maintain and publish public logs of all data scraped, allowing creators and rights holders to verify whether their works have been collected.

Opaque data scraping practices undermine copyright protections and make it impossible for creators to safeguard their works. Without strict transparency requirements, rights holders have no meaningful way to challenge the unauthorized use of their content.

25 What is a proportionate approach to ensuring appropriate transparency?

Please give us your views.:

Transparency must be a legal obligation, not a voluntary best practice. To ensure proportionality and effectiveness, transparency standards should be legally binding, with clear regulatory frameworks and oversight. An independent regulatory body should be established to monitor compliance and enforce transparency obligations.

AI developers should be required to disclose training data before deploying their models, providing full documentation of all sources used. This pre-emptive approach would prevent developers from releasing AI systems that have been trained on unauthorized data. It would also establish a culture of accountability, making it clear that transparency is a foundational requirement, not an afterthought.

This approach balances the interests of creators, AI developers, and the public while ensuring that the burden of transparency falls on those profiting from AI models, not on the creators whose works are being exploited.

26 Where possible, please indicate what you anticipate the costs of introducing transparency measures on AI developers would be.

Please indicate the anticipated costs of transparency measures.:

The cost of implementing transparency measures should be considered a standard compliance expense for AI developers, much like licensing fees for streaming services or broadcasting rights. Given the scale of investments AI developers already make in computing power, engineering talent, and data acquisition, the additional cost of ensuring transparency is proportionate and justified.

Importantly, the burden of transparency must fall on AI developers - not on creators through opt-out mechanisms - especially since AI companies' business models often rely on copyrighted material. Transparency is a basic requirement for ethical and legal AI development, and its costs should be treated as part of the normal cost of doing business.

27 How can compliance with transparency requirements be encouraged, and does this require regulatory underpinning?

Please give us your views.:

Yes, strong legal and regulatory frameworks are essential to ensure compliance with transparency obligations. To encourage and enforce transparency, the government should make transparency a legally binding requirement for AI developers. Establishing an independent regulatory body to monitor and enforce compliance is crucial. This body should have the authority to conduct regular audits of AI developers' training data, ensuring that proper documentation and licenses are in place.

Penalties for non-compliance should be significant enough to act as a deterrent. This could include substantial fines, usage bans on AI models trained without proper disclosures, or legal action for persistent violations. A public reporting system would also help, allowing creators to challenge unauthorized use or undisclosed data scraping. Without such regulatory oversight, AI developers will have little incentive to operate transparently or respect copyright law.

28 What are your views on the EU's approach to transparency?

Please give us your views.:

The EU AI Act's proposed Code of Practice (https://digital-strategy.ec.europa.eu/en/policies/ai-code-practice) is a positive step but falls short in several critical areas. While it rightly mandates that general-purpose AI model providers produce technical documentation and summaries of their training data, the emphasis on "summaries" leaves too much room for vague or incomplete disclosures. Full datasets—not summaries—should be disclosed to ensure accountability.

Another shortcoming is the lack of strong enforcement mechanisms within the EU's approach. Without clear penalties for non-compliance and a robust framework for monitoring AI developers, the risk remains that transparency obligations will be sidestepped or diluted.

The UK has an opportunity to learn from these gaps and take a more rigorous approach to transparency. Legally binding and enforceable transparency requirements, backed by strong regulatory oversight and penalties for non-compliance, will ensure that AI developers cannot hide behind ambiguous disclosures or voluntary standards.

Regardless of which copyright framework the government adopts, transparency must be legally mandated, enforceable, and monitored by an independent body. Without these safeguards, rights holders will remain powerless to protect their work, enforce their rights, or seek fair remuneration. The UK has the opportunity to lead in ethical AI governance. Strong, enforceable transparency laws are the foundation for that leadership and essential to ensuring that creators can maintain control over their works in an increasingly AI-driven world.

Wider clarification of copyright law

29 What steps can the government take to encourage AI developers to train their models in the UK and in accordance with UK law to ensure that the rights of right holders are respected?

Please give us your views:

To attract AI developers to train their models in the UK while ensuring that the rights of rights holders are respected, the government should establish a clear, enforceable legal framework that balances the needs of AI innovation with robust protections for creators. A transparent, well-regulated environment will not only safeguard creators' rights but also offer legal certainty for AI developers, making the UK an attractive jurisdiction for ethical AI

development.

The government should focus on three key areas:

1. Strengthen and Enforce Copyright Protections:
The government must ensure that existing copyright laws are enforced, requiring AI developers to obtain explicit licenses before using copyrighted materials for training. A permissions-based system, where rights holders control how their works are used, offers both legal clarity and protection for creators.

2. Encourage Streamlined Licensing Pathways:
Simplifying the licensing process will encourage AI developers to operate within UK law. The government can support the development of centralised licensing platforms or collective licensing schemes that make it easier for AI developers to acquire legal access to training data. These platforms should facilitate negotiations between rights holders and AI companies while ensuring fair remuneration for creators.

3. Offer Incentives for Ethical AI Development:
The UK could provide financial incentives, such as tax reliefs or grants, to AI developers who commit to transparent, ethical practices and adhere to licensing requirements. These incentives would encourage developers to train models domestically while respecting copyright protections.

By combining clear legal obligations with supportive infrastructure and incentives, the UK can position itself as a leader in ethical AI development while safeguarding the interests of creators.

30 To what extent does the copyright status of AI models trained outside the UK require clarification to ensure fairness for AI developers and right holders?

Please give us your views:

The copyright status of AI models trained outside the UK urgently requires clarification to prevent jurisdictional loopholes that undermine the rights of UK-based creators. Without clear guidance, AI developers could bypass UK copyright laws by training models in jurisdictions with weaker protections and then commercialising those models in the UK, effectively exploiting UK creators without accountability.

The government should establish that any AI model used or commercialised within the UK is subject to UK copyright law, regardless of where it was trained. This approach would align with the principle that the act of distribution and commercialisation is what triggers legal obligations, ensuring that rights holders are protected even if their works were scraped or used abroad.

Clarifying this point would also create legal certainty for AI developers, who would understand the requirements for entering the UK market. It would prevent unfair competition by ensuring that all developers - whether domestic or international - operate under the same legal standards when distributing AI-generated outputs in the UK.

In the future, the government should advocate for international harmonisation of copyright laws in collaboration with the World Intellectual Property Organization (WIPO). The most effective way to address the copyright status of AI models trained outside the UK is to ensure that member countries of international copyright treaties adopt consistent legal frameworks. Harmonising legal approaches to AI at an international level is essential to prevent jurisdictional loopholes, protect the rights of creators globally, and provide legal clarity for AI developers operating across borders.

31 Does the temporary copies exception require clarification in relation to AI training?

Please give us your views:

Yes, the temporary copies exception requires urgent clarification in the context of AI training. Originally intended to cover transient, incidental copies made during lawful use (such as caching or streaming), the exception is now being exploited by AI developers to justify large-scale data scraping and copying during the training process.

The LAION v. Kneschke case in Germany (https://www.euipo.europa.eu/en/law/recent-case-law/germany-hamburg-district-court-310-o-22723-laion-v-robert-kneschke) highlighted this issue, where the Hamburg District Court ruled that the temporary copies created during the curation of the LAION dataset were permissible under Germany's temporary copies exception for non-commercial scientific research. However, the court made it clear that this exception applied only to the creation of the dataset and did not extend to the commercial use of AI models trained on that data. This distinction is crucial, as the large-scale duplication, storage, and use of entire datasets for commercial AI training goes far beyond the narrow, transient purpose originally intended by the temporary copies exception.

