BRAID Researchers' Response to the UK AI & Copyright Consultation

This response was prepared by researchers in the Bridging Responsible AI Divides (BRAID) Programme. BRAID is a UK-wide programme dedicated to integrating Arts and Humanities research more fully into the Responsible AI ecosystem, as well as bridging the divides between academic, industry, policy and regulatory work on responsible AI (https://braiduk.org/).

Views are our own - as researchers/academics - and do not reflect those of our institutions, the organisations partnered with us in our projects, the BRAID programme as a whole, or other BRAID researchers. Throughout this response there may be varying views from the responding researchers, and we have flagged these where relevant, as well as points where there has been explicit agreement. Researchers who have not contributed to a particular answer do not take a view on that specific point.

The researchers responding are (in alphabetical order by surname):

Prof. Nick Bryan-Kinns is Professor of Creative Computing at the University of the Arts London and a BRAID Research Fellow. His research explores new approaches to interactive technologies for the Arts and the Creative Industries through Creative Computing and Human-Centred AI. He is a Fellow of the Royal Society of Arts and the British Computer Society, and Senior Member of the Association of Computing Machinery. Bryan-Kinns' BRAID fellowship with the BBC explores how eXplainable AI (XAI) approaches can be deployed in creative practice within the BBC and Creative Industries more widely to address Responsible AI (RAI) concerns of transparency, intelligibility, and accountability.

Prof. Mercedes Bunz is Professor in Digital Culture and Society at King's College London and a BRAID Research Fellow. Her research explores the calculation of meaning through machine learning systems. She co-leads the Creative AI Lab, a collaboration with the Serpentine Gallery, London. Her BRAID fellowship explores AI art beyond the gallery: exploring the capacity of cultural institutions to impact tech policy.

Dr Clementine Collett is a BRAID Research Fellow at the University of Cambridge and Principal Investigator on the project 'Co-designing responsible technology and policy for the impact of generative AI on the writing and publishing of the novel' in partnership with the Institute for the Future of Work. This fellowship aims to develop policy and design recommendations surrounding generative AI and its impact on novel writing and publishing.

Dr Bahareh Heravi is an Associate Professor (Reader) in AI and Media at the Institute for People-Centred AI at the University of Surrey and a BRAID Research Fellow. She specialises in the use of AI in journalism and media, ethical and responsible AI, and data storytelling. Heravi is a BRAID Fellow with the BBC, researching Responsible AI Literacy in News and Journalism. She has extensive international experience working with journalists and news media organisations on the use of data and AI for reporting and storytelling.

Dr Kate Miltner is a Lecturer (Assistant Professor) in Data, AI, and Society in the Information School at the University of Sheffield and Co-Director of the University of Sheffield Digital Society Network. Her research explores issues of power and inequality in digital systems,

institutions, and cultures. Miltner is a Co-I on the BRAID-funded scoping project Framing Responsible AI Implementation & Management (FRAIM), which brings together cross-sector perspectives on organisational RAI policy and process to scope key stakeholders, shared values, and actionable research needs for building the evidence base on implementing and managing RAI.

Dr Caterina Moruzzi is a Chancellor's fellow in Design Informatics, School of Design, University of Edinburgh, and a BRAID Research Fellow. Her research lies at the intersection between the philosophy of art, human and artificial creativity, and the philosophy of Artificial Intelligence, focusing on modes of shared agency and creativity between humans, data, and technology. As BRAID Research Fellow, she leads a collaboration with Adobe to promote the responsible integration of AI tools into creative practices.

Dr Anna-Maria Sichani is a BRAID Research Fellow and Research Associate in Digital Humanities at Digital Humanities Research Hub, School of Advanced Study, University of London. Her fellowship 'Responsible data, models and workflows: Responsible AI digital skills provision for the cultural heritage community', in partnership with The Turing Institute, is looking at embedding responsible AI literacy skills across the cultural heritage community to empower informed, responsible and ethical use of AI and machine learning.

Dr Bev Townsend is a BRAID Research Fellow at the York Law School, University of York. Her work focuses on the regulation, law, and ethics of emerging technologies, including socio-technical resilience, autonomous systems, robotics, data, and related medical technologies. Her BRAID fellowship, in partnership with Microsoft Research, seeks to explore socio-technical harm and medical AI adoption and to understand how to identify, reduce, and mitigate such harm through policy intervention.

Dr Paula Westenberger is a BRAID Research Fellow, Senior Lecturer in Intellectual Property Law at Brunel University of London, and member of the Brunel Centre for AI: Social and Digital Innovation. Her BRAID project 'Responsible AI for Heritage: copyright and human rights perspectives', in partnership with RBG Kew, focuses on how intellectual property law, particularly copyright, can responsibly advance the use of AI in heritage research.

B.4 Policy options

Question 1. Do you agree that option 3 is most likely to meet the objectives set out above?

No. The researchers named in this answer do not believe that option 3 is likely to meet the objectives.

Dr Clementine Collett explains some of the issues with option 3:

Uncertain impact:

An opt out model, as proposed in option 3, compromises on protecting the creative industries, without recognition of its evident growth and contribution to our economy and culture. Instead,

the proposed data mining exception with rights reservation prioritises the technology industry, the growth and impact of which is uncertain. Simultaneously, this proposal might damage our creative economy, endanger livelihoods, and harm elements of our national identity and soft power.[1] Generative AI has been shown to detrimentally affect writing and translation, for example.[2] Collett's work at the University of Cambridge has found in an initial survey of around 150 authors and publishers that some have already lost work from the use of AI in the industry, with particular impacts on translators, ghostwriters, and copywriters.

False premise:

The Government should consider that the options they have presented around AI and Copyright are built on the assumption that more data will produce better and more capable large language models (LLMs). This is not necessarily the case, as has been shown.[3] The Government should be aiming to innovate models which use less data and therefore less energy; as we know, scaling models also scales environmental impact.[4] In addition, models which use less data would still be able to cater to supporting basic AI functions that people rely on so frequently without compromising on control of IP.

Practical and technical issues:

There is a lack of clarity about how an opt out model would work in practice and no impact assessment to consider its implications. There is uncertainty about whether reserved works could still be accessed and used as training data. For example, models could be trained before the opt out is actioned, or there might be unforeseen ways to get around machine-readable reservations, or it might be difficult to reserve material from all web crawlers.

Moreover, we see from the EU that this opt out model outlined in articles 3 and 4 of the Copyright in the Digital Single Market (CDSM) Directive in 2019 is not effective as it is currently unclear how this opt out works in practice.[5] This does not strengthen rights holders' position and also has been found to potentially hinder the advancement of AI developments in the EU.[6]

[1] Mammen et al. (2024). *Creativity, Artificial Intelligence, and the Requirement of Human Authors and Inventors in Copyright and Patent Law.* https://www.cs.ox.ac.uk/tom.melham/pub/Mammen-2024-CAI.pdf

[2] Teuloff, Ole; Einsiedler, Johanna; Kässi, Otto; Braesemann, Fabian; Mishkin, Pamela; Maria del Rio-Chanona, R. (2025). "Winners and losers of generative AI: Early Evidence of Shifts in Freelancer Demand". *Journal of Economic Behavior & Organization.* https://doi.org/10.1016/j.jebo.2024.106845

[3] Bender, Emily; McMillan-Major, Angelina; Gebru, Timnit; Shmitchell, Shmargaret (2021). "On the Dangers of Stochastic Parrots: Can Language Models be Too Big?" FAccT '21, ACM, NY, New York. https://doi.org/10.1145/3442188.3445922

[4] Ibid.

