



decade.ac.uk

11 February 2025

DECaDE – UKRI Centre for the Decentralised Digital Economy

DECaDE’s Response to the UK Government Open Consultation on Copyright and Artificial Intelligence

1. Executive summary of DECaDE’s response

DECaDE’s position is that the UK can play a leading role in shaping balanced AI copyright policy. By supporting emerging standards and decentralized technologies, the UK can uphold creator rights while stimulating an innovative environment for AI research and commercial development.

1. Support for a Balanced “Opt-Out” Exception (Option 3)

DECaDE supports Option 3 (data mining exception with rights reservation). We believe this best balances innovation opportunities and rightsholder concerns, ensuring that the UK remains an attractive location for AI research and development while respecting the interests of creators. By permitting text and data mining (TDM) with the ability for rightsholders to opt out, creators may reserve rights in a manner aligned with emerging EU frameworks, giving the UK legislative alignment and competitive positioning.

2. Importance of Granularity and Transparency

We recommend a more granular approach to “AI opt-out” than current EU TDM provisions, enabling rightsholders to distinguish between types of AI usage (e.g. training generative models vs. AI-driven search, classification, or recommendation). This avoids an overly broad blanket opt-out that may inadvertently stifle other uses including content discoverability.

- **Machine-Readable Declarations:** DECaDE advocates for an open, patent-free technical standard—like the Coalition for Content Provenance and Authenticity (C2PA)—to encode provenance data including authorship and AI opt-out signals in a machine-readable form.
- **Encouraging Preservation of Metadata:** We urge the government to incentivise platforms to preserve these provenance markers. The stripping of provenance metadata by content distribution platforms remains a critical barrier to scalable, enforceable opt-out frameworks.

3. Licensing Models and Emerging Compensation Frameworks

While current large-scale licensing practices exist, we see an opportunity for new, decentralized models that could automatically compensate creators for reuse of their work in AI training. The ORA framework co-developed by DECaDE and some of its commercial partners, for example, attaches rights and payment information to content, enabling transparent and fair compensation. We encourage the government to work with academia and industry to explore these emerging solutions for enabling creator value from AI in our Decentralized Creative Economy.

4. Proportionate Transparency for Generative AI

Disclosing every detail of an AI model’s training corpus may compromise commercial secrets. We instead recommend that AI developers attest to using properly licensed data. Further, they could disclose the licenses under which training data was used, rather than specific sources. This helps build trust and clarity over permissible use without deterring investment and innovation in UK-based AI development. It can also help identify inadvertent re-training on synthetic data which can reduce the quality and commercial value of AI models.

5. Government Leadership and Regulatory Underpinning

We believe government legislation should remain at a high level—requiring that AI opt-out and provenance be communicated via open standards—rather than prescribing a specific technology or standard. At the same time, a clear regulatory framework underpinning transparency and rights reservation would ensure compliance. We also encourage the government to lead by example, embedding content provenance metadata its own published media or in public broadcasters.

2. Introduction to DECaDE

DECaDE is the UKRI Centre for the Decentralized Digital Economy, a multi-disciplinary research centre led by the **University of Surrey** in partnership with the **University of Edinburgh** and the **Digital Catapult**. DECaDE is funded 2020-2026 by the UKRI/EPSC.

DECaDE’s mission is to explore how decentralized platforms and data centric technologies such as AI and Distributed Ledger Technology (DLT) can help create value and in our future digital economy in which everyone is a producer and consumer of digital goods and services. DECaDE studies these questions primarily through the lens of the creative industries, which have shifted from monolithic content producers to a decentralized model where individuals and smaller production houses also increasingly produce and consume content disseminated via online platforms. As a multidisciplinary academic research centre, DECaDE brings together technical expertise in AI, DLT and Cyber-security, with business, law, and human factors / design. DECaDE co-creates its research with over 30 commercial partners including several in the creative sector including Adobe and the BBC, and has worked extensively with the cabinet office and Scottish government. More information on DECaDE and its partners is at decade.ac.uk [1].

DECaDE focuses on provenance of assets within digital supply chains, and has researched and commercialized multiple data provenance technologies to trace the provenance of media with a view to tackling fake news and misinformation online. For example, DECaDE researchers have developed digital watermarking and fingerprinting technologies to identify and trace media (durable content credentials) cited in a recent report by the international security agencies [2]. They have contributed to the development of the Coalition for Content Provenance and Authenticity (C2PA) which develops an eponymous international standard for expressing media provenance, commonly referred to as Content Credentials [3] and broadly adopted by industry (e.g. Meta, Google, Adobe, Microsoft, Amazon, Tiktok, LinkedIn, Samsung, and others) as well as being fast-tracked as an ISO standard under ISO/TC 171/SC 2. As an early member of the Content Authenticity Initiative (CAI) DECaDE’s technologies have been integrated by multiple industry partners to signal use of Generative AI and help trace the origins and actions performed on media in order to enable users to make more informed trust decisions about online content [4].

DECaDE has been exploring applications of provenance beyond transparency and authenticity, to help create value in new ways within the decentralized creative economy. For example, DECaDE co-

developed the ORA framework with its industry partners; a prototype technical solution to identify the ownership and rights associated with content [5]. Building upon the C2PA open standard, these technologies demonstrate ways for creators to attach rights information (including Copyright, and AI training opt-out signals [7]) as well as receive automated royalty payments for the re-use of content including its use to training Generative AI [5,6]. These technologies have been built and workshopped with creative rightsholders [8,9].

