

**Wikimedia Foundation’s Responses
to the
United States Copyright Office
Request for Comments on Artificial Intelligence and Copyright
Docket No. 2023-6
30 October, 2023**

Introduction

The Wikimedia Foundation appreciates the opportunity to comment on the US Copyright Office’s Request for Comments on Artificial Intelligence and Copyright. The Wikimedia Foundation is the nonprofit organization that hosts and supports a number of free, online, collectively-produced resources, including Wikipedia. The Foundation’s objective is to help create a world where every human being can share freely in the sum of all knowledge.

Copyright has always been particularly important to that mission because copyright rules set the limits on what constitutes free knowledge. The Foundation has used free culture and free and open source software copyright licenses since its inception and is currently the world’s largest repository of Creative Commons licensed material. This also places the Wikimedia Foundation and the projects it hosts in an important role with relation to generative AI: freely licensed technology, including various machine learning (ML) tools, help to support the quality of the Wikimedia projects and aid volunteer editors in working more effectively and efficiently.¹ For example, virtually all of the content on Wikipedia is created, edited, and verified by volunteers. Volunteers, in conjunction with Foundation staff, also build and maintain automated tools to help them maintain the projects’ high quality, such as tools to detect vandalism or assess the relative “health” of articles so that volunteer editors can prioritize their efforts. At the same time, the freely licensed high quality information and media on Wikipedia and the other Wikimedia projects forms one of the most important bases for training generative AI programs.²

Because of this importance, we wish to