[5] Keller, Paul, and Warso, Zuzanna (2023). "Defining Best Practices for Opting Out of ML Training". Open Future Policy Brief 5. https://openfuture.eu/wp-content/uploads/2023/09/Best-_practices_for_optout_ML_training.pdf

[6] Ziaka, Gina Maria (2014). "The text and data mining opt-out in Article 4(3) CDSMD: Adequate veto right for rightsholders or a suffocating blanket for European artificial intelligence innovations?" *Journal of Intellectual Property Law & Practice.* 19(5). https://doi.org/10.1093/jiplp/jpae025

Compromising on control:

Some technology companies may be keen to avoid giving granular details of the data used to train their models to protect their own IP, even though the proposed text and data mining exception would give them access to large amounts of IP which might have been used prior to opt out, without full and informed control of works, or even despite it. Access to creatives' work should not undermine authors' legitimate rights to control their own IP. An opt out model compromises on this control. It would require creatives to be responsible for knowing about, and acting upon, these changes. The administrative burden of opting out is huge as it must be done on all copies of artists' work, past and future, which is also unfair on less resourced creatives.[7] Some authors may not realise that they need to opt out, may not know how or why it is important to do so, or they may not be able to opt out. Indeed, evidence shows that there is generally low take-up of opt out models, but this does not tally with how we see creators' and rights holders' feel about how their content should be licensed and used.[8] The ALCS survey on over 13,500 authors showed that 91% of authors feel they should be asked permission if their work is used to train AI.[9]

Caterina Moruzzi seconds Collett's view and adds that the disadvantages of an opt-out approach to text and data mining practices (TDM) include the following: 1. It puts the burden of opting out on creators and rights holders, assuming that opt-out options are known and accessible by right holders who wish to exert this right, 2. It may expose them to inadequate licences that do not align with their intent, assuming that right holders possess the legal and computer literacy to understand the details of these licences. 3. Fundamentally, it assumes that by default content created by right holders can be used for TDM, impacting on creator's moral rights.

Bahareh Heravi adds that another issue with the opt-out approach is how it is integrated into various platforms, and the extent to which the right holders can have control over that. For example, if an artist shared their visual artwork on a platform like Instagram, would there be an option for them to opt-out, or to set mechanisms to reserve their rights, e.g. for licensing or remuneration in the case of access for AI training?

Kate Miltner seconds Collett's, Moruzzi's, and Heravi's views and adds that even when individuals are eligible to make 'take-down' requests under current legal frameworks (and have the resources and know-how to do so), these requests are not always honoured or even acknowledged despite legal obligations/requirements.[10] Close attention would need to be paid to enforcement mechanisms, particularly in lop-sided power relations between individual creators/artists and multinational corporations. Opt-out mechanisms are the tool of choice for model or platform owners, some of whom have created (or are in the process of creating) opt-out tools for their individual models (e.g., OpenAI's Media Manager), but it is unclear how

[7] Newton-Rex, Ed (2024). "The insurmountable problems with generative AI opt-outs." https://static1.squarespace.com/static/5cc5785816b6406e50258c5c/t/67368c12cc35b5469feb0bfd/1731628050768/The+insurmountable+problems+with+generative+AI+opt-outs.pdf
[8] Ibid.
[9] ALCS (2024). "A Brave New World? A survey of writers on AI, remuneration, transparency and choice". https://d16dqzv7ay57st.cloudfront.net/uploads/2024/12/A-Brave-New-World-_ALCS_AI_Report.pdf
[10] Harlan, E. and Brunner, K (2023). 'We Are All Raw Material for AI', *BR* https://interaktiv.br.de/ki-trainingsdaten/en/

these tools operate or what compliance mechanisms will be put in place.[11] If dominant model operators and owners are voluntarily introducing opt-out mechanisms, it begs the question as to what an opt-out mechanism adds to copyright protections for artists and creators.

Nick Bryan-Kinns seconds Collett's, Moruzzi's, Heravi's and Miltner's points. In addition, in Bryan-Kinns' view the intent of putting creative content on the internet prior to the mass commercialisation of deep learning models in the early 2020s was not for it to be used without permission to train AI models that had not yet been invented. Moreover, much work on the internet is made available with specific copyright notices or terms displayed on the web pages already and it is unclear the extent to which these notices of copyright would be respected or ignored in the opt-out approach. In addition, point 80 in the consultation highlights a major problem with the opt-out model - the inability of rights holders to prevent the use of other copies of a work (where rights are not reserved) in AI model training. An opt-in approach would remove this problem.

Paula Westenberger flags also the unclear and problematic "lawful access" requirement, which she further analyses in answer to Question 28 below. Westenberger is also concerned about the potential impact of the proposed exception in relation TDM research, in view of the current issues with s29A CDPA, which are discussed in Question 28. As discussed in her forthcoming paper, "the boundaries between any new commercial TDM exception and the non-commercial research exception must be carefully and clearly delineated."[12]

C.1 Exception with rights reservation

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

The researchers named in Question 1 do not support the introduction of an exception along the lines outlined, in view of the comments above.

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?

In Caterina Moruzzi's view, granularity is a crucial aspect. There is an opportunity to adopt a more granular approach in the Copyright legislation that responds to the creators' need for a more nuanced control over the use of their works. For example, the majority of existing opt-out mechanisms are binary, allowing to either allow training or not, without the possibility for right holders to have more granular choices on whether to allow their work to be findable in search results, and/or allowing their content to be used for training.

This granularity should however be transparent and accessible. Studies conducted with expert users involved in different aspects of the creative media supply chain [Moruzzi et al., under

[11] Knibbs, K. (2024). "OpenAI's New Tool Will Give Artists Control Over Their Data—But It's Unclear How". *Wired*. https://www.wired.com/story/openai-olive-branch-artists-ai-algorithms/

[12] Westenberger, P. and Farmaki, D. (2025). "Artificial Intelligence for Cultural Heritage Research: the Challenges in UK Copyright Law and Policy". Available at SSRN: https://ssrn.com/abstract=5153757 or http://dx.doi.org/10.2139/ssrn.5153757.

review], evidence the need for copyright and provenance systems and regulations to consider how to balance the need for granular control and the additional burden that granularity may pose for cognitive load. Without clear, accessible, and transparent indication of how AI tools may use their work, the potential benefits of a more granular approach would be diminished.

Collett suggests the Government should consider an opt in model with underpinning transparency and auditing measures, along with a strengthening of copyright law as discussed in option 1 of the consultation. This should include clarification regarding legacy contracts where rights have been transferred to another party, historical infringement, full transparency regarding data reporting, centralised enforcement mechanisms, and regulation on dynamic licensing and fair remunerative standards, for example through a royalty-based model.[13] [14] This should be developed in collaboration with the creative industries. This would not rule out innovation from technology companies, but it would ensure that there is less compromise on creative's control over their own work and ensure there is fair, remunerated, and informed licensing of these works. Regulation does not mean no room for innovation and growth, indeed, AI companies are often more open to regulation than commonly perceived.[15] This provides them with a stamp which they can use to assure their clients, and it enables them to avoid reputation and financial costs of being sued.