References

- [1] **DECaDE: UKRI/EPSCRC Centre for the Decentralized Digital Economy.** <https://decade.ac.uk/>
- [2] **Content Credentials: Strengthening Multimedia Integrity in the Generative AI Era.** <https://media.defense.gov/2025/Jan/29/2003634788/-1/-1/0/CSI-CONTENT-CREDENTIALS.PDF>
- [3] **Coalition for Content Provenance and Authenticity (C2PA).** Technical Specification, v2.1. <https://c2pa.org/specifications/specifications/2.1/index.html>
- [4] **To Authenticity, and Beyond! Building Safe and Fair Generative AI upon the Three Pillars of Provenance.** J. Collomosse and A. Parsons. IEEE Computer Graphics and Applications (IEEE CG&A). 2024. <https://bit.ly/provtriad>
- [5] **EKILA: Synthetic Media Provenance and Attribution for Generative Art.** K. Balan and S. Agarwal and S. Jenni and A. Parsons and A. Gilbert and J. Collomosse. In Proc. CVPR Workshop on Media Forensics (CVPRW). 2023. <https://arxiv.org/abs/2304.04639>
- [6] **ProMark: Proactive Diffusion Watermarking for Causal Attribution.** V. Asnani, J. Collomosse, T. Bui, X. Liu and S. Agarwal. In Proc. Computer Vision and Pattern Recognition (CVPR). 2024. <https://arxiv.org/abs/2403.09914>
- [7] **DECORAIT - DECentralized Opt-in/out Registry for AI Training.** K. Balan and A. Black and S. Jenni and A. Gilbert and A. Parsons and J. Collomosse. In Proc. Conf. Visual Media Production (CVMP). 2023. <https://arxiv.org/abs/2309.14400>
- [8] **ORAgen: Exploring the Design of Attribution through Media Tokenisation.** F. Liddell, E. Tallyn, E. Morgan, K. Balan, M. Disley, T. Koterwas, B. Dixon, C. Moruzzi, J. Collomosse and C. Elsdén. In Proc. ACM Designing Interactive Systems (DIS). 2024. <https://decade.ac.uk/articles/oragen-exploring-the-design-of-attribution-through-media-tokenisation/>
- [9] **ORAgen: Emerging Futures for Tokenisation and Digital Media Rights.** DECaDE/Digital Catapult Technical Report.. <https://www.digicatapult.org.uk/publications/post/oragen-emerging-futures-for-tokenisation-and-digital-media-rights/>

3. Response to the Open Consultation Questions

3.1 Copyright and Artificial Intelligence

Question 1: 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?

Yes

Question 2: Which Option do you prefer and why?

DECaDE prefers **Option 3**, which uses an opt-out model balancing the interests of rightsholders with the UK’s global competitiveness. An opt-in requirement would severely limit access to the large data volumes needed for commercial AI development, reducing the attractiveness of the UK for both academic research and spin-out ventures. Smaller operators, such as universities and startups, would struggle to obtain enough content for their projects if they had to secure explicit permissions upfront. Option 3 also aligns the UK with the EU Copyright Directive (2019) TDM exception, which will help to keep the country competitive in AI innovation by avoiding the risk of talent and research relocating to jurisdictions with potentially more permissive legislation around re-use of data in AI training.

3.2 Our proposed approach: Exception with rights reservation

Question 3: Do you support the introduction of an exception along the lines outlined above?

Yes

Question 4: If so, what aspects do you consider to be the most important? If not, what other approach do you propose and how would that achieve the intended balance of objectives?

Two key considerations are **granularity** and **transparency**. In terms of granularity, an opt-out framework should recognize that “AI training” spans different processes—ranging from full-scale generative model creation to smaller-scale fine-tuning or analytic tasks like search and recommendation. A single, broad opt-out of ‘AI’ can unintentionally block beneficial uses, such as content discoverability through AI-driven search – which is often desired by rightsholders. This is the case with existing European Legislation (EU Copyright Directive 2019, Article 4) which is a blanket opt-out of all ‘data mining’ including AI. Perversely this could even prevent content being indexed by search registries for the purpose of discovering opt-out. There is an opportunity for the UK to take a more granular approach in its Copyright legislation that allows rightsholders to specify which types of AI processes they wish to exclude—for instance, prohibiting the training of large generative models while still permitting analytic uses.

Transparency involves making these choices clear and verifiable at both the **site level** (applying an overarching rule to an entire domain) and the **asset or “unit” level** (embedding metadata in each file). To minimize friction for rightsholders and developers, any approach should adopt open standards that enable automated identification and enforcement of opt-out markers. One such standard, already being applied by industry [Forbes, 2024] for unit-level protection is the Coalition for Content Provenance and Authenticity (C2PA) [C2PA] which also offers the abovementioned granular opt-out. The C2PA is already widely adopted across industry, and is being fast-tracked as an ISO standard under ISO/TC 171/SC 2 for ISO adoption early 2025. The Joint Photo Expert Group (JPEG) standards committee responsible for the common JPEG image format, has already incorporated the C2PA into their own JPEG Trust standard (ISO/IEC 21617-1:2025). At the site-level, there are defacto standards such as the convention of placing a ‘robots.txt’ file on a website indicating which areas of the site may be scraped. Some rightsholders are critical that this ‘throws the baby out with the bathwater’ in that they often desire content discovery through search, but wish to opt-out of AI training. Early discussions are being had around an ‘ai.txt’ alternative, or standards for indicating site-level opt-out such as TDMRep [TDMRep]. In any case all opt-out standards are new and emerging, and no single approach covers well both unit and site level opt-outs. DECaDE does not therefore advocate for the naming of

a particular standard within UK Copyright legislation, but rather for the encouragement of open standards for machine-readable opt-out.