This legal grey area allows AI developers to argue that their large-scale scraping activities are lawful, even when they directly infringe upon copyrighted works. Without clarification, this exception risks undermining the core principles of copyright law by enabling mass exploitation of creative works without permission or remuneration.

32 If so, how could this be done in a way that does not undermine the intended purpose of this exception?

Please provide further comments:

Clarification of the temporary copies exception should focus on reaffirming its original intent - to facilitate lawful, transient uses - while explicitly excluding large-scale, non-transient copying for AI training purposes.

The government could amend the exception to include specific language that limits its application to copies that are:

• Incidental and transient: The copy must be purely technical, created as part of a lawful process (e.g., streaming or caching) and automatically deleted afterward.
• Non-exploitative: Copies made under the exception should not be used for commercial purposes or as part of processes that result in derivative works, such as AI model training.

To further clarify, the legislation could state that copies made during AI training, which involve the systematic scraping, duplication, and storage of copyrighted works for analysis and model development, fall outside the scope of the temporary copies exception. This would close the loophole exploited by some AI developers while preserving the original purpose of the exception for lawful, everyday digital processes.

This approach ensures that the exception continues to facilitate legitimate uses without opening the door to mass infringement under the guise of technical necessity.

Encouraging research and innovation

33 Does the existing data mining exception for non-commercial research remain fit for purpose?

Please give us your views:

Yes, the existing data mining exception for non-commercial research remains fit for purpose when used as originally intended—to facilitate research that serves the public good, such as in areas like public health, climate change, or scientific discovery. This exception is crucial for enabling academic institutions, researchers, and non-profit organizations to access large datasets for non-commercial research that benefits society.

However, the TDM exception for non-commercial purposes has effectively become a loophole for AI companies to acquire vast amounts of copyrighted content without permission. AI companies already avoid paying for training data by scraping copyrighted works under the non-commercial research exception, only to later commercialise their models and outputs without compensating rights holders. In 'LAION v Robert Kneschke' (https://www.euipo.europa.eu/en/law/recent-case-law/germany-hamburg-district-court-310-o-22723-laion-v-robert-kneschke), the Regional Court of Hamburg upheld the use of non-commercial TDM for dataset creation and ruled that the downstream use of LAION's freely available datasets was irrelevant to legal assessment, exposing this significant loophole. This case exposes the exploitation of the non-commercial TDM exception, which creates an unfair, unregulated marketplace where AI developers gain an unfair advantage over music creators.

To address this, clearer legal definitions and stricter enforcement are needed to ensure that the non-commercial research exception is used solely for its intended purpose. One potential solution is to implement legally binding transparency obligations that require AI developers to declare the intended use of data at the outset of training, along with mechanisms to prevent the downstream commercialisation of models trained under this exception without securing the necessary licenses.

34 Should copyright rules relating to AI consider factors such as the purpose of an AI model, or the size of an AI firm?

Please give us your views:

In principle, copyright rules should be consistent and not depend on the purpose of the AI model or the size of the company using it. A license is a license, and the value of copyrighted material remains the same regardless of whether a small startup or a major corporation is using it.

That said, if a collective licensing system were introduced for AI training - similar to existing models in the music industry - then the cost structure could be scaled based on the size and/or revenue of the AI firm, much like TheMusicLicence operates for businesses that publicly communicate music. This approach would make licensing more accessible to smaller developers while ensuring that larger, well-funded companies pay proportionally higher fees, reflecting their greater commercial potential and resource use.

Such a model would promote fairness, ensuring that all AI developers - regardless of size - contribute to compensating creators while lowering the financial barrier for smaller firms and non-commercial projects. This would also encourage broader participation in ethical AI development without compromising the rights and remuneration of content creators.

CGW Policy Option 0: No legal change, maintain the current provisions

35 Are you in favour of maintaining current protection for computer-generated works? If yes, please explain whether and how you currently rely on this provision.

Yes

Please give us your views:

Yes, the current protection for computer-generated works (CGWs) under Section 9(3) of the Copyright, Designs and Patents Act 1988 (CDPA) should be maintained - at least for now. While this provision has not been widely tested in court, this likely stems from a lack of clarity around its application, particularly in the context of modern AI-generated works, rather than any inherent flaw in the legislation. Maintaining Section 9(3) poses no immediate harm and could offer valuable flexibility as AI technologies continue to evolve.

UK copyright law has traditionally focused on protecting works resulting from human intellectual creation, but Section 9(3) provides a useful degree of adaptability by attributing authorship of CGWs to the person who made the "arrangements necessary" for their creation. Though this phrasing is

admittedly vague, it presents an opportunity for the provision to be interpreted as offering a form of related rights, distinct from traditional copyright. This interpretation could serve several key functions in the evolving landscape of AI-generated content.

Firstly, maintaining Section 9(3) could incentivize responsible licensing practices. If AI developers understand that CGWs can receive legal protection under UK law, they may be more motivated to license copyrighted works for training datasets, rather than relying on unauthorised scraping. This could help foster a fairer ecosystem where both AI developers and content creators' benefit.

Secondly, there's potential for future frameworks to use Section 9(3) as a mechanism to reward original creators. A share of royalties generated by AI-created outputs could be directed back to the artists, musicians, or writers whose works were used to train the AI models, ensuring a more equitable distribution of value along the content supply chain.

Thirdly, maintaining legal protection for CGWs could support innovation. Users and companies would be more inclined to adopt AI tools if they knew the outputs they generated would enjoy legal protection, thereby promoting the responsible use of AI technologies. It would also help maintain clear traceability and accountability in cases of copyright infringement.

Removing Section 9(3) would introduce several risks. It could create significant legal uncertainty around AI-generated works, potentially pushing them into the public domain and leading to market saturation with unowned, unregulated content. This lack of clarity could also open the door to copyfraud, with bad actors falsely claiming authorship over AI-generated works that fall into the public domain. Additionally, eliminating the provision could limit licensing opportunities by discouraging collaboration between rights holders and AI developers, who may be less motivated to work together if AI outputs lack clear legal protection.

When Section 9(3) was introduced, Lord Young of Graffham described it as "the first copyright legislation anywhere in the world which attempts to deal specifically with the advent of artificial intelligence." Given that foresight, removing the provision now would be premature. Instead, the government should focus on clarifying Section 9(3) to reflect the realities of today's AI landscape while promoting fair remuneration and responsible licensing practices.

For instance, one area that requires clarification is the threshold for "substantial human input" in AI-assisted works. Establishing a clear standard would help differentiate between works that qualify for traditional copyright protection and those that fall under Section 9(3), reducing ambiguity for creators and AI developers alike.

Rather than eliminating this legal tool, the government has an opportunity to refine and modernise it, positioning the UK as a global leader in AI copyright regulation—a step that could provide long-term benefits for both the creative and technology sectors.

36 Do you have views on how the provision should be interpreted?

Please give us your views:

Yes. Section 9(3) of the Copyright, Designs and Patents Act 1988 (CDPA) should be interpreted as a related right, distinct from traditional copyright, aligning with the fundamental principles of copyright law - primarily to protect and reward human creativity - while also accommodating the realities of AI-generated content.

1. Distinguishing AI-Generated and AI-Assisted Works

Section 9(3) should draw a clear boundary between fully AI-generated works and AI-assisted works.