Heravi seconds Moruzzi's and Collett's view.

Bryan-Kinns seconds Moruzzi's and Collett's views. Moreover, in Bryan-Kinns' view, an opt in mechanism would likely be less technically complex than an opt out mechanism as described in the consultation. An opt in mechanism may also produce the environment for a more entrepreneurial approach to content creation for AI model training by encouraging creative practitioners to create content specifically to be licensed to AI model training - this could in turn become a market differentiator for AI model training in the UK and provide a much needed source of income to creative industries.

Miltner seconds Bryan-Kinns, Moruzzi and Collett's views and adds that the encouragement of creative practitioners to create content licensed specifically for AI model training could also address concerns around 1.) algorithmic monoculture (the homogenisation of outcomes resulting from algorithmic production and curation) and 2.) the lack of diverse representation within AI training datasets. If creators (particularly those from minoritised or marginalised communities) were incentivised to contribute a wide-range of self-representational content to training datasets, this could help address some of the harmful and (incorrect) biases currently represented within major AI models.[16] Furthermore, diversifying training data through appropriately remunerated and consensual means could help break the recursive loop of AI-

[13] Lucchi, Nicola (2023). "ChatGPT: A Case Study on Copyright Challenges for Generative Artificial Intelligence Systems". *European Journal of Risk Regulation.* 15(3). https://www.cambridge.org/core/journals/european-journal-of-risk-regulation/article/chatgpt-a-case-study-on-copyright-challenges-for-generative-artificial-intelligence-systems/CEDCE34DED599CC4EB201289BB161965.

[14] Jacques, Sabine, and Flynn, Mathew (2024). "Protecting Human Creativity in AI-Generated Music with the Introduction of an AI-Royalty Fund".*GRUR International. 73(12).* https://doi.org/10.1093/grurint/ikae134

[15] Collett, C. (2024). *Framing gender bias in the design of AI recruitment technology*. [Doctoral Thesis, University of Oxford]

[16] See, for example, Jake Elwes' Zizi Project: https://www.jakeelwes.com/project-zizi-2019.html

generated content production, where AI-generated outputs feed back into those same models at scale, potentially solidifying the dominance of a narrow range of homogenous styles and genres.[17]

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.

In Bryan-Kinns' view there would be negative influence on new generations of artists and creative practitioners graduating from our universities. The requirement for an opt out would cause concern and confusion about how to opt out for inexperienced creative practitioners with little or no legal or rights representation. This would likely lead to a reluctance to publicly promote and distribute their creative works making it harder to generate interest in their works, build reputation and audience, and earn income from their creative works. Bahareh Heravi and Clementine Collett share this view.

Anna-Maria Sichani notes that the proposed TDM exception with rights reservation, by supporting measures on transparency, could partially address Galleries, Libraries, Archives and Museums (GLAM) institutions' growing concerns around use of their open collections for profit-driven AI without consent or compensation. Since late 2022, the rise of commercial AI companies scraping open GLAM collections, among other openly available datasets, for model training has led institutions to reconsider open data policies such as the closing down of National Library of the Netherlands' collections for commercial Generative AI companies[18]. Historical collections — often out of copyright but structurally biased — risk being repurposed for profit-driven AI without consent or compensation, fueling debates on balancing openness with creators' rights.

A rushed and unconsidered implementation of the new proposed TDM exception could create uncertainty, fear and resistance among GLAM institutions and professionals. Many of these institutions have long advocated for open access and the responsible use of digital collections, yet they also face increasing concerns about the unchecked use of their data by AI companies. Without clear guidance, standardised rights reservation mechanisms, and strong enforcement of transparency obligations, institutions may feel pressured to restrict access to their collections altogether to avoid potential misuse by AI companies. This could not only hinder advancements in digital humanities and AI-driven research but also set back the Open GLAM movement, which has championed accessibility and openness for digital cultural heritage over the past two decades. Heravi and Westenberger share this concern.

Westenberger adds to the point on rights reservation for GLAM's the contrasting view that:

> "cultural heritage institutions are not entitled to make the decision to apply a data mining opt-out. They can only do so if copyright exists, and they hold the copyright. But even if they do, it is worth wondering whether opting out supports their objectives. In a way, blocking the possibility to use cultural heritage data seems counter to the

[17] Miltner, K. M., & Highfield, T. (2024). The Possibilities of 'Good' Generative AI in the Cultural and Creative Industries. *The British Academy.* https://www.thebritishacademy.ac.uk/publications/the-possibilities-of-good-generative-ai-in-the-cultural-and-creative-industries/

[18] https://www.kb.nl/en/ai-statement

> mission of publicly funded cultural heritage institutions. Isn't contributing trustworthy qualitative information and fighting misinformation and bias (in algorithms) more in line with their objectives?"[19]

Westenberger is also concerned about the potential impact of the proposed exception in relation to TDM research, in view of the current issues with s29A CDPA (TDM for non-commercial research), which are discussed in Question 28 and her forthcoming paper.[20] Westenberger's view is that we should avoid the situation that research uses (that should be covered by the non-commercial research exception, which cannot be overridden by contract) end up being regulated under a commercial TDM exception with opt out, as this could have the possible result of biasing and affecting the quality of research as a result of incomplete datasets. The boundaries between any new commercial TDM exception and the non-commercial research exception must be carefully and clearly delineated.[21] S. 29A must be clarified and expanded before any new TDM exception is included in legislation, as discussed in Question 28 below.

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

Collett argues that a rights reservation model would be technically difficult to enforce. It has been argued that in the case of unit-based opt out schemes, metadata put on top of content can easily be removed, and besides, there is no way to add metadata to text.[22] Some models might also be trained before an opt out is actioned, there may be ways to get around machine-readable reservations, or the author might not have known to opt out of reserving material from all web crawlers.[23]

Additionally, it is difficult to protect downstream versions of the work.[24] This is because most opt out schemes are location-based, for instance, robots.txt which contains information about how the web domain can be accessed by web-crawlers. However, the work of a creative is likely to appear in many more places over which they have no control, e.g. various social media and web-based/digital platforms, or in the case of a book, could even be illegally copied onto the web without their knowledge.[25] [26] This, alongside the high-level summary approach to transparency proposed by the Government would mean that some copyrighted works which have been opted out could still be used within AI model training.

---

[19] Matas A (2024). "AI 'opt-outs': should cultural heritage institutions (dis)allow the mining of cultural heritage data?". https://pro.europeana.eu/post/ai-opt-outs-should-cultural-heritage-institutions-dis-allow-the-mining-of-cultural-heritage-data

[20] Westenberger, P. and Farmaki, D. (2025) "Artificial Intelligence for Cultural Heritage Research: the Challenges in UK Copyright Law and Policy". Available at SSRN: https://ssrn.com/abstract=5153757 or http://dx.doi.org/10.2139/ssrn.5153757.

[21] Ibid.