References

[TDMRep] TDM Reservation Protocol (TDMRep). W3C Community Report. February 2024. Available at: <https://www.w3.org/community/reports/tdmrep/CG-FINAL-tdmrep-20240202>

[C2PA] Coalition for Content Provenance and Authenticity, Technical Specification v2.1. Publisher: Linux Foundation. Available at: <https://c2pa.org>

[Forbes 2024] Adobe Debuts Free Web App To Fight Misinformation And Protect Creators. Forbes Magazine. October 2024. <https://www.forbes.com/sites/moorinsights/2024/10/16/adobe-debuts-free-web-app-to-fight-misinformation-and-protect-creators/>

Question 5: 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.

Introduction of an AI opt-out exception would enable creatives to have increased confidence when distributing their content, as there would be a legal support for any opt-out signals carried with that work. As we discuss later, content provenance standards such as the C2PA can not only enable opt-out declarations, but can also be extended include ownership and payment details, opening avenues for creators to monetize AI-driven reuse of their content—particularly in the emerging decentralized creative economy [Collomosse 2024]. However, it is important to recognise the burden that this additional granularity may place on some creatives to understand and anticipate how different AI tools may use and exploit their work. Without clear visibility of potential benefits, some rightsholders may simply opt out of all AI uses by default, forfeiting opportunities to earn revenue from permissible or desirable AI activities.

[Collomosse 2024] **To Authenticity, and Beyond! Building Safe and Fair Generative AI upon the Three Pillars of Provenance.** J. Collomosse and A. Parsons. IEEE Computer Graphics and Applications (IEEE CG&A). 2024 <https://bit.ly/provtriad>

Question 6: What action should a developer take when a reservation has been applied to a copy of a work?

DECaDE believes that developers should respect any clear, machine-readable markers indicating that a work is off-limits for AI training or analysis. At a minimum, they should honour any metadata attached to a data item (so called 'unit level' opt-out) that indicates the reservation using a widely adopted open standard capable of expressing the same. Such opt-out may be specified using an open standard such as C2PA (Content Credentials) At a minimum site-level opt-out instructions (e.g., “no-scrape” flags) on websites similar to ‘robots.txt’ in the context of search engine scraping, should also be honoured.

Challenges emerge when content is copied from one site to another containing no such site-level markers, or when unit-level metadata attached to individual assets is stripped away by content platforms (such as social media sites) during distribution online. Although it is difficult to recover from this situation for site-level opt-out, unit-level opt-out can be reinforced in two main ways. 1) Metadata standards can employ watermarking technologies to improve the persistence of metadata through platforms, for example Durable Content Credentials for C2PA [Adobe 2024, NSA 2025]; 2) Registry services may be established to lookup missing provenance information using visual matching

technologies such as content fingerprinting, for example DECORAIT [Balan 2023b] describe a way to harness decentralized lookup technologies for this purpose.

References

[Balan 2023b] **DECORAIT - DECentralized Opt-in/out Registry for AI Training**. K. Balan et al. Proc. Conf. Visual Media Production (CVMP). 2023. <https://arxiv.org/abs/2309.14400>

[Adobe 2024] **Durable Content Credentials**. Content Authenticity Initiative Blog post. <https://contentauthenticity.org/blog/durable-content-credentials>

[NSA 2025] **Joint guidance on content credentials and strengthening multimedia integrity in the generative artificial intelligence era**. January 2025. US National Security Agency. <https://media.defense.gov/2025/Jan/29/2003634788/-1/-1/0/CSI-CONTENT-CREDENTIALS.PDF>

Question 7: What should be the legal consequences if a reservation is ignored?

This is potentially one of the most significant, but also most contentious aspects of the proposed regime. Traditional copyright law has a range of possible remedies. The most common is monetary damages that compensate the rightsholder for quantifiable losses. These however will frequently be of little use in the context of training AI models, as the loss for the creator will often be difficult to quantify, and also often be de minimis. The role of collecting societies might mitigate this to a degree, though only a small percentage of rightsholders will be represented by them. In the US, class actions are a powerful tool to combine large numbers of small infringements into litigation that can achieve also the desired deterrent effect, but the creation of such a procedural tool would be a significant change for UK law. Without additional remedies, this approach therefore give larger commercial entities an undue advantage and allows them to treat these damages simply as a business expense, effectively granting themselves a de facto license even though a rightsholder has explicitly indicated opt-out.

In addition to compensatory damages, UK copyright law also allows the award of “additional” damages under the EU Enforcement Directive 2004/48/EC and the UK IP Enforcement Regulations 2006. These act as “punitive damages” that also have a deterrent effect, and protect not just the individual claimant, but society at large. Because of that dual nature though, they can raise rule of law issues (criminal law rationales in private law litigation, with its lower evidential burdens etc) . In the past, as far as we can tell, they have been used if at all only very sparingly by the courts. For the question of AI training, they may provide a more equitable solution that gives big tech companies the right incentives to ensure opt-outs are taken seriously and respected. While in principle, the law already provides for this option, it may be necessary to provide a clearer steer, possibly through an amendment, when additional damages are not just permitted but expected to be awarded.