• AI-generated works, produced without meaningful human input, should fall under Section 9(3) as a related right. Works generated by text-to-music platforms such as Suno and Udio do not meet the threshold of human intellectual creativity and thus should not receive the same level of protection as human-created works.
• AI-assisted works, where human creators use AI as a tool in the creative process, such as a plugin integrated into Digital Audio Workstations (DAWs) should continue to be protected under traditional copyright law, with the human creator recognized as the author.

This distinction maintains the integrity of copyright law by ensuring that works reflecting human creativity receive appropriate protection and rewards.

2. Clarifying "Arrangements Necessary"

A central element in interpreting Section 9(3) as a related right is the term "arrangements necessary" regarding ownership of AI-generated content. "Arrangements necessary" should be clarified to reflect the complex nature of AI-generated works. Based on a joint ownership model, this term could include:

• AI developers who design and refine the algorithms, train the AI models and supervise the AI's output, contributing to the nature and characteristics of the generated outputs.
• Users who input prompts and parameters, and/or directly influence the outcome of the AI system.
• Human creators whose works are used in training datasets, recognizing their indirect but critical contribution to the AI's outputs.

This broader interpretation ensures that ownership attribution reflects all parties who significantly contribute to the creation of AI-generated works. A joint ownership model would fairly remunerate the human authors whose works were used to train the AI system, while also incentivising users to engage with AI platforms, boosting subscription income for developers. At the same time, it ensures that AI developers are financially compensated for their role in enabling creative output, allowing them to reinvest in licensing more human content for future training. This model also promotes accountability and traceability, supporting copyright enforcement and enabling action against infringement in AI-generated outputs.

3. Implementing a Related Right Framework

As proposed in my thesis, Section 9(3) should classify AI-generated works as related rights - similar to existing protections for sound recordings or broadcasts - rather than granting them full copyright. This approach offers several key benefits:

• Shortened protection terms (potentially less than the current 50 years) would help prevent market oversaturation and reduce the risk of a volatile creative economy flooded with AI-generated content.
• Since related rights provisions already exist within the CDPA (e.g., for sound recordings and broadcasts), this approach would avoid the need for significant reform. Instead, it would focus on clarifying Section 9(3), ensuring that the legal framework evolves in line with technological advancements while maintaining consistency with established copyright principles.

4. Supporting Licensing and Fair Remuneration

Section 9(3) should also be interpreted to promote fairer economic participation for human creators. As outlined in my thesis, a two-tier licensing and royalty-sharing framework could ensure that all stakeholders are fairly compensated. Here's how this would work:

• AI developers would pay upfront licensing fees to use copyrighted works in training datasets. These fees would be managed through existing Collective Management Organizations (CMOs) such as PRS for Music or PPL. CMOs have long been responsible for efficiently managing rights and distributing royalties in the music industry, and their existing infrastructure makes them ideally suited to handle AI training licenses.
• CMOs could facilitate licensing by administering collective licenses for AI developers, granting legal access to large datasets of copyrighted works for training purposes. AI developers would pay standardised fees, which CMOs would then distribute among rights holders based on usage data, ensuring transparency and fairness.
• Royalties generated from AI outputs would be partially redirected to the original creators whose works were used to train the AI models. This system would establish a sustainable value chain where both AI developers and human creators benefit.
• Joint ownership models would formalize the distribution of economic rights, with human creators, AI developers, and users sharing revenue based on their respective contributions to the AI-generated output.

Real-World Example: Licensing and Royalty Sharing in Practice

Consider an AI-generated song produced by a text-to-music platform like Suno, which uses a dataset partially sourced from works licensed through PRS for Music. Under this framework:

The AI developer pays an upfront licensing fee to PRS, granting legal access to the musical works used for training.

When a user generates a new track using the platform, any revenue from commercializing the AI-generated song (e.g., streaming royalties or licensing for ads) would be shared among:

• The AI developer (for creating the platform and enabling the output) (e.g., 33%).
• The user who input prompts and guided the creative process (e.g., 33%).
• The original human creators whose music was used in the training dataset, compensated through PRS (e.g., 34%).

PRS for Music would manage the collection and distribution of royalties, ensuring that all parties receive their fair share based on transparent reporting mechanisms.

This example illustrates how existing licensing frameworks, when adapted to AI, can ensure that creators are fairly compensated, while also providing legal clarity and simplicity for AI developers.

5. Balancing Innovation and Protection

Interpreting Section 9(3) in this way strikes the right balance between fostering AI innovation and safeguarding the economic and moral rights of human creators. By distinguishing between AI-generated and AI-assisted works, clarifying the "arrangements necessary", and embedding fair licensing and royalty-sharing mechanisms, the UK can maintain a flexible, future-proof legal framework.

This approach would also position the UK as a global leader in AI copyright regulation, ensuring that legal protections evolve in tandem with technological advancements without compromising the core objective of copyright law: to reward and protect human creativity.

CGW Policy Option 1: Reform the current protection to clarify its scope

37 Would CGW legislation benefit from greater legal clarity, for example to clarify the originality requirement? If so, how should it be clarified?

Yes

Please give us your views:

The originality requirement is a cornerstone of copyright law, ensuring that protection is granted only to works resulting from human skill, labour, and judgment. However, the current Computer-Generated Works (CGW) provision under Section 9(3) of the Copyright, Designs and Patents Act 1988 (CDPA) lacks sufficient clarity, leading to confusion around authorship, ownership, and the application of originality standards in the context of AI-generated content.

Why Greater Legal Clarity is Necessary

At its core, copyright law is designed to reward and protect human creativity. AI-generated content - produced by algorithms and artificial neural networks that replicate and remix patterns from existing human works - lacks the independent intellectual effort that the originality standard demands. Without clear legal definitions, the current framework struggles to differentiate between works that deserve full copyright protection and those that do not.

To address this, legal reforms should explicitly affirm that fully AI-generated works, created without meaningful human input, do not meet the originality threshold and should be excluded from full copyright protection. These works should instead fall under a distinct related right framework, as proposed in previous responses. Works that currently fall under the scope of related rights - such as sound recordings and broadcasts - do not need to meet the originality requirement. By interpreting Section 9(3) as a related right, this removes the originality requirement from AI-generated. However, clarity is needed to distinguish between fully AI-generated works that do not attract copyright (only a related right) and AI-assisted works that are created with meaningful human contribution (to be protected by copyright).

Distinguishing AI-Assisted and Fully AI-Generated Works

One of the key challenges is distinguishing between AI-assisted and fully AI-generated works:

• AI-assisted works, where a human plays a substantial creative role—such as composing music using AI plugins integrated into Digital Audio Workstations (DAWs)—should continue to be protected under traditional copyright law, with the human recognized as the author.
• Fully AI-generated works, where the output is produced autonomously or with minimal human input (e.g., a simple text prompt), should not qualify for full copyright protection. Instead, these works would fall under a related right regime, granting limited protections to AI developers and users while maintaining the integrity of human-centric copyright principles.

This distinction ensures that the intellectual effort of human creators remains the basis for full copyright protection, while also providing a structured framework for handling AI-generated content.

Protecting Original Creators in AI Training

Another critical area requiring legal clarity is the role of original creators whose works are used to train AI systems. AI models are heavily reliant on large datasets composed of human-created content. Granting copyright or related rights to AI-generated outputs without recognizing the contributions of original creators undermines the core purpose of copyright law—to incentivise, protect, and reward human creativity.

Legal reforms should support licensing frameworks that ensure original creators are fairly remunerated when their works contribute to AI-generated outputs.

• AI developers are required to pay upfront licensing fees when using copyrighted materials in their training datasets.
• In addition to training licenses, a royalty-sharing mechanism could direct a portion of royalties from commercially successful AI-generated works back to the original creators whose content was used in training.