[22] Newton-Rex, Ed (2024). "The insurmountable problems with generative AI opt-outs". https://static1.squarespace.com/static/5cc5785816b6406e50258c5c/t/67368c12cc35b5469feb0bfd/1731628050768/The+insurmountable+problems+with+generative+AI+opt-outs.pdf

[23] Ibid.

[24] Ibid.

[25] Ibid.

[26] Reisner, Alex (2023). "These 183,000 books are fuelling the biggest fight in publishing and tech". *The Atlantic.* https://www.theatlantic.com/technology/archive/2023/09/books3-database-generative-ai-training-copyright-infringement/675363/

In contrast, with an opt in system, this could be controlled via a regulated, centralised repository where authors can see if their works have been used. This would not work, however, for an opt out scheme, where the directory was scanned by automatic content recognition (ACR) for opted out content, as content could easily be modified and therefore go undetected.[27]

Bryan-Kinns and Miltner second Collett's view.

Sichani argues that AI developers and practitioners seeking to employ datasets for model training often face significant ambiguity regarding rights reservations. In the absence of a robust regulatory framework, many rights holders have resorted to what can be described as "activist positions" to assert their opposition, including public statements, blanket opt outs in press releases, industry-led efforts, registry-based opt outs and reservations of rights in website terms and conditions. While these measures can create awareness and encourage ethical AI practices more widely, they do not provide guaranteed legal protection against unauthorised use of their content for commercial AI model training.

From the AI developer's perspective, ethical considerations might encourage the recognition of rights reservation, but the reality is that without a legal framework and a standardised and machine-readable mechanism, there is no reliable way to detect and respect these reservations at scale. Most AI training data is nowadays collected through automated processes—crawlers, scrapers, and large-scale data aggregation tools. If an opt out is not expressed in a way that these automated systems can recognise and process, it is effectively invisible.

Bryan-Kinns, Heravi and Miltner second Sichani's view.

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.

Sichani highlights ongoing efforts to develop protocols and opt out mechanisms for Text and Data Mining (TDM) exceptions, categorised into registry-based solutions, metadata standards, and robots.txt-type technologies. Notable among these are the World Wide Web Consortium's (W3C) TDM Reservation Protocol (TDMRep)[28] and the Content Authenticity Initiative (CAI) which has introduced a C2PA technical specification (version 1.3)[29] allowing rights holders to signal a "do not train" intent for different types of machine learning (ML). Additionally, Spawning's ai.txt[30] file offers a promising approach by enabling selective access restrictions to site content and integrating opt out data from external sources, warranting further exploration.

[27] Newton-Rex, Ed (2024). "The insurmountable problems with generative AI opt-outs." https://static1.squarespace.com/static/5cc5785816b6406e50258c5c/t/67368c12cc35b5469feb0bfd/1731628050768/The+insurmountable+problems+with+generative+AI+opt-outs.pdf
[28] https://www.w3.org/community/reports/tdmrep/CG-FINAL-tdmrep-20240510/
[29] https://c2pa.org/specifications/specifications/1.3/specs/C2PA_Specification.html
[30] https://spawning.ai/ai-txt

Although different opt out solutions could coexist—combining domain-level and metadata-level reservations—these technologies and efforts do not yet fully meet the needs and expectations of all stakeholders and their implementation will not be necessarily straightforward or easy. Challenges include the proliferation of multiple copies of works ("duplicates" or altered versions) across the web, requiring right-reservation to be expressed for each copy hosted on sites beyond the right holder's control. Additionally, ensuring consistency in opt out expressions for joint works—such as collections or works of a similar type—is crucial to making the TDM exception workable at scale, a crucial parameter for cultural heritage institutions with large collections. While creators and other rights holders will be entrusted with the ability to reserve their rights, the validity of such machine-readable reservations will be difficult to assess unless there are measures in place to prevent false claims by persons or entities other than the actual rights holders. An issue that is not yet very well addressed is to what extent right-reservation mechanisms can be applied in retrospect through a "machine unlearning" approach[31], meaning for content that has already been infringed and used as training data for AI training. Another key concern for the opt out to be exercised and observed in its fullness remains the lack of transparency of the training data which makes it difficult to verify AI providers' respect of the right-reservations duly expressed by right holders. Related to this, Keller & Warzo (2023) highlight that there is often a time lag between opt out decisions and their practical effect as there can be significant time delays of up to 1.5 years between the initial selection of training data and the release of a trained model.[32] Since the adoption of universally followed best practices for opting out copyrighted works from ML training would only be reflected in trained models after a significant delay, this raises concerns among creators and rightholders, especially as once ML models are trained they cannot (or it is extremely complex to) 'unlearn' information derived from specific works present in their training data.

C.2 Technical standards

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

Sichani claims that standardisation may be the optimal choice for opt out mechanisms, however its implementation may be challenging due to the varying readiness and capacity of right holders, the fast-moving fluidity of ML technologies and regulatory standards, and the absence of a global standardisation body to enforce widespread adoption. However, at a minimum, machine-readable rights reservations protocols should ensure convergence toward minimum requirements and specifications: should be based on open standards, freely available, and designed to be user-friendly and well-documented for creators and other right holders, maintaining backward compatibility, and balancing clarity, the ease of expressing opt outs, and granularity towards both the expression of opt out reservations on different levels of content (per record/work, collection of works, website level etc.) and the ability to differentiate (or express different policies) for different ML uses (search and discovery, either traditional indexing or new AI-based search, reuse etc).

[31] Jie Xu, Zihan Wu, Cong Wang (2024). " Machine Unlearning: Solutions and Challenges" https://arxiv.org/pdf/2308.07061

[32] Paul Keller & Zuzanna Warso (2023). "Defining best practices for opting out of ML training", Open Future https://openfuture.eu/wp-content/uploads/2023/09/Best-_practices_for_optout_ML_training.pdf

Question 10. How can compliance with standards be encouraged?

In Sichani's view, it is essential to establish technical standards and best practices for expressing opt outs in compliance with robust legislation. These practices should cater to both the supply side—offering clear guidance to creators and rights holders on expressing and specifying their desired opt outs—and the demand side—encouraging ML developers to recognise and uphold these reservations with transparency. Additionally, for some cases, machine-readable opt outs could complement collective remuneration schemes, encouraging rights holders to opt out as a pathway to compensation.

In Bryan-Kinns' view, compliance with technical standards on rights management by technology companies can only be encouraged through opt in mechanisms which companies necessarily need to engage with to undertake AI model training. The proposed exception approach (opt out mechanisms) do not encourage companies to engage with technical standards for opt out descriptions beyond simply ignoring content which has some opt out attached to it.

Collett and Miltner second Bryan-Kinns' view.

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

In Mercedes Bunz's view, the government's understanding of the constrained functioning of /robot.txt is an important point. The possibility of artists and content producers to make use of a more granular, machine-readable control is to be supported. A separate '*do not train' file* besides /robots.txt will allow crawled content and works to be indexed in search, so they remain findable – important for artists and content producers; at the same time, the file indicates that the material is not allowed to be used for commercial AI training. Spawning.AI's /ai.txt mentioned by the government seems indeed a promising solution adopted by some AI companies but also commercial players such as Shopify, but also by content hosts such as Wordpress, Squarespace, Wix, etc.