While monetary damages are one remedy for unauthorised use, they are only one of the rights currently available to right holders. As important, and in practice used more often, is injunctive relief combined with the notice and takedown regime: in addition to damages for the past infringement, the right holders can prevent the continuation of the infringement, for instance by requesting that infringing content is removed from a website. Applied to the AI scenario, the rights holder might now be over-protected. Even if their contribution was only minimal, they could request that further infringing use be prevented by retraining the model, with the unlicensed material removed. This could result in significant costs for the provider.

Finally, a data scraper ignoring a machine-readable opt-out or other digital method that indicates the copyright status of an object could be seen as the circumvention of a Technological Protection Measures (TPMs). While under current law, section 29A provided a TPM exception for copies for text and data analysis, this is restricted to non-commercial research only. Under sections 296 to 296ZF of the amended Copyright Designs and Patents Act 1988 (CDPA), this circumvention would be a separate tort from that of copying the underlying asset, resulting in both civil and potentially criminal sanctions.

The combination of these remedies presents a paradox when applied to AI training: rights holders are both over- and under-protected. If for example hobby photographer who has put a picture on their website and protected it with an appropriate, machine-readable opt-out finds that despite the opt out, it was used to train a model, would; 1) Have next to no effective remedy for compensatory damages (de minimis or no economic loss); 2) But could, in theory, require the entire model to be taken offline and retrained; 3) While at the same time, depending on the technical nature of the chosen opt out, might instigate both civil and criminal sanctions against the data scraper.

We are not convinced that this is a desirable outcome for any of the parties; nor can it be remedied by mere evolving interpretation/clarification of the law, it will require legislative intervention.

DECaDE therefore recommends a “graded response” were sanctions are linked to the scale of the infringement overall (so that it matters not just if the current claimant’s rights were violated, but the scale of the unauthorised use across the board), the degree to which disregarding the opt-out was intentional or grossly negligent, and the overall profits that the model developer. These can then range from additional (punitive) damages that reflect the scale and severity of the abuse, right up to the duty to retrain the model in cases where a significant percentage of the training data was unauthorised.

Question 8: 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?

DECaDE emphasizes the importance of adopting a machine-readable format for rights reservation. AI training often involves large-scale reuse of content, making automated detection and enforcement of opt-out signals essential. Current AI methods—such as NLP—are not sufficiently reliable to parse free-text declarations at scale. Moreover, interoperability across different sources, tools, and platforms requires an open, patent-free technical standard. Therefore DECaDE advocates for the use of open technical standards to communicate AI opt out, and the government should work with rightsholders, industry, academia and standards organizations to determine the properties of such a standard, including that it is accessible and understandable to diverse creatives, who will have varying degrees of AI literacy. DECaDE suggests that one promising standard for unit-level AI opt-out is the technical standard of the Coalition for Content Provenance and Authenticity (C2PA) in that it meets these criteria is already gaining traction within the technology sector for AI opt-out.

3.3 Technical standards

Question 9: Is there a need for greater standardisation of rights reservation protocols?

Currently there are emerging standards for rights reservation, example being TDMRep [TDMRep] (at the site level) and C2PA [C2PA] (at the per asset, or unit level). In addition a number of open rights description languages (ORDL) have been developed for applications beyond rights reservation. In order to achieve scalable opt-out, there will be a convergence a small number of these, driven by organic adoption of standards by industry. DECaDE does not believe that the UK government should

proscribe the use of any preferred standard, since such standards are all emerging and the landscape of AI progress is fast moving making any such technical recommendation in legislation at risk of becoming quickly obsolete. Rather, DECaDE instead believes that as a matter of substantive law, the UK government should legislate more in the abstract, requiring that an open and free technical standard be used to indicate rights reservation.

References

[TDMRep] TDM Reservation Protocol (TDMRep). W3C Community Report. February 2024. Available at: <https://www.w3.org/community/reports/tdmrep/CG-FINAL-tdmrep-20240202>

[C2PA] Coalition for Content Provenance and Authenticity, Technical Specification v2.1. Publisher: Linux Foundation. Available at: <https://c2pa.org>

Question 10: How can compliance with standards be encouraged?

Standards have become an increasingly important regulatory tool. If they are required by law, legal certainty and safe interoperability is enhanced, but potentially at the cost of flexibility and innovation.

A possible solution is used in next-generation technology regulation such as the EU AI Act. There, certified compliance with standards creates an evidentiary (rebuttable) presumption of compliance with the substantive provisions of the law. It is rebuttable, which discourages “tick box compliance” that at best adheres to the letter, but not the spirit of the law. And it allows organisations for which the prescribed standard is not appropriate to find other solutions, though these have then the evidential burden to demonstrate the substantive compliance of their approach with the law.

A similar approach could be taken here. Appropriate industry standards such as the C2PA could create for both creators and model trainers presumptions of compliance. A rights holder who uses one of these certified standards will then be deemed to have communicated their decision in an appropriate way, and a model developer who can formally verify that their data collection method will always recognise such standards will be deemed to have complied (again subject to rebuttal, but with the burden of proof now on the rightsholder)

Question 11: Should the government have a role in ensuring this and, if so, what should that be?

Legislation could clarify whether stripping rights metadata—including opt-out signals—is tantamount to circumventing digital rights management (DRM), potentially making it a criminal offense. Such clarity would provide a strong deterrent against metadata removal. Beyond that, compliance may be encouraged by explicitly requiring platforms to preserve provenance or site-level metadata wherever it is present. While technical solutions such as watermarking or fingerprinting (e.g. Durable Content Credentials [NSA 2025]) can help identify content that has lost its metadata, they cannot entirely overcome its removal suggesting that legislation as well as technology should be combined to address the issue.