Addressing Market Saturation and Copyfraud

Without clearer legal definitions and distinctions, there is a risk of market saturation with AI-generated works, leading to a devaluation of human-created content and complicating copyright enforcement. Ambiguities in authorship also create opportunities for copyfraud, where individuals or companies falsely claim authorship over AI-generated works that lack meaningful human input.

Proposed Legal Clarifications

To address these challenges, the following legal clarifications are recommended:

Define Originality in the Context of AI:
• Affirm that originality requires substantial human intellectual effort.
• Clearly state that fully AI-generated works - produced without meaningful human input - do not meet the originality threshold for copyright protection and should fall under a distinct related right.
• AI-assisted works, where humans use AI tools creatively (e.g., using AI-powered plugins in DAWs), should retain full copyright protection.

Clarify Ownership:

• Specify who qualifies as the author in AI-assisted works, focusing on the human individuals who contribute significant creative input.
• Exclude individuals who merely input simple prompts into AI systems from being granted authorship rights.

Support Licensing Mechanisms:
• Support the mandatory licensing frameworks already in place for AI training datasets.
• Implement royalty-sharing models to ensure that human creators whose works were used in training are fairly compensated.

Balancing Innovation and Human-Centric Copyright

These proposed clarifications would reaffirm the human-centric foundation of copyright law while accommodating the evolving landscape of

AI-generated content. By clearly defining the originality requirement, clarifying authorship, and supporting fair licensing mechanisms, the UK can:

• Protect and reward human creativity.
• Foster ethical and responsible AI development.
• Prevent market oversaturation and protect the value of human-made content.
• Set a global standard for addressing the complexities of AI-generated works.

By refining the CGW provision under Section 9(3), the UK has the opportunity to lead the way in AI copyright regulation, ensuring that human creators remain at the heart of legal protection, ownership, and fair remuneration, even as technology continues to evolve.

38 Should other changes be made to the scope of CGW protection?

Please give us your views:

Changes should be made to the scope of Computer-Generated Works (CGWs) protection under Section 9(3) of the Copyright, Designs and Patents Act 1988 (CDPA) to better reflect the evolving landscape of AI-generated content, while safeguarding the rights and economic interests of human creators. These changes would ensure that the legal framework continues to incentivize human creativity while accommodating the realities of AI innovation.

1. Reframe CGW Protection as a Related Right

Rather than granting full copyright protection to CGWs, the law should reframe these rights as a related right - similar to the protections granted to sound recordings or broadcasts - with a shorter duration (e.g., less than the current 50 years). This approach would provide limited legal protection for AI-generated works while maintaining the integrity of copyright law, which is fundamentally designed to reward human creativity.

A related right would:

• Allow AI developers and users to commercially exploit AI-generated works under a defined legal framework.
• Limit the duration of protection to avoid market oversaturation and the devaluation of human-created content.
• Enable licensing and royalty-sharing frameworks to fairly compensate human creators whose works were used in AI training datasets.

This model ensures that AI-generated works can enter the market legally while preserving incentives for human creativity.

2. Explicit Distinction Between AI-Generated and AI-Assisted Works

The law should clearly differentiate between fully AI-generated works and AI-assisted works to avoid legal ambiguity and maintain consistency in the application of copyright protections.

• AI-generated works, created with minimal or no human input (e.g., works produced by text-to-image or text-to-music platforms like Suno or Udio), should fall under the proposed related right framework. These works lack the substantial human creativity required for full copyright protection.
• AI-assisted works, where a human plays a substantial and demonstrable role in shaping the creative output (e.g., using AI-powered plugins within Digital Audio Workstations (DAWs)), should continue to receive full copyright protection. In these cases, the human creator is the author.

A clear legal threshold must be established to define “substantial human input.” This could include criteria such as:

• The degree of creative control the human exercises over the AI tool.
• The complexity and originality of the human's contribution.
• The extent to which the AI system operates autonomously versus following human guidance.
• AI-generated v AI-assisted could be determined by the service or tool offered.

By establishing a distinction, the law would prevent AI-generated works from being treated on the same legal footing as human-created works, preserving the integrity of copyright protections.

3. Limitations on Exclusive Rights for AI Developers

If CGWs are granted legal protection (even under a related rights framework), AI developers should not be awarded exclusive rights over AI-generated content if their systems were trained on copyrighted materials. Granting full rights to AI developers without compensating the original creators whose works were used to train AI systems would fundamentally undermine the purpose of copyright law, which is to reward human creativity.

4. Support Licensing and Royalty Mechanisms

To ensure fair compensation for human creators, the law should support existing licensing and potential royalty-sharing mechanisms tied to CGWs. These mechanisms would help maintain a sustainable value chain in the creative economy, even as AI-generated content becomes more prevalent.

Proposed Licensing and Royalty Framework:

Licensing Input Data:
• AI developers must pay upfront licensing fees to access copyrighted materials for training purposes.
• These fees could be managed by existing Collective Management Organizations (CMOs), such as PRS for Music or PPL, ensuring transparency and fairness.

Royalty Sharing for Successful AI Outputs:
• When an AI-generated work achieves commercial success (e.g., streamed on platforms like Spotify or used in advertisements), a portion of the royalties should be automatically redirected to the original creators whose works contributed to the AI's training.
• CMOs could oversee the collection and distribution of these royalties, ensuring that human creators are fairly compensated.

Transparency Requirements:
• AI developers must disclose the training datasets used to create AI-generated works, allowing rights holders to verify whether their content was utilized.
• This would support accountability and help enforce licensing obligations.

5. Balancing Innovation with Human-Centric Copyright Protections

The proposed changes to the scope of CGW protection would create a legal framework that fosters AI innovation while safeguarding the rights and economic interests of human creators. By reframing CGWs as related rights, introducing clear distinctions between AI-generated and AI-assisted works, and supporting licensing and royalty-sharing mechanisms, the UK can:

• Encourage ethical AI development while preventing unauthorized use of copyrighted content.
• Reward human creators for their contributions to AI training datasets.
• Limit market oversaturation of AI-generated works, preserving the value of human-made content.
• Maintain the integrity of copyright law while adapting to the realities of emerging technologies.

These changes would position the UK as a global leader in AI copyright regulation, setting a precedent for balancing technological advancement with the fundamental goal of copyright law: to protect and reward human creativity.

39 Would reforming the CGW provision have an impact on you or your organisation? If so, how? Please provide quantitative information where possible.

Minor positive impact

Please give us your views:

Reforming the Computer-Generated Works (CGW) as a related right, as outlined in previous responses would have a minor positive impact on music creators. AI-generated works would continue to disrupt the music market and directly compete with human-created music, however, with safeguards in place human creators can be fairly remunerated for their contribution to the development of AI systems and they can continue to create knowing that the requirement to licence is enforced and that accountability, traceability and liability are in place should an AI-generated input infringe on human-created works.