While this standard naturally needs to be agreed upon on a global level, there are supportive mechanisms this government can initiate. For example, writing the ability to opt out into law will strengthen the content producer's position as it has clarified the legal situation; at the moment, the practice to opt out is mainly supported by EU's Text and Data Mining Article 4 exception. In general here and elsewhere, future UK law should be thought in conjunction with technical standards and see /ai.txt's role as one mechanism that allows to respect the wishes of artists and content producers.

But creating a standard does not mean a standard is adopted. Here, the government should support technical and legal research into *do-not-train* files as well as support social research and design thinking that studies AI literacy of 'do not train' files. Studying the understanding and adoption of this technical solution is essential in order to have technical solutions not just available but also accessible for artists and content creators and to bridge technical divides. Creative Commons has shown that such a thing is possible for licenses usually linked to complex legal language; now that code can be law, as Lawrence Lessig once stated, translating similar efforts to technical standards is a challenge that needs the government's

support. Steering activities through research funding that links arts and humanities with computational efforts is one effective measure that should accompany the technical standard.

Sichani claims that making technical standards for opt outs a legal requirement rather than an optional choice would ensure compliance, but the government should also support adoption by appointing a credible entity/body to provide authoritative guidance on implementing machine-readable rights reservations effectively while enhancing accessibility and transparency in these issues across sectors and stakeholders. Relatively small but impactful interventions might include 1) explanatory, user friendly online resources (guides) of data sources, protocols, and standards for machine-readable rights reservations that comply with legislation that would guide rights holders in opting out of ML training and help developers understand compliance obligations, 2) community support events throughout the course of the introduction and implementation of the standards in order to familiarise and support all stakeholders.

In Bryan-Kinns' view opt out rights management requires revocable rights management per creative output. This cannot be achieved with a robots.txt file but instead would likely require addition of rights metadata to creative outputs and/ or registration of the rights in a centralised ledger such a blockchain. The technical complexity and administrative overhead of such an approach would restrict the use of opt out rights to rights holders with sufficient technical and administrative capacity such as major music recording rights holders, likely preventing opt outs from being a feasible option for individual creative practitioners, and small and medium sized creative organisations.

In Heravi's view, whatever the technical approach and standard, eventually there is a need for enforcement measures to ensure platforms provide simple and clear implementations of the opt in or opt out mechanisms, to enable all content holders to confidently and safely share their content while reserving their rights, without fear of being unlawfully used for training, and with providing mechanisms for licencing and remuneration, if applicable.

C.3 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?

Collett argues that current practice relating to the licensing of copyrighted works for AI training does not meet the needs of creators and performers. There has been much evidence that authors' work has already been used without their permission and without any licensing or remuneration.[33] In the recent ALCS survey, only 7% of those who knew their works had been used to train AI gave permission for this use, while 91% felt that they should be asked for their permission. Many authors are unaware of the extent to which their work is used, and the ALCS

[33] Milmo, Dan (2025). "Zuckerberg approved Meta's use of 'pirated' books to train AI models, authors claim". *The Guardian.* https://www.theguardian.com/technology/2025/jan/10/mark-zuckerberg-meta-books-ai-models-sarah-silverman

found that 77% of authors don't know their works are being used to train AI or not. 92% said they would want to receive compensation for any historic use of their work to train AI.[34]

This historical infringement is something which must be addressed, particularly as an opt out model, as proposed by the Government, would be introduced after many AI companies have started training or releasing their models. Even if these companies did have to re-train their models, they might do it on synthetic data which is, itself, trained on copyrighted works.[35]

Bryan-Kinns, Heravi, and Miltner second Collett's view. Miltner also adds that not only do artists, organisations and trade unions overwhelmingly prefer opt in models where artists can choose for their work to be included in datasets,[36] but artist collectives have already started developing "non-infringing" datasets in response to current industry practices (e.g., Spawning's PD12m and the upcoming Public Diffusion set).[37] Artists have also resorted to using dataset "poisoning" tools (e.g., Nightshade, Kudurru) that block crawlers and/or cause damage to AI models in response to their work being non-consensually included in AI training datasets.[38] [39]

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

Bryan-Kinns and Heravi: Yes.

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

In Bryan-Kinns' view the government should be supporting organisations such as Mechanical-Copyright Protection Society (MCPS) and Performing Right Society Limited (PRS) to collect and distribute royalties for the use of creative content in training commercial AI models e.g. through an opt in rights model managed through such organisations. Bryan-Kinns notes that the Head of Research of the Creative Computing Institute, University of the Arts London, has collaborated with senior policy advisors and called for the formation of a new agency whose remit would include operating in collaboration with existing collections agencies to remunerate rights holders.

Sichani notes that a welcome development in data aggregation services has been introduced in the recently published AI Opportunities Act. The UK government is prioritising the creation of a UK National Data Library (NDL) as a digital infrastructure for collecting and managing

[34] ALCS (2024). "UK AI regulation: principles and the way forward". https://d16dqzv7ay57st.cloudfront.net/uploads/2024/12/ALCS-AI-Position-Paper.pdf

[35] Newton-Rex, Ed (2024). "The insurmountable problems with generative AI opt-outs". https://static1.squarespace.com/static/5cc5785816b6406e50258c5c/t/67368c12cc35b5469feb0bfd/1731628050768/The+insurmountable+problems+with+generative+AI+opt-outs.pdf

[36] Miltner, K. M., & Highfield, T. (2024). "The Possibilities of 'Good' Generative AI in the Cultural and Creative Industries". *The British Academy.* https://www.thebritishacademy.ac.uk/publications/the-possibilities-of-good-generative-ai-in-the-cultural-and-creative-industries/

[37] https://source.plus/

[38] Heikkila, M. (2023). "This new data poisoning tool lets artists fight back against generative AI". *MIT Technology Review.* https://www.technologyreview.com/2023/10/23/1082189/data-poisoning-artists-fight-generative-ai/

[39] Knibbs, K. (n.d.). "A New Tool Helps Artists Thwart AI—With a Middle Finger". *Wired*. https://www.wired.com/story/kudurru-ai-scraping-block-poisoning-spawning/

public sector data as a strategic resource. By establishing a coherent data collection and access policy, along with data library and exchange services and the necessary infrastructure, the NDL aims to responsibly unlock both public and private datasets to drive innovation. It will actively incentivise and reward researchers and industry for curating and making private datasets accessible. Following Wellcome's successful Technical White Paper Challenge[40] last year, there were existing proof-of-concept models that can be further explored and consulted to advance this effort.

Sichani argues that cultural heritage data should be central to this strategy, with copyright-cleared assets from institutions like The National Archives, the Natural History Museum, and the British Library serving as training datasets for international licensing and AI development. Political will must align with sector-specific needs and requirements, and in this vein the National Data Library goals and operation should be further aligned with current AHRC-UKRI investments in digital research infrastructure for the arts and humanities, including the Digital Heritage Science Data Service,[41] the Museum Data Service,[42] and core Digital Research Services under the RICHeS programme, as well as the CoSTAR Network for creative industries.[43]

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

Paula Westenberger's research highlights that the heritage sector (in its full breadth and diversity) and researchers in such contexts should be taken into further consideration in AI and copyright policy, as further developed in her forthcoming paper[44] and Questions 28 and 29 below.