In addition to imbuing demonstrable compliance with relevant industry standards with some form of legal recognition, as described above, the government or one of its agencies could also play a role in curating and certifying some machine-readable licenses or the smart contracts that operationalise them. The main disadvantage of the “opt-out” approach is that it potentially creates a significant burden on creators and other rights holders. They will be used to license terms in ordinary language, which means they need to develop reasonable levels of computer literacy to make full and informed use of the automated, fine-grained licenses that we described above as a favoured model. Potentially, this puts them at a significant disadvantage vis-a-vis technology companies that have significant in-

house resources and also exposes them potentially to fraudulent or inadequate licenses that do not faithfully express their intent. A curated database of machine-readable license terms that certifiable translate correctly natural language terms into code could be of help here - such a database has been suggested by the French Ministry of Justice. These curated translations of terms into code would have to a) ensure that the code represents correctly the meaning of typical license terms and b) are legally valid (so for instance no smart contract that transfers full copyright, as this would fall foul of the legal “in writing” requirement for copyright transfer).

References

[NSA 2025] **Joint guidance on content credentials and strengthening multimedia integrity in the generative artificial intelligence era.** January 2025. US National Security Agency. <https://media.defense.gov/2025/Jan/29/2003634788/-1/-1/0/CSI-CONTENT-CREDENTIALS.PDF>

3.4 Contracts and licensing

Question 12: Does current practice relating to the licensing of copyright works for AI training meet the needs of creators and performers?

DECaDE believes that current practices could be substantially improved to meet the evolving needs of creators and performers, particularly around compensation for content used in AI model training. With the rise of generative AI, emerging provenance standards can now track not only which models generate a piece of content but also which data sources contributed to it. This makes it possible to identify specific creator inputs responsible for a given generation. Micropayment systems could be combined with provenance metadata, allowing creators to effectively hold an equity stake in AI models and receive ongoing, automated royalties. As AI models become more specialized—and in some cases, developed directly by artists themselves—having clear provenance and licensing metadata may increase the value of contributors’ work. This may also help mitigate “model collapse” caused by poor quality or synthetic training data, by incentivizing the inclusion of high-quality, authorized data. The ORA (Ownership, Rights, Attribution) framework [Collomosse 2024, Balan 2023a], and some commercial startups (e.g. bria.ai), have shown how granular attribution and compensation can be implemented in practice. Nonetheless, creators’ acceptance of such systems varies widely by scene and medium, with some viewing metadata and licensing choices as integral to their creative process [Liddell 2024, DECaDE 2025]. Addressing this diversity—and ensuring that compensation mechanisms genuinely benefit the people producing the original works—will be essential to any successful licensing reform.

The experience with NFTs indicates that sometimes, attempts were made to transfer through smart contracts full copyright, rather than merely creating a license of use. There could be scope of reforming the law of copyright assignation, to clarify or change the requirement that these have to be “in writing”. Some AI operators may prefer full assignation of right rather than a mere non-exclusive license to use.

[Collomosse 2024] **To Authenticity, and Beyond! Building Safe and Fair Generative AI upon the Three Pillars of Provenance.** J. Collomosse and A. Parsons. IEEE Computer Graphics and Applications (IEEE CG&A). 2024. <https://bit.ly/provtriad>

[Balan 2023a] **EKILA: Synthetic Media Provenance and Attribution for Generative Art.** K. Balan and S. Agarwal and S. Jenni and A. Parsons and A. Gilbert and J. Collomosse. In Proc. CVPR Workshop on Media Forensics (CVPRW). 2023. <https://arxiv.org/abs/2304.04639>

[Asnani 2024] **ProMark: Proactive Diffusion Watermarking for Causal Attribution**. V. Asnani, J. Collomosse, T. Bui, X. Liu and S. Agarwal. In Proc. Computer Vision and Pattern Recognition (CVPR). 2024. <https://arxiv.org/abs/2403.09914>

[Lidell 2024] **ORAgen: Exploring the Design of Attribution through Media Tokenisation**. F. Liddell, E. Tallyn, E. Morgan, K. Balan, M. Disley, T. Koterwas, B. Dixon, m C. Moruzzi, J. Collomosse and C. Elsdon. In Proc. ACM Designing Interactive Systems (DIS). 2024 <https://decade.ac.uk/articles/oragen-exploring-the-design-of-attribution-through-media-tokenisation/>

[DECaDE 2025] **ORAgen: Emerging Futures for Tokenisation and Digital Media Rights**. DECaDE/Digital Catapult **Technical Report**.. <https://www.digicatapult.org.uk/publications/post/oragen-emerging-futures-for-tokenisation-and-digital-media-rights/>

Question 14: Should measures be introduced to support good licensing practice?

DECaDE believes that through technical innovations new decentralized markets will emerge for creative content licensing re-use including re-use for generative AI. DECaDE believes that such innovations will emerge through market forces and do not require new specific measures to be introduced through legislation, though as noted above, the government can promote best practice, and also give technologically less sophisticated users reassurances through officially certified templates for automated licensing.

Question 15: Should the government have a role in encouraging collective licensing and/or data aggregation services? If so, what role should it play?