Reforming the Computer-Generated Works (CGW) provision without proper safeguards would have a significant negative impact on music creators, affecting their livelihoods, income streams, and ability to maintain creative control over their works:

• AI-generated music directly competes with human-made works across streaming platforms, advertising, and sync licensing, reducing opportunities for music creators to license their works and earn royalties.
• Reframing the CGW provision without clear distinctions between AI-generated and AI-assisted works risks eroding copyright protections for music creators.
• Without enforceable licensing frameworks, AI developers could continue using copyrighted works in training datasets without permission or fair remuneration, undermining the fundamental rights of music creators.
• A surge in AI-generated music could lead to market oversaturation, making it harder for human-made music to stand out on streaming platforms and in licensing markets.
• As streaming algorithms prioritize low-cost, scalable content, music creators may face reduced exposure, leading to lower streaming revenues and diminished career sustainability, particularly for independent and emerging artists.
• Current practices around AI training lack transparency, leaving music creators unaware when their works are used to train AI models. Without mandatory disclosure requirements, creators cannot enforce their rights or claim fair compensation.
• Music is an expression of human emotion, culture, and identity. The unchecked proliferation of AI-generated content risks diminishing the diversity and authenticity of the global music landscape, replacing unique human voices with algorithmically generated works.

Carefully managed reform of the CGW provision could positively impact the music creators whose livelihoods are threatened by generative AI. Strong legal safeguards, transparent licensing frameworks, and fair remuneration models will support the music creators who stand to lose both economically and culturally in the face of expanding AI technologies.

CGW Policy Option 2: Remove specific protection for CGWs

40 Are you in favour of removing copyright protection for computer-generated works without a human author?

No

Please give us your views:

I am in favour of maintaining the current provisions for CGWs under Section 9(3) of the CDPA. It currently does no harm and holds potential future benefits as AI technologies and creative processes evolve.

Rather than removing the CGW provision, the government should focus on clarifying its scope. The provision could be adapted to serve as a related right, distinct from traditional copyright, offering limited protections for fully AI-generated works without equating them to human-created content.

This would:

• Encourage responsible AI development by incentivising AI developers to licence training data knowing that resulting works can be legally exploited under a structured framework.
• Support human creators by allowing a share of royalties from AI-generated works to be directed back to the original creators whose works were used in AI training datasets.
• Promote creativity and innovation by giving users of AI tools clarity on the legal status of AI outputs, while ensuring that human creativity remains central in copyright law.

41 What would be the economic impact of doing this? Please provide quantitative information where possible.

Please provide further comments:

Removing the Computer-Generated Works (CGW) provision from the Copyright, Designs and Patents Act 1988 (CDPA) could have far-reaching economic consequences, leading to legal uncertainty around AI-generated works and potentially harming both music creators and the broader creative economy. Without a clear legal framework governing AI-generated content, the risks to the market, creators, and the industry's long-term sustainability are significant.

Key Economic Risks of Removing the CGW Provision

1. Market Saturation with Unregulated Content
Without clear protections, AI-generated works would fall directly into the public domain, enabling companies to exploit them without oversight or accountability. This would lead to an influx of low-cost, low-quality content flooding the market, devaluing human-created music and making it harder for artists to compete. The oversaturation of unowned content would dilute streaming platforms and licensing markets, leading to lower royalties and reduced income for music creators.

2. Reduced Incentives for Collaboration and Innovation
A lack of legal clarity would likely discourage AI developers and music creators from collaborating on innovative projects, stifling new creative processes that blend human artistry with AI technologies. Without proper protections, licensing opportunities for AI-assisted works could diminish leading to lost revenue streams that could otherwise benefit both sectors.

3. Increased Risks of Copyfraud and Legal Disputes
Without defined ownership rights, bad actors could exploit the legal vacuum by falsely claiming authorship over AI-generated works, leading to copyfraud and a rise in costly legal disputes. This would create instability in the creative economy and further erode trust between creators, developers, and rights holders.

4. Weakened Copyright Enforcement and Accountability
Without clear attribution of ownership, copyright enforcement becomes nearly impossible. Music creators would lose the ability to trace the use of their works in AI-generated content, limiting their ability to hold infringers accountable or seek fair compensation.

Quantitative Evidence: The Economic Stakes

The potential economic fallout of removing the CGW provision is significant, given the size and contribution of the music industry to the UK economy. In 2023, PRS for Music reported record-breaking revenues, collecting £944 million, a 12.8% increase from the previous year (https://www.prsformusic.com/about-us/track-record/2023-financial-results). This growth spanned multiple income streams:

• Public performance: £188.2 million.
• Broadcast: £110.7 million.
• Online: £360.3 million.

According to a 2024 report by UK Music, synchronisation income (licensing music for films, TV, and games) generated £40 million in 2023, further diversifying income streams for creators (https://www.ukmusic.org/wp-content/uploads/2024/11/TIM-Report-2024-reduced.pdf).

Streaming revenues have also shown steady growth, becoming one of the most important income sources for the music industry. However, the integration of AI-generated music into streaming platforms could dilute royalty pools, leading to reduced earnings for human artists.

The Musicians' Union (MU) reports that its royalties department collects between £1.5 million and £2 million annually for musicians, covering secondary uses and further exploitation of recordings (https://musiciansunion.org.uk/working-performing/recording-and-broadcasting/musician-royalties/music-streaming-royalties). AI-generated content could significantly undercut this income if human-created works are replaced by AI outputs.

Despite these revenue streams, a 2019 government-commissioned report revealed that 62% of musicians earned £20,000 or less annually from music activities (https://www.gov.uk/government/publications/music-creators-earnings-in-the-digital-era/executive-summary). With additional economic pressures from Brexit and Covid-19, the rise of AI-generated music could further decimate musicians' incomes unless clear legal protections are in place.

42 Would the removal of the current CGW provision affect you or your organisation? Please provide quantitative information where possible

Significant negative effect

Please give us your views:

The removal of Section 9(3) could create legal uncertainty around AI-generated works, limiting future opportunities to develop frameworks that fairly compensate creators. Maintaining the provision preserves flexibility for future reforms that could benefit human creators and the wider music industry, particularly as AI technologies evolve. Removing it now would be premature and could hinder efforts to establish fair licensing and remuneration systems for AI-generated content.

Infringement and liability relating to AI-generated content

43 Does the current approach to liability in AI-generated outputs allow effective enforcement of copyright?

No

Please give us your views:

No. The current approach to liability in AI-generated outputs is inadequate and fails to provide effective enforcement mechanisms for copyright holders. While the law acknowledges that AI-generated content can infringe copyright if it reproduces a substantial part of a protected work, the lack of transparency in AI training datasets and model outputs makes enforcement extremely challenging.

1. Lack of Transparency and the "Black Box" Problem

A major barrier to enforcement is the opacity surrounding how AI models are trained and how they generate content. Without full disclosure of the training datasets, creators often cannot determine whether their works were included in AI training, making it almost impossible to prove infringement. AI models trained on vast datasets containing copyrighted works can unintentionally reproduce protected material. This results in the "black box" problem, where the internal workings of AI systems are so opaque that it becomes nearly impossible to trace how specific outputs were generated or whether they are derivative of existing works. Even when infringement occurs, the lack of traceability prevents creators from effectively enforcing their rights.

2. Ambiguity Around Liability

The current legal framework lacks clarity on who should be held accountable when an AI-generated output infringes copyright. This creates a liability gap that undermines enforcement and leaves creators with limited legal recourse. AI developers should be held responsible when their models produce infringing content, especially if they have trained systems on copyrighted data without obtaining proper licenses or failed to implement safeguards against infringement. Users who intentionally prompt AI systems to replicate copyrighted material should also bear shared liability. A shared liability model would ensure that both AI developers and users are incentivized to respect copyright laws.

3. The Case for Shared Liability

To strengthen copyright enforcement, a shared liability approach should be adopted:

AI developers would be liable if they:
• Train models on copyrighted works without licenses.
• Fail to implement sufficient technical safeguards to prevent infringing outputs.
• Do not disclose training datasets, hindering transparency and traceability.