The heritage sector is diverse, formed not only by libraries, archives, galleries and museums, but also natural heritage institutions, heritage sites, historic houses, intangible heritage stakeholders, and others, from different sizes and regions. Different institutions house different materials, and each kind of collection bears special characteristics that can have different and complex copyright implications. It is thus important to engage with the sector in its full diversity to understand the different copyright issues facing these stakeholders.

Bespoke and non-commercial AI models used as tools for research or collections management in heritage contexts present different copyright law implications - particularly regarding "non-commercial research" - compared to commercial generative AI tools. The current policy focus on commercial generative AI presents a missed opportunity to infuse crucial considerations on ethics, bias and cultural impact into AI and copyright regulation, that heritage practitioners are well equipped to reflect on. Heritage stakeholders (including researchers) can thus provide important insights to shape better policies and regulation.

[40] https://wellcome.org/what-we-do/our-work/uk-data-library
[41] https://hsds.ac.uk/
[42] https://museumdata.uk/
[43] https://www.costarnetwork.co.uk/about
[44] Westenberger, P. and Farmaki, D. (2025). "Artificial Intelligence for Cultural Heritage Research: the Challenges in UK Copyright Law and Policy". Available at SSRN: https://ssrn.com/abstract=5153757 or http://dx.doi.org/10.2139/ssrn.5153757.

Understanding the risk-averse nature of the heritage sector is also an important consideration, as this reflects on attitudes towards legal uncertainties, which Westenberger explores further below in Questions 28 and 29 and in her forthcoming paper.[45]

Sichani and Heravi endorse this view.

C.4 Transparency

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

The researchers named in the answer to this question agree that AI developers should disclose the sources of their training material.

Bev Townsend argues that greater transparency is an essential requirement in AI development and AI developers should be compelled to disclose the source of their training data/material. Transparency is instrumental in enabling AI developers and individual's (as data and copyright holders) to navigate the permissions, limitations, and extent of the data/material used. However, there is a need for conceptual clarity, articulating the actual level of transparency and understandability required, and for better understanding how transparency will be operationalised in practice. A recent study [Townsend et al. under review] found that the UK public believe that disclosure and transparency are significant drivers of user trust and that the lack of transparency is a barrier to safe and ethical AI adoption. Transparency should allow for appropriate and reliable means to support traceability and explainability regarding both the lawful and ethical acquisition and collection of the data/material/content and regarding the provenance or source of data/material/content used to test, train, and validate AI models. Without disclosure requirements and transparency obligations and clearly stated inclusion and exclusion criteria in the compilation of the datasets it is difficult for copyright holders to ascertain the origin of the data/material/content and whether and to what extent material has been used or manipulated to generate specific output. Requirements obliging sufficient and appropriate transparency for, specifically systems at high-risk of data and right-holder violations, and unethical use, should compel developers to inform right holders about the system's use of their data and content, and with their compliance with relevant law (such as data protection and copyright law). This should include the requirement to label synthetic and artificially generated or manipulated AI output and any modified right holders' material/content. However, a level of granularity in transparency is required providing rights holders the ability to make informed decisions about their data/material/content and its repurposing, secondary, and extended use.

In Collett's view, while AI developers might not be keen to disclose the sources of their training material in a granular fashion -- as they might feel it endangers their IP -- detailed, high-level lists of training data must be called for and prioritised to ensure sufficient transparency and rightful control over creative works. This should be required by transparency regulations in relation to copyright law, which should be updated to account for generative AI developments. There should also be the ability to see these sources upon request and to request freedom of information, similar to the GDPR.

[45] Ibid.

Bryan-Kinns seconds the views of Townsend and Collett. Bryan-Kinns notes that in addition, consideration needs to be given to the use of foundational AI models trained before the revision of copyright laws. Specifically, models trained before the revision of copyright law may have been trained on copyright material which the rights holder did not wish to be used in AI training which after the revision of copyright laws are explicitly not to be used for AI model training. In such cases it needs to be clear whether legacy AI models trained with potentially copyright materials prior to the copyright law revisions can be used legally without infringing copyright laws.

Sichani suggests that AI developers must disclose the sources of their training material through transparent documentation, enforced by law and supported by standardised templates and aggregation mechanisms. Transparent documentation is vital for responsible AI, ensuring compliance, accountability, and ethical data use. Beyond legal enforcement, structured documentation enhances data reuse and streamlines AI pipelines, especially for foundational models. Since most AI development relies on fine-tuning pre-trained models, improving transparency around fine-tuning data is as crucial as for training data. Transparency around training data is essential also for rights holders to determine if their data has been used, opt out, or seek remuneration, while also supporting research verification. Finally, lawmakers and regulators require access to training data to ensure compliance, and greater transparency further fosters users' trust in AI systems.

Heravi and Miltner share the views of Sichani, Bryan-Kinns, Collett and Townsend.

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

In Bryan-Kinns' view the level of granularity needed is at the level of rights given. For example, if a whole collection is licensed for AI training then this collection should be reported as used in training whereas if an individual creative output is used in the AI training then this would need to be reported as used in the training. Ideally the list should be machine readable and searchable and human readable.

Sichani underlines that since most AI development relies on fine-tuning pre-trained models, transparency around fine-tuning data is as critical as for training data. However, this aspect often remains opaque. Initiatives like the Data Provenance Explorer[46] highlight the need for transparency also in fine-tuning data level. Ensuring consistent and standardized disclosure of training, testing, and fine-tuning data enables regulators and stakeholders to easily interpret and compare model development processes.

Heravi and Miltner second Bryan-Kinns and Sichani's view.

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

In Townsend's view, an amendment to the Data (Use and Access) Bill, introduced in Parliament on 24 October 2024, places new responsibilities on the ICO to regulate the transparency of web crawler use. Aligned with this, and minimally, operators of web crawlers

[46] https://www.dataprovenance.org/

should observe UK copyright law and should be transparent about the acquisition, origin, and purpose of the data/content/material.

Bryan-Kinns and Collett second Townsend's view.

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

In Bryan-Kinns' view a proportionate approach to transparency is to provide a full list of all collections and individual creative outputs used in training AI models. This is not onerous on the company training the AI model given that they programmatically collect all this information when training the AI model and will need to ascertain that they have the rights to use the content for training - currently this information is typically discarded in the training process. This list should be machine readable and searchable and human readable.

Collett, Townsend and Miltner second this view.

Heravi seconds this view and adds that the disclosure can get blurred when the content is used from secondary sources, for example images from social media platforms or news from news aggregator or social media platforms. There is a need to ensure such platforms can carry the preferences of the original content creator, or provide clear technical and non-technical features for clarifying such rights and preferences. For example, disclosing that 'Instagram' was used as a source or news items used for training from Facebook or Google news aggregators is not acceptable. Provenance is very important.

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

In Bryan-Kinns' view, compared to the enormous financial and energy cost of training large AI models, the cost of retaining and making available a sizeable list of sources of training data will be negligible.

Collett, Heravi and Miltner second this view.

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

In Townsend's view, while regulatory underpinning provides an imperative objective to good transparency practice, practical tools, templates, and codes of practice should be developed to provide guidance and direction. In addition, clear paths for redress and sanction should be provided by regulators.