The government should encourage the development of decentralized licensing models that leverage open provenance standards, allowing creators to manage rights and compensation for AI training. DECaDE believes that decentralized platforms for creative rights management are a promising way to create a registry for licensing, including AI opt-out information. Such information may be expressed via open provenance standards and stored within a registry linked to digital watermarking or fingerprinting services. While Extended Collective Licensing (ECL) exists as a mechanism to simplify large-scale licensing, it carries risks to the UK competitive positioning for AI training if mass opt-outs occur without meaningful rightsholder engagement. Government should work with industry and academia to explore the development of scalable decentralized registries to this end, and consider emerging ideas around the incentivization of data collection re-use for AI training (see for example prototypes discussed in [Balan 2023a, Balan 2023b, Collomosse 2024]).

References

[Balan 2023a] **EKILA: Synthetic Media Provenance and Attribution for Generative Art**. K. Balan and S. Agarwal and S. Jenni and A. Parsons and A. Gilbert and J. Collomosse. In Proc. CVPR Workshop on Media Forensics (CVPRW). 2023. <https://arxiv.org/abs/2304.04639>

[Balan 2023b] **DECORAIT - DECentralized Opt-in/out Registry for AI Training**. K. Balan et al. Proc. Conf. Visual Media Production (CVMP). 2023. <https://arxiv.org/abs/2309.14400>

[Collomosse 2024] **To Authenticity, and Beyond! Building Safe and Fair Generative AI upon the Three Pillars of Provenance**. J. Collomosse and A. Parsons. IEEE Computer Graphics and Applications (IEEE CG&A). 2024. <https://bit.ly/provtriad>

Question 16: Are you aware of any individuals or bodies with specific licensing needs that should be taken into account?

DECaDE advises that special consideration should be made for the cultural sector and specifically to the policies and protocols associated with the digitisation and sharing of cultural heritage. This includes the consideration of indigenous data sovereignty and the growing recognition to find alternative data policies and licensing mechanisms for cultural heritage that reflect the values of the source communities (see for example Anderson & Christen, 2013). Doing so will bring a more encompassing understanding of copyright principles that including cultural and social conditions.

Crown copyright, Parliamentary copyright and open justice licenses may require minor tweaks. For open justice licenses e.g. currently excludes “computational analysis of the Information (including indexing by search engines)”. it is also questionable if in the context of AI training, the mandated “Use the current version of the Information “ is appropriate in all cases.

References

[Anderson 2013] **Chuck a Copyright on It.** J. Anderson and K. Christen. <https://sustainableheritagenetwork.org/digital-heritage/chuck-copyright-it%E2%80%9D-dilemmas-digital-return-and-possibilities-traditional-knowledge>

3.5 Transparency

Question 17: Do you agree that AI developers should disclose the sources of their training material?

We want to start with an important distinction: There is a difference to be made between “upfront disclosure” - where training data is made available outside the context of litigation - and “responsive disclosability” that becomes relevant after an infringement has been claimed.

There is no reason that we can see that *copyright* law requires a new upfront disclosure duty, and such a duty might be harmful given the commercial sensitivity of that data (although there may be exceptions in the case of safety critical systems). Open disclosure of the sources of data used to train an AI model may therefore make the UK less attractive to AI innovators.

However, if a rightsholder claims as part of litigation that a work created by an AI a) was trained on their work as input and b) has created as an output a work that is sufficiently similar to the training work to constitute a copyright violation, the AI provider may have to disclose the training data to refute condition a). There are of course tried and tested methods to disclose information safely and confidentially during litigation. There may also be new technological tools that can in a situation as the above formally verify that the works of the claimant have not been used, without disclosing any of the training data (we note though that the government is also currently consulting on the future regulation of computer-generated evidence, which could put constraints on such an automated verification).

While there are therefore no good reasons to require *for copyright* reasons upfront disclosure of training data, there can be other regulatory regimes that require such disclosure for a number of objectives. For some safety critical systems, it may be necessary to have appropriate reassurances about the reliability of the system. For AIs used in the justice system, there may be a need for disclosure to prevent biased decisions etc. Under the AI Whitepaper of the previous government, some of these future disclosure rules may come from the appropriate sectoral regulator, and may not have the form of primary legislation. It should be clear that while copyright law does not demand disclosure, it should

also not be possible to evoke copyright law to refuse disclosure when the relevant sectoral rules and regulations, including delegated law-making by regulators, so require.

Question 18: If so, what level of granularity is sufficient and necessary for AI firms when providing transparency over the inputs to generative models?

DECaDE believes instead that AI developers should attest to the use of appropriately licensed data, or perhaps to disclose the licenses of the data used to train the AI model. This will be helpful to determine the purposes for which the AI model may be used, and define the boundaries of liability between model creators and users. This will increase confidence in the re-use of AI models both commercially and for open source re-use. However, it is important to note that some artists may be willing to consent for their work to be included in an AI model, but not for this to be closed. In effect, creatives should not be implicitly forced to disclose that their work has supported an AI model.

Question 19: What transparency should be required in relation to web crawlers?

DECaDE does not believe it is helpful to indicate how data was acquired, since it is the data itself rather than its distribution that defines the capabilities, and eventual licensing, of an AI model trained on that data.

Question 20: What is a proportionate approach to ensuring appropriate transparency?

DECaDE believes that most Copyright concerns around AI pertain to generative AI in which derivative works are made from training data, rather than AI in general. As such, model transparency should be a consideration only for generative AI. DECaDE believes that emerging provenance standards should be leveraged to provide a means for transparently communicating the provenance of a model, specifically the licenses under which the data was used to train the model.

Question 21: Where possible, please indicate what you anticipate the costs of introducing transparency measures on AI developers would be.