Users would be liable if they:
• Intentionally prompt AI systems to reproduce or mimic specific copyrighted works.
• Use AI tools for purposes that clearly infringe upon existing rights.

This approach balances responsibility between the creators of AI systems and their users, ensuring that both parties are accountable and that rights holders have clearer paths to enforcement.

The current approach to liability in AI-generated outputs is insufficient and fails to provide creators with the tools necessary to enforce their rights effectively. The lack of transparency, combined with ambiguous liability frameworks, leaves significant gaps in copyright enforcement. Adopting a shared liability model, along with mandatory transparency obligations for AI developers, would provide a more robust legal framework that protects the rights of music creators and ensures that AI-generated outputs respect existing copyright laws.

44 What steps should AI providers take to avoid copyright infringing outputs?

Please give us your views:

AI providers must adopt a proactive and comprehensive approach to prevent copyright infringement and ensure their models do not reproduce protected works without permission. Protecting the rights of creators requires a multi-layered strategy that incorporates licensing, technical safeguards, transparency, and accountability.

1. Obtain Proper Licenses for Training Data

The most critical step AI providers must take is to secure appropriate licenses for any copyrighted material used in training datasets. Licensing is the foundation of the music industry and other creative sectors, ensuring that creators are fairly compensated for the use of their works. AI developers

should be held to the same standards as other commercial users of creative content. All copyrighted works used for AI training must be licensed through Collective Management Organizations (CMOs), such as PRS for Music or PPL, ensuring that rights holders are included in the value chain. Licensing not only ensures compliance with copyright law but also promotes a fair and sustainable market for both AI developers and music creators.

2. Implement Robust Prompt and Output Filtering Systems

To prevent AI systems from generating infringing outputs, providers should integrate advanced filtering mechanisms at both the input (prompt) and output stages.

Prompt Filtering:

• AI models should detect and block user prompts that attempt to generate content based on specific copyrighted works or styles.
• For example, prompts like "create a song that sounds like [specific artist]" should trigger filters that prevent the model from producing infringing content.

Output Filtering:
• AI-generated outputs should be scanned using content recognition algorithms to detect unauthorized reproductions of copyrighted material.
• This includes preventing the generation of near-identical copies or derivative works without appropriate licensing.
• Filters should be dynamic and regularly updated to adapt to evolving risks and the continuous influx of new copyrighted works.

3. Ensure Full Transparency in Training and Outputs

Transparency is a cornerstone of responsible AI development. AI providers must fully disclose their training practices and provide clear documentation of the data used to train their models. AI developers should maintain public logs of their training datasets, enabling rights holders to verify whether their works were used without consent. Disclosing this information would allow for accountability and help creators protect their rights.

Detailed records of how AI systems process data and generate outputs should be maintained, providing an audit trail in the event of copyright disputes. This level of transparent reporting will foster trust between AI developers and rights holders while enabling more effective copyright enforcement.

4. Establish Clear Redress and Takedown Mechanisms

To further protect rights holders, AI providers must implement accessible and efficient processes for reporting and removing infringing content.

Rights Holder Reporting Tools:
• A streamlined system should be in place for rights holders to flag infringing outputs.
• These tools must be user-friendly, allowing creators to easily report violations and request takedowns.

Swift Takedown Procedures:
• Upon receiving valid reports, AI providers should be required to promptly remove infringing content.
• A clear timeline for takedowns should be established to ensure timely action.

Appeals Process for Users:
• To balance the rights of users, AI providers should offer a transparent appeals process for instances where content has been wrongly flagged or removed.
• This ensures fairness while maintaining strong protections for rights holders.

5. Enforce Strong Penalties for Non-Compliance

To incentivize compliance and accountability, AI developers must face meaningful penalties for failing to uphold copyright standards. Companies that fail to license training data, implement filtering systems, or respond to infringement reports should face substantial fines and legal repercussions. Persistent violators could face restrictions on deploying their AI systems until compliance is achieved. A shared liability model should be introduced, where both AI developers and users bear responsibility for infringing outputs. AI developers would be liable if their models are trained on unlicensed data or if safeguards against infringement are insufficient. Users should be held accountable if they intentionally prompt AI systems to produce infringing content. This model ensures that all parties involved are incentivized to respect copyright laws.

Conclusion

The current liability framework for AI-generated outputs is insufficient, primarily due to a lack of transparency, weak enforcement mechanisms, and unclear attribution of responsibility. AI providers must take proactive steps to prevent copyright infringement by:

• Licensing all training data.
• Implementing prompt and output filtering systems.
• Ensuring full transparency in AI training and output processes.
• Establishing robust takedown and redress mechanisms.
• Accepting liability for non-compliance, supported by meaningful legal penalties.

By adopting these measures, AI providers can ensure that their technologies operate within legal frameworks that protect music creators and uphold the integrity of copyright law.

AI Output labelling

45 Do you agree that generative AI outputs should be labelled as AI generated? If so, what is a proportionate approach, and is regulation required?

Yes

Please give us your views:

Yes. It is essential that AI-generated outputs are clearly labelled to promote transparency, uphold the rights of creators, and allow audiences to distinguish between human-made and AI-generated content. From both an academic and music practitioner perspective, labelling is a critical safeguard, not only for the creative industries but also for the integrity of public discourse and consumer trust.

The Need for Labelling:

• Transparency and Trust: Labelling AI-generated content ensures that consumers can distinguish between authentic, human-created works and algorithmically produced outputs. This is crucial for maintaining public trust in creative sectors like music, journalism, and visual arts.
• Combating Misinformation: The proliferation of deepfakes, synthetic voices, and manipulated content presents serious risks, particularly when social media platforms like Meta are reducing their efforts to combat misinformation. Without clear labelling, AI-generated materials could easily be used to mislead the public, eroding trust in news, politics, and even cultural content.
• Protecting Creative Industries: Unlabelled AI-generated content can undermine the value of human creativity. Without clear markers, AI-generated music could be misrepresented as human-made, misleading listeners and potentially diverting royalties away from human creators.

A Proportionate Approach:

• Automatic Labelling: AI developers should integrate metadata tagging and digital watermarking directly into their models. These labels should be applied at the point of creation and be tamper-proof, preventing users from removing or altering them post-production.
• Universal Application: Labelling standards should apply across all content types—from music and images to video and text—to ensure consistency and transparency across platforms.
• Visible, audible and Machine-Readable: Labels should not only be embedded in metadata but also presented in a visible or audible format so that consumers can easily identify AI-generated content without specialized tools.

The Role of Regulation:

• Mandatory Labelling Standards: Voluntary labelling will be insufficient. The government must introduce mandatory regulations that require AI developers to label all outputs.
• Enforcement Mechanisms: Strong penalties for non-compliance—such as fines or restrictions on the deployment of AI systems—are essential to deter bad actors who may attempt to remove or falsify labels.
• Consumer Protection: Regulation will also help protect consumers, ensuring that they are not misled by unlabelled AI content in advertising, music, news, or other creative works.

Clear labelling of AI-generated outputs is vital to safeguarding public trust, protecting creative rights, and maintaining the integrity of information in the digital world. Regulation is essential to ensure these standards are met and consistently enforced.

46 How can government support development of emerging tools and standards, reflecting the technical challenges associated with labelling tools?

Please give us your views:

The government has a critical role in supporting the development of labelling tools that are effective, tamper-proof, and capable of working across different forms of AI-generated content. Addressing the technical challenges requires a collaborative approach, balancing innovation with the need for transparency and creator protection.