In Bryan-Kinns' view this requires regulatory underpinning. Compliance can be ensured through auditing and support for litigation against companies suspected of copyright infringements around AI model training.

In Sichani's view, clear documentation of training sources must be enforced by law and supported by standardised templates and aggregation mechanisms. Several existing initiatives provide guidance, best practices and some already-made and tested templates for

responsible AI documentation that can be used as a starting point, including standardised data cards[47] and datasheets[48] and AI model cards[49]. For datasets from cultural heritage institutions, specialised documentation approaches have been proposed, such as Datasheets for Digital Cultural Heritage Datasets[50], that include specific fields to GLAM collections. Also the concepts of Data envelopes[51] designed to capture ethical and contextual considerations specific to cultural heritage datasets.

Furthermore, initiatives like the AI Security Institute's[52] efforts to design mandatory reporting requirements and standards are a welcome step. Clear documentation of training sources could be reinforced through a registry or certification system for AI companies adhering to best practices. A certification model, such as a 'Fairly Trained' license[53], could help establish industry standards.

Heravi, Collett and Miltner second the views of Townsend, Bryan-Kinns, and Sichani.

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

In Bryan-Kinns' view the EU's approach to transparency is welcomed in principle but lacks detail on what "sufficient detail" is which makes the approach hard to evaluate and implement. A better approach would be to require that all AI training sources are reported to remove ambiguity. Collett, Townsend, Sichani, Heravi and Miltner second this view.

C.6 Encouraging research and innovation

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

Westenberger's view is that the current TDM exception in s 29A is not fit for purpose for non-commercial research, particularly in heritage contexts, and that the exception should be clarified and expanded. The response below summarises some of the arguments made in her

[47] Pushkarna, Mahima et al. (2022). "Data Cards: Purposeful and Transparent Dataset Documentation for Responsible AI." *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency* https://doi.org/10.1145/3531146.3533231

[48] Gebru, T., Morgenstern, J., Vecchione, B., Vaughan, J.W., Wallach, H.M., Daumé, H., & Crawford, K. (2018). "Datasheets for datasets". *Communications of the ACM, 64*, 86 - 92. https://arxiv.org/abs/1803.09010

[49] Mitchell, M., Wu, S., Zaldivar, A., Barnes, P., Vasserman, L., Hutchinson, B., Spitzer, E., Raji, I.D., & Gebru, T. (2018). "Model Cards for Model Reporting". *Proceedings of the Conference on Fairness, Accountability, and Transparency*. https://arxiv.org/abs/1810.03993

[50] Alkemade, H., Claeyssens, S., Colavizza, G., Freire, N., Lehmann, J., Neudecker, C., Osti, G. and van Strien, D. (2023). "Datasheets for Digital Cultural Heritage Datasets", *Journal of Open Humanities Data*, 9(1), p. 17. https://doi.org/10.5334/johd.124.

[51] Mrinalini Luthra and Maria Eskevich (2024). "Data-Envelopes for Cultural Heritage: Going beyond Datasheets". In *Proceedings of the Workshop on Legal and Ethical Issues in Human Language Technologies @ LREC-COLING 2024*, p.52–65, Torino, Italia. ELRA and ICCL. https://aclanthology.org/2024.legal-1.9/

[52] https://www.aisi.gov.uk

[53] https://www.fairlytrained.org

forthcoming paper 'Artificial Intelligence for Cultural Heritage Research: the challenges in UK copyright law and policy'.[54]

As noted in the *Living With Machines* project book, the exception has proved difficult to use in innovative research involving diverse datasets, limiting its effectiveness in supporting national priority research in the intersection of technology and culture.[55]

The main issues, as argued in the above-mentioned Westenberger's forthcoming paper, are:

- Definition of lawful access is unclear;
- Lawful access requirements may restrict the enjoyment of the exception and may represent contractual override in practice (which is not permitted by the exception);
- Prohibition to transfer copies in CDPA, s 29(A)(2)(a) is unclear and unsuitable for current collaborative forms of AI research: organisations holding relevant data often partner with those possessing computational resources/expertise needed for AI projects. The TDM research exception should support this practice by allowing transfer of datasets between research partners in such cases;
- Web mining: it is unclear the extent to which content placed online falls under lawful access;
- Risk-averse attitudes and lack of resources in the sector: this can essentially deter organisations from engaging in important AI research, especially in view of the lack of clarity of the exception;
- Issues of bias: relying on copyright permissions, or the copyright status of materials, to select research datasets may compromise AI research quality;
- "Non-commercial" research: boundaries between commercial and non-commercial can become blurred (including in the heritage sector, where public-private partnerships can occur in digitisation projects). Focusing on non-profit organisations/uses may be an alternative solution.

The issues above need to be clarified and resolved before any new commercial TDM exception with opt out as proposed in section C.1 is introduced. As argued in Westenberger's forthcoming paper, opt outs "raises concerns about potential biases, omissions, and incomplete datasets that could skew and compromise AI research. This is particularly problematic if this provision was to apply in research and heritage contexts, which could become the case in light of the uncertainties regarding the current non-commercial research TDM exception and risk-averse attitudes of heritage stakeholders."[56]

Sichani seconds Westenberger's view.

---

[54] Westenberger, P. and Farmaki, D. (2025). "Artificial Intelligence for Cultural Heritage Research: the Challenges in UK Copyright Law and Policy". Available at SSRN: https://ssrn.com/abstract=5153757 or http://dx.doi.org/10.2139/ssrn.5153757 .

[55] Ahnert R, Griffin E, Ridge M, and Tolfo G. (2023). *Collaborative Historical Research in the age of big data.* Cambridge University Press p. 28. Available at: https://www.cambridge.org/core/elements/collaborative-historical-research-in-the-age-of-big-data/839C422CCAA6C1699DE8D353B3A1960D

[56] Westenberger, P. and Farmaki, D. (2025). "Artificial Intelligence for Cultural Heritage Research: the Challenges in UK Copyright Law and Policy". Available at SSRN: https://ssrn.com/abstract=5153757 or http://dx.doi.org/10.2139/ssrn.5153757

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?

Paula Westenberger believes that yes, the purpose of an AI model should be considered. It is important to note, however, that looking at purpose alone might not provide enough clarity depending on how the purpose is framed: see the discussion above and in her paper on the blurred boundaries between commercial versus non-commercial purposes.[57] It might be important to also look at the beneficiary of the provision. In Westenberger's view, heritage and research uses that are in the public interest, and public heritage and research institutions, deserve particular consideration.

Other exceptions are specifically tailored to the needs and characteristics of the heritage and research sectors, so it is not something unfamiliar to copyright law to look at specific purposes or beneficiaries, particularly when discussing exceptions. What seems important is that those rules are clearly crafted, and in close consultation with the relevant stakeholders, considering the practical issues facing the relevant sector. Suggesting "a more rigorous and empirically focused framework through which to assess … the drafting of copyright exceptions", and focusing specifically on cultural institutions.[58]

Westenberger sees as problematic that this question is framed as enquiring on the size of the "AI firm". In a discussion on s 29A, aimed at TDM for non-commercial research, the question should, importantly, enquire on the size of the research or heritage organisation. Consultation documents tend to overemphasise AI firms and creative industries, failing to address important public sector organisations and stakeholders such as those in heritage.