DECaDE has not quantitative estimates however, combining machine readable rights declaration languages (e.g. ORDL) with provenance standards (e.g. C2PA) may produce automatable ways to aggregate data licensing information from training data in order to automate data governance and any transparency requirements around licenses of data used. Adhering to open standards and protocols will likely carry lower cost burden than creating and curating central repositories of AI opt-out signals. We also note here that AI developers, including some of the most innovative, in many cases may be individuals, including artists, researchers and creators themselves and consideration should be taken as to whether transparency costs and measures create unreasonable barriers to entry for innovations in this field.

Question 22: How can compliance with transparency requirements be encouraged and does this require regulatory underpinning.

DECaDE believes regulatory underpinning would be essential to ensure adoption of any transparency measures.

Question 23: What are your views on the EU's approach to transparency?

DECaDE acknowledges the intent behind Article 53(1)(d) but cautions against broad disclosure requirements that could expose commercially sensitive information and discourage AI development in the UK and EU. A layered approach—providing general details on model providers and dataset characteristics while ensuring opt-out compliance through machine-readable metadata—would better

balance transparency with innovation. Rather than exhaustive dataset disclosure, AI developers should attest to using properly licensed data and document compliance with rights reservations. The UK should align with EU principles while ensuring its approach remains pragmatic, enforceable, and globally competitive.

3.6 Wider clarification of copyright law

Question 24: 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 rightsholders are respected?

Clarity on the legal expectations around due diligence for opt-out would encourage confidence in training models in the UK. The creation of granular opt-out legislation (for example allowing a user to opt-out of some but not all forms of AI training) would contrast with coarser, more binary legislation (EU Copyright Act, Article 4) and so might enable greater freedom for model trainers to operate within the UK.

Question 26. Does the temporary copies exception require clarification in relation to AI training?

If the use of digital objects for training purposes were to fall under the TCE, training may not be an infringing act, and no opt out would be possible. So by inverse inference, the fact that this law creates an opt-out regime (and hence implicitly recognises that a right of the creator is potentially violated) means that model developers cannot rely on the TCE alone. However, there may be acts prior to the training that fall under the TCE. For instance, if a rights statement is digitally embedded in an object, the crawler has to make a temporary copy just to read that statement. Or it may be desirable to make a temporary copy of an asset in order to perform e.g. a visual feature extraction to index it within a registry for recording opt-out. For these cases linked to establishing training opt-out rather than training per se, a clarification would be helpful.

3.7 Encouraging research and innovation

Question 28. Does the existing data mining exception for non-commercial research remain fit for purpose?

Yes. The current exception for non-commercial research aligns with EU standards and remains appropriate. Introducing an opt-out for non-commercial use would impose additional constraints on academic research—limits that do not exist in the EU—and could thereby undermine innovation and the UK's competitive standing.

Question 29. Should copyright rules relating to AI consider factors such as the purpose of an AI model, or the size of an AI firm?

No, DECaDE believes that copyright rules relating to the size of an AI firm could encourage AI start-ups to sell or move outside of the UK when they grow, adding to the challenges of the UK tech sector at present.

Yes, DECaDE believes that copyright rules should consider the purpose of an AI model. Please see DECaDE's response to Q4. It is desirable to create granular levels of opt-out, for example enabling a rightsholder to opt out of certain AI uses and not others.

3.8 AI Outputs

Question 30. 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

From the recent consultation by the IPO, it seems that Section 9(3) is very rarely used in practice. The reason in our view is that it is very difficult to conceptualise a situation where a work is generated without *any* relevant human input. As a result, there are many works where a creator uses an AI, often to a significant extent, to create a work - but hardly any where there is no identifiable author. This seems to be a sound approach that works well in practice

Question 31. Do you have views on how the provision should be interpreted?

As above, given that the creativity threshold for copyright is generally not very demanding, an interpretation that puts high bars on “computer-generated works without an author” remains appropriate.

Question 35. Are you in favour of removing copyright protection for computer-generated works without a human author?

As noted above, there seem to be very few cases where Section 9(3) has ever been used. In almost all situations, it is better and more straightforward to consider AI as a creative tool rather than creative in its own right. On this basis, it should be permitted to copyright the results of applying AI creatively, but not to copyright content that is autonomously created by an AI model without manual creative input. Removing or keeping the protection therefore is unlikely to make a major change in practice.

3.9 Infringement and liability relating to AI-generated content

Question 38. Does the current approach to liability in AI-generated outputs allow effective enforcement of copyright?

Current enforcement mechanisms are limited because AI-generated content lacks reliable attribution i.e. to prove which model made an image, or which data is most responsible for a generated image. AI outputs that closely resemble training data pose risks of copyright infringement, yet proving infringement is difficult without robust provenance tracking or data attribution. Whilst the former is achievable with emerging standards (C2PA), data attribution technologies remain in an early stage of research.

A second, less discussed problem is the treatment of independent creation. Copyright, unlike a patent, does not protect against independent parallel creation. In the past, it was difficult (though not impossible) to plead independent creation if two works of sufficient complexity were identical or near identical. It is simply extremely unlikely that a second person would write all three volumes of the Lord of the Rings, without ever having been exposed to the original - much more unlikely than the “balance of probability standard requires. With AI-generated works, the situation is different, two users of the same AI model, by using only slightly different prompts, will often create the same or near-identical outputs. In this situation, it will be difficult to prove for either that their work really was independently created, and not a copied from the work of the other. It is difficult to predict just how much of a problem this will become, and if we will see disingenuous use of the “I did not copy, I simply used an AI that must have had similar training data” defence.