1. Supporting Research and Development

The government should actively fund research and development into advanced labelling technologies that can securely mark AI-generated content. This includes investing in digital watermarking and cryptographic verification methods to create tamper-proof labels embedded directly into AI outputs. To further enhance traceability and accountability, the government could also explore the use of blockchain technology. Blockchain can provide immutable records of AI-generated works, making it easier to verify the origins of content and ensure that labels remain intact throughout the content's lifecycle.

2. Establishing Standardized Labelling Through Industry Collaboration

Creating universal labelling standards requires strong collaboration between AI developers, creative industries, and consumer rights organizations. It's essential that key stakeholders—such as music creators, visual artists, and writers—are involved in shaping labelling frameworks to address the unique needs of each creative sector. These labelling standards must work consistently across all forms of content, whether it's music, images, video, or text. This ensures that regardless of the medium, AI-generated content remains transparent and identifiable for both consumers and rights holders.

3. Launching Pilot Programs and Real-World Testing

The government should sponsor pilot programs to test emerging labelling tools in real-world environments. These programs would help identify technical challenges, refine labelling solutions, and ensure that they are both effective and user-friendly. Testing should cover a wide range of platforms, from streaming services to social media networks, to ensure labelling tools function reliably across diverse use cases. This approach will help fine-tune labelling technologies before broader implementation, reducing the risk of technical failures or gaps.

4. Promoting International Collaboration for Global Standards

Since AI-generated content circulates globally, the UK should work closely with international bodies to align its labelling standards with global best practices. Consistency across borders is crucial for effective enforcement and transparency. The UK should aim to develop interoperable labelling tools that align with frameworks used in the EU, US, and other major markets. This will ensure that AI-generated content is properly labelled and recognizable no matter where it's created or distributed, promoting global transparency and helping to protect the rights of creators worldwide.

47 What are your views on the EU's approach to AI output labelling?

Please give us your views:

The EU's AI Act takes an important step toward enhancing transparency and accountability in AI-generated content, but it falls short in several key areas, especially from the perspective of creative industries like music. While the framework focuses on labelling AI outputs to support copyright enforcement, it overlooks critical aspects that would help both rights holders and consumers navigate the growing landscape of AI-generated content.

Where the EU Model Falls Short

One major limitation of the EU's approach is its focus on machine-readable labels, which are useful for rights holders and enforcement agencies but offer little value to consumers. In sectors like music, where AI-generated tracks can closely mimic human-made works, it becomes nearly impossible for listeners to know whether a song was created by a person or an algorithm. To address this, visible or audible cues - such as on-screen notifications, audible notifications or watermarks - are essential so that consumers can easily identify AI-generated works.

Another shortcoming is the weak enforcement mechanisms. Without strong penalties for non-compliance, there's a real risk that AI developers might bypass labelling requirements. Without consistent enforcement, labelling systems lose their effectiveness, and rights holders are left without the tools they need to protect their work.

Opportunities for the UK to Lead

The UK has a unique opportunity to build upon the EU's framework and create a more comprehensive, consumer-friendly approach. Stronger enforcement mechanisms should be introduced, ensuring that AI developers face meaningful penalties if they fail to properly label AI-generated content. This would help maintain accountability and ensure compliance across the industry.

Additionally, the UK could implement a dual labelling system that serves both rights holders and consumers. This would include machine-readable labels for copyright enforcement alongside consumer-facing labels—such as visible tags or audible indicators—so users can easily recognize AI-generated content.

To further strengthen its position, the UK should promote cross-sector collaboration in developing labelling standards that work across music, visual arts, literature, and other creative sectors. Establishing consistent and transparent labelling protocols would support ethical AI development and solidify the UK's reputation as a leader in responsible AI governance.

Conclusion

While the EU's approach represents a positive first step, it remains incomplete. The UK has the chance to refine this framework, placing a stronger emphasis on consumer protection, rights holder transparency, and robust enforcement. By going beyond the EU's model, the UK can set a new global standard for ethical AI development—balancing innovation with the need to protect the integrity of creative industries and the rights of human creators.

Digital replicas and other issues

48 To what extent would the approach(es) outlined in the first part of this consultation, in relation to transparency and text and data mining, provide individuals with sufficient control over the use of their image and voice in AI outputs?

Please give us your views:

The measures proposed in the consultation - focused on transparency and text and data mining (TDM) - fall short of giving individuals meaningful control over how their name, image, voice, or likeness is used in AI-generated content. Transparency on its own is a reactive tool; it only helps individuals realize that their voice or image has been used without permission after the fact. This approach offers no preventative protection, leaving creators, especially musicians, vulnerable to misuse. In the music industry, AI technologies can now replicate voices or instrumental performances with minimal data, allowing unauthorized copies to circulate before the original artist is even aware.

A major contributor to this issue is the use of TDM techniques. AI developers routinely scrape the internet, pulling vast amounts of data - including music recordings, voices, images, and videos - often without securing explicit consent from the creators or individuals involved. When TDM exceptions are applied broadly to commercial AI development without strong licensing or consent requirements, it essentially legalizes the extraction of personal data, leaving musicians and other creators exposed to exploitation.

The proposed opt-out system is not a workable solution. Expecting musicians to locate every instance of their recorded voice or performance online and apply opt-outs to each is simply unrealistic. The same issue applies to individuals outside the creative industries. Consider whether every public figure - including politicians - should be responsible for opting out of every news broadcast, social media post, or public appearance to avoid being digitally cloned. Clearly, this is an impractical expectation.

Existing legal protections also fail to bridge this gap. The UK's tort of passing off, for example, was designed to protect commercial goodwill, not individuals' likenesses or voices. This law makes it difficult for artists to take action when AI-generated content mimics their voice or style without directly misrepresenting itself as their work. Even if a deepfake song copies an artist's voice but is clearly labeled as AI-generated, it may not violate existing laws—despite misleading listeners and potentially harming the artist's reputation.

To provide real control over personal identity in AI-generated content, the UK needs to introduce legal safeguards that go beyond transparency and opt-outs. There should be explicit consent requirements before any person's voice, image, or likeness can be used in AI-generated works. Adopting personality rights, similar to the proposed NO FAKES Act in the United States, would ensure that artists and public figures have legal protection against unauthorized replication of their identities in AI-generated media.

49 Could you share your experience or evidence of AI and digital replicas to date?

Please give us your views:

The music industry has already seen several high-profile cases where AI-generated content has been used to replicate artists' voices and styles without permission, highlighting the urgent need for stronger legal protections.

One striking example is an Instagram video titled "Rihanna's Most Expensive Purchases," where an AI-generated voice mimicked Rihanna, falsely portraying her discussing luxury goods. This use of an unauthorized voice clone misled viewers, prompting Rihanna to publicly denounce the misuse. The case drew attention to the growing problem of AI-generated content impersonating artists without consent.

Another notable incident occurred in 2023 when an anonymous artist, known as Ghostwriter, released "Heart on My Sleeve," an AI-generated track that replicated the voices of Drake and The Weeknd. The song quickly went viral, racking up over 600,000 Spotify streams, 275,000 YouTube views, and 15 million views on TikTok before it was taken down. It is estimated that the Spotify streams alone could have generated over £8,000. Interestingly, the song wasn't removed due to the cloned vocals—which current copyright laws don't protect—but because it contained an unauthorized sample. This loophole reveals a troubling gap in legal protections: artists currently lack rights over the unauthorized use of their voices in AI-generated works.