More proactive and substantial engagement with the heritage and research sectors is required in copyright policy making. Public consultations such as this, albeit a great opportunity to contribute to discussions, may not necessarily be accessible/feasible to many such stakeholders. Consultations so far are being drafted in a manner that in my view does not appropriately invite heritage sector contributions (for example the "AI firms" emphasis in this question), and provide too short of a response timeframe (including over the festive period). As such, it does not allow sufficient time for stakeholder engagement events and research to be carried out that could reveal important evidence to the questions posed. Therefore, Westenberger believes that post-consultation engagement strategies, including targeted events such as roundtables with appropriate and diverse representation, would be important to try and capture the views and evidence of the heritage sector in its full breadth and diversity.

See also Westenberger's response to Question 16 above. Sichani endorses Westenberger's view.

[57] Ibid.
[58] Hudson E. (2020). *Drafting Copyright Exceptions: From the Law in Books to the Law in Action*, Cambridge University Press, p.13.

D.2 AI outputs Policy options

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.

Bryan-Kinns and Heravi: No.

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

Bryan-Kinns: Yes. In Bryan-Kinns' view there is no situation in which an artistic work is generated by a computer where there is no human author. In all cases, there is a "person by whom the arrangements necessary for the creation of the work are undertaken" (s 9(3) CDPA) - whether it refers to those who are writing a user prompt or training the AI model or contributing to writing the code of the AI model or simply starting the AI model and providing it with electricity to allow it to continue to function.

Question 36. What would be the economic impact of doing this? Please provide quantitative information where possible.

Bryan-Kinns: None. In Bryan-Kinns' view the removal of the copyright protection for CGW would have no effect as they do not and cannot exist.

Question 37. Would the removal of the current CGW provision affect you or your organisation? Please provide quantitative information where possible.

Bryan-Kinns: No.

D.4 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?

In Bryan-Kinns' view it mostly does, but there needs to be clarification of whether users of AI models intentionally infringe copyright when using AI models. Specifically, the consultation notes that "Depending on the circumstances, both the user and provider of a model may be liable for infringement of copyright". This needs to be clarified for the situation where a user unintentionally and unknowingly generated copyright infringing content using an AI model. In such situations there needs to be clarity that the AI service would be solely liable for copyright infringement and not the user. Heravi shares this view.

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

In Bryan-Kinns' view, AI providers should not train their models on copyright content that they do not have rights to use in the training of their models. Heravi seconds this view.

D.5 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?

In Caterina Moruzzi's view, labelling AI outputs as AI generated can be valuable to reduce the spread of misinformation[59] and to track the creation, sharing, and publication of creative work. The potential value of provenance data stretches further than their contribution to reducing disinformation and preventing audiences from being misled. It can be integral to ensuring proper attribution, understanding creative intent, and supporting the proper management of rights of creative workers.[60] However, exposing Generative AI use to users in a binary way ("AI", "Not AI") may not be sufficient. Different user groups involved in the creation and consumption of content may have different tolerances and needs in regards to the level of granularity they can access in respect to the provenance and origin of content.[61] In addition, a simple binary labelling system may have unintended consequences, as users may overgeneralise the meaning of provenance indicators, and raising awareness of misinformation may increase scepticism of factual information.[62] A recent study [Moruzzi *et al.*, under review] has highlighted the relevance of considering contextual and situated information when designing provenance signals. In particular, users might engage with content differently and have different demands according to the context of use. This all calls for acknowledging the multifaceted nature of provenance and provenance needs, which should be reflected in labelling systems that allow for tailoring information to various user groups.

In addition, it is necessary to consider how rich provenance data and labels may challenge traditional notions of ownership, creativity, and recognition. From a creator's perspective, the integration of rich provenance labels into their work can introduce unwanted elements of surveillance and also reduce some of the aesthetic and artistic characteristics that are essential to the work.[63] Careful consideration of the granularity of the information that is provided is crucial both for increased transparency for users, and for increased creative agency and control for creators.

Bryan-Kinns, Heravi and Miltner second Moruzzi's view.

---

[59] Francis II, Errol et al. (2024). "Usable News Authentication: How the Presentation and Location of Cryptographic Information Impacts the Usability of Provenance Information and Perceptions of News Articles". In *Proceedings of the CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI '24). Association for Computing Machinery, New York, NY, USA, Article 830, 20 pages. https://doi.org/10.1145/3613904.3642331

[60] Liddell, Ella et al. (2024). "ORAgen: Exploring the Design of Attribution through Media Tokenisation". In Companion Publication of the 2024 ACM Designing Interactive Systems Conference (IT University of Copenhagen, Denmark) (DIS '24 Companion). Association for Computing Machinery, New York, NY, USA, 229–233. https://doi.org/10.1145/3656156.3663693

[61] Wittenberg, Chloe et al. (2024). "Labeling AI-Generated Content: Promises, Perils, and Future Directions. An MIT Exploration of Generative AI". https://doi.org/10.21428/e4baedd9.0319e3a6

[62] Feng, Kevin et al. (2023). "Examining the Impact of Provenance-Enabled Media on Trust and Accuracy Perceptions". Proc.ACMHum.-Comput.Interact.7,CSCW2,Article270(oct2023). https://doi.org/10.1145/3610061

[63] Kronfeldner, M. (2009). "Creativity Naturalized". *The Philosophical Quarterly* 59, 237 (07 2009), 577–592. https://doi.org/10.1111/j.1467-9213.2009.637.x

D.6 Digital replicas and other issues

Question 44. Could you share your experience or evidence of AI and digital replicas to date?

Bahareh Heravi argues that digital replicas are being explored in news organisations for use in translation, simultaneous interpretation add-ons, and similar applications. For example, a Spanish politician gives a speech in Spanish, and a news organisation offers an option to watch or listen to their speech in English. AI tools already trained on the politician's voice and image, as well as potentially on their English-speaking voice from past instances where they have spoken in English, can generate an English version of their speech, in which the politician could appear to be speaking English—with their own Spanish accent and accurate lip-syncing. This can be highly realistic, making it hard for viewers to distinguish whether the politician originally spoke in English. While not intended as a deepfake, it still creates a digital replica. Moreover, subtle translation errors could alter meaning, yet the audience might assume the politician said those exact words. In this case, "supposed evidence" is no longer proof, and this can cause severe consequences. Clear labelling or requiring explicit user action (e.g., clicking a button for the English version) could help. However, if recorded and reshared, provenance may be lost, making it unclear whether the footage is AI-generated. Heravi highlights that this is already implemented at experiment level at some news organisations.

About BRAID

BRAID is a UK-wide programme dedicated to integrating Arts and Humanities research more fully into the Responsible AI ecosystem, as well as bridging the divides between academic, industry, policy and regulatory work on responsible AI.

It represents a six-year, £15.9 million investment in enabling responsible AI in the UK, funded by the Arts and Humanities Research Council (AHRC) from 2022 to 2028 (grant number AH/X007146/1).

Working in partnership with the Ada Lovelace Institute and BBC, its team brings together expertise in human-computer interaction, moral philosophy, arts, design, law, social sciences, journalism, and AI. BRAID is extended by a network of interdisciplinary researchers and partnering organisations through the delivery of funding calls, community building events, and a series of programmed activities.