Question 39. What steps should AI providers take to avoid copyright infringing outputs?

It is possible to run safety checks on the outputs of generative AI models to establish whether generated content is very close to training data example, using similar technology to that employed by reverse image search engines. Such technologies may be helpful to AI providers avoid any penalties applicable for copyright infringing outputs and so the adoption of such technologies is likely to be driven by market forces rather than requiring legislation in its own right.

3.10 AI output labelling

Question 40. Do you agree that generative AI outputs should be labelled as AI generated? If so, what is a proportionate approach, and is regulation required?

Arguably, the potential for mis-use of generative AI outputs suggests value in such a label to fight misinformation, or to prevent certain types of fraud. On the other hand, there can be negative perceptions of creative value associated with overt generative AI use. DECaDE believes a balanced approach is therefore required, and suggests that legislation mandating labelling is limited only to certain categories of platform such as social media or news sites.

References

[Meta 2024] **Labelling AI Generated Images on Facebook and Instagram.** N. Clegg. <https://about.fb.com/news/2024/02/labeling-ai-generated-images-on-facebook-instagram-and-threads/>

Question 41. How can government support development of emerging tools and standards, reflecting the technical challenges associated with labelling tools?

The government should support industry collaboration and standardization efforts for AI-generated content labelling, particularly through open metadata frameworks such as C2PA and IPTC standards. This could include funding research and pilots that integrate these technologies across major digital platforms, ensuring robust adoption. Additionally, the government could convene regulatory sandboxes to test and refine labelling technologies in real-world applications, particularly in sectors such as journalism, social media, and creative industries.

To address technical challenges, AI-generated content labels must be persistent, machine-readable, and resistant to removal. The government should explore incentives for platform compliance and mechanisms to enforce metadata retention on content distribution platforms such as social media.

Question 42. What are your views on the EU's approach to AI output labelling?

The EU's approach to AI output labelling, which includes mandatory AI-generated content disclosure for foundation models, aims to improve transparency and mitigate misinformation risks. However, a one-size-fits-all approach risks overburdening AI developers and stifling innovation.

DECaDE supports context-dependent labelling, where AI-generated content is transparently marked, but without discouraging legitimate creative uses. Open standards such as C2PA already enable granular labelling of AI-assisted content, distinguishing between fully AI-generated media and AI-enhanced media. The UK should take a balanced approach, requiring provenance-based AI labelling for critical sectors (e.g., news and social media), while ensuring flexibility in creative and commercial applications.

3.11 Digital replicas and other issues

Question 43. 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?

Transparency is needed for creators to understand how their work is used in AI training and seek remuneration. A level of granularity in transparency regarding the specific ways in which their work is used in AI training is important for creators to assess the impact of such reuse and make informed decisions about whether to permit it, allowing them to maintain control and agency over their creative output.

3.12 Other emerging issues

Question 45. 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?

AI can generate derivative works both from its training data, or at inference time from data passed through it for adaption (for example, an AI model may be customized to create specific artistic styles, specific people or specific brand logos). To improve clarity, the government should establish legal guidance on liability, defining when AI-generated content constitutes a derivative work requiring licensing.

Question 46. 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?

Synthetic data poses risks to AI model quality through *Model Autophagy Disorder (MAD)* [Alemohammad 2023] which is a form of AI model collapse, where excessive AI-generated content in training datasets degrades performance over time. This issue is particularly pressing as AI-generated media increasingly populates the open internet. Identifying and distinguishing synthetic from human-created data is therefore commercially significant and vital for maintaining model reliability. Recent research explores how synthetic data presence can be mitigated during training [Alemohammad 2024] however this still leads to cost inefficiencies and potential quality issues.

To address this, content provenance standards such as C2PA should be strengthened to help AI developers track and filter synthetic data sources. Granularity in attribution is also crucial—rather than binary labels like *AI-generated* or *human-generated*, provenance tools should capture the degree of AI modification. The government should support open, interoperable provenance standards and ensure their adoption across digital platforms, enabling creators to trace how their work is used and mitigating the risks of synthetic data pollution in AI training.

References

[Alemohammad 2023] **Self-consuming generative models go MAD**. S. Alemohammad, J. Casco-Rodriguez, L. Luzi, A. I. Humayun, H. Babaei, D. LeJeune, A. Siahkoohi, and R. Baraniuk. <https://arxiv.org/abs/2307.01850>

[Alemohammad 2024] **Self-Improving Diffusion Models**. S. Alemohammad, A. I. Humayun* , A. Agarwal, J. Collomosse, and R. Baraniuk. <https://arxiv.org/pdf/2408.16333>

Question 47. What other developments are driving emerging questions for the UK’s copyright framework, and how should the government respond to them?

Decentralized technologies such as distributed ledgers technology (DLT), or colloquially ‘blockchain’, enable the creation of secure decentralized databases, without relying upon the trust of any individual organization or third party. Many commercial use cases have explored DLT for supply chain traceability, and there are clear parallels between this and the creative supply chain for digital content.

Emerging standards such as C2PA provide a machine readable foundation to describe content ownership, provenance and rights. However this information is always removable from assets, suggesting a need for a decentralized database to recover it. DECaDE therefore believes that an important piece of the puzzle in solving general copyright and attribution for creative content involves DLT, which in turn raises questions around the representation of digital identity in such systems, and so legislative questions around platform governance that are not immediately obvious when considering of AI and Copyright.

END