The issue extends beyond vocalists. Instrumentalists have also been targeted. There have been instances where colleagues of mine have been hired to record individual notes and articulations on their instruments over several days. These recordings were then used to train AI systems that can replicate the nuances of the performer's style - again, often without the artist's consent or proper compensation. This practice allows AI models to produce new recordings or instrumental models that sound convincingly human, potentially negating the need to pay humans to perform on soundtracks, studios to host recording sessions, and engineers and producers to refine the tracks.

These examples demonstrate how AI technologies can exploit artists and creators without their approval. Without new legal frameworks that specifically address voice cloning, performance replication, and the use of personal likenesses, musicians and performers will continue to be vulnerable. The UK must consider implementing stronger protections, similar to the NO FAKES Act proposed in the United States, to safeguard the identities and creative outputs of artists and public figures against unauthorized AI-generated replicas.

Other emerging issues

50 Is the legal framework that applies to AI products that interact with copyright works at the point of inference clear? If it is not, what could the government do to make it clearer?

No

Please give us your views:

The current legal framework lacks clarity regarding how copyright law applies when AI models interact with copyrighted works at the point of inference (i.e., in real-time as new content is generated). Any use of copyrighted works at the point of inference requires a valid licence, just as it does during training. AI developers must be held accountable if their systems access or process copyrighted works without appropriate permissions.

Transparency requirements should be extended to cover inference-stage interactions. AI providers should be mandated to disclose whether their models retrieve or reference copyrighted content during inference and implement safeguards to prevent unlicensed use. The government should introduce clear legal obligations, supported by regulatory oversight, to ensure compliance and accountability in this area.

51 What are the implications of the use of synthetic data to train AI models and how could this develop over time, and how should the government respond?

Please give us your views:

Synthetic datasets are still derived from original human-created works. This means they may indirectly carry elements of copyrighted material, raising concerns about the legitimacy of their use in AI training.
The government should not assume that synthetic data eliminates copyright concerns. AI developers should be required to demonstrate that their synthetic datasets are not built upon or derived from unlicensed copyright content. This could be achieved through transparency obligations and third-party audits to verify the provenance of synthetic datasets.

The government should monitor developments in this area and consider implementing specific legal frameworks to address emerging risks, ensuring that creators' rights remain protected.

52 What other developments are driving emerging questions for the UK's copyright framework, and how should the government respond to them?

Please give us your views:

The UK's copyright framework faces mounting pressure from rapid technological advancements, particularly around AI development, streaming economics, and international inconsistencies in copyright law. Several critical issues are emerging that require government attention to ensure that creators' rights are protected in this shifting landscape.

1. The Need for International Harmonisation

One of the most pressing challenges is the lack of a unified international approach to AI and copyright. AI development operates on a global scale, but stark differences between UK, US, EU, and other international copyright laws create a fragmented legal environment. This inconsistency opens legal loopholes that AI developers can exploit, allowing them to train AI systems on copyrighted content in jurisdictions with more permissive laws and then commercialise the outputs globally without facing uniform legal consequences.

This problem is particularly evident in the case of downstream commercialisation, where AI models are trained under non-commercial research exceptions but later used for profit without compensating rights holders. The LAION v. Kneschke case in Germany highlighted this loophole, where the court ruled that the downstream use of non-commercially curated datasets was legally permissible, even when used for commercial purposes later (https://www.euipo.europa.eu/en/law/recent-case-law/germany-hamburg-district-court-310-o-22723-laion-v-robert-kneschke). This sets a dangerous precedent and highlights the need for coordinated legal responses.

The UK government must play an active role in global copyright reform, collaborating with bodies like the World Intellectual Property Organization (WIPO) to promote international harmonisation. Establishing cross-border standards on issues such as training data licensing, dataset transparency, and commercialisation pathways would help close existing legal loopholes and ensure stronger protections for creators worldwide. A coordinated global effort is the only way to hold AI developers accountable and safeguard the economic interests of music creators in an interconnected market.

2. The Economics of Streaming and AI's Disruptive Impact

The UK music industry continues to grapple with unresolved issues in streaming economics, which have only been compounded by the rise of AI-generated content. In its 2021 inquiry, the Digital, Culture, Media and Sport (DCMS) Committee called for a "complete reset" of music streaming, highlighting fundamental flaws in the royalty distribution system (https://publications.parliament.uk/pa/cm5803/cmselect/cmcumeds/874/report.html). Despite recommendations to explore fairer payment models - such as equitable remuneration akin to the broadcast model - the government has so far resisted implementing systemic changes, citing the risks to the current streaming economy.

These unresolved issues are now colliding with the disruptive force of AI-generated music. As AI systems increasingly produce music that competes with human-made content, artists face yet another threat to their already precarious income streams. Projections suggest that AI-generated tracks could account for up to 20% of total music streaming revenues by 2028, diverting even more income away from human musicians (https://info.xposuremusic.com/article/how-ai-generated-music-could-impact-music-catalog-valuations). Without legal clarity on how AI-generated works are treated within existing streaming models, this could exacerbate the inequalities that the DCMS report sought to address.

To future-proof the industry, the government must integrate AI-related concerns into ongoing reforms around streaming economics. Policies must be introduced that differentiate between AI-generated and human-created music, possibly applying separate royalty frameworks to ensure that human creators are not undercut by algorithmically produced content. Additionally, there must be clear guidelines around how AI-generated music is labelled, monetized, and integrated into streaming platforms to avoid flooding the market with low-cost, high-volume AI content.

3. Addressing Copyright Gaps: Downstream Commercialisation and TDM Loopholes

Another critical area that demands attention is the growing issue of downstream commercialisation and text and data mining (TDM) loopholes. The existing TDM exception for non-commercial research is increasingly being exploited by AI developers who scrape vast amounts of copyrighted material under the guise of research, only to later commercialise the trained models or their outputs. This creates an unregulated marketplace where the original rights holders—musicians, songwriters, and performers—receive no compensation for their contributions, despite their works being integral to the AI's success.

The UK government must act decisively to close these loopholes. One approach would be to introduce mandatory licensing schemes for TDM activities, requiring AI developers to secure licenses before using copyrighted works, even for non-commercial purposes. Additionally, retrospective licensing could be mandated for cases where non-commercial TDM leads to downstream commercialisation, ensuring that rights holders are compensated when their work contributes to profitable AI outputs.

4. Preparing for Future AI Innovations

The current focus on generative AI in music is only the beginning. Emerging technologies like autonomous generative AI, capable of creating complex works without direct human prompts, will pose even greater challenges for copyright law. As AI tools become more advanced, the line between AI-generated and AI-assisted content will blur, complicating questions of authorship, ownership, and licensing.

The government needs to adopt a proactive rather than reactive approach, continuously engaging with industry stakeholders, creators, and technologists to understand the evolving landscape. Establishing regular consultation forums focused on the intersection of AI, copyright, and creative industries would help policymakers stay ahead of technological developments and craft legislation that balances innovation with robust protections for creators.

A Call for Proactive Reform

The UK's copyright framework is at a pivotal moment. The combination of international legal fragmentation, AI-driven commercialisation loopholes, and unresolved issues in streaming economics is creating a perfect storm that threatens the rights and livelihoods of creators. Without decisive action, the creative industries—particularly music—face the risk of significant economic and cultural harm.

The government must take a multi-faceted approach: leading efforts for global copyright harmonisation, reforming streaming royalty models to address the influx of AI-generated content and closing TDM loopholes that allow unlicensed data scraping and downstream commercialisation. Only by addressing these intersecting challenges can the UK maintain a copyright framework that supports both technological innovation and the sustainability of its creative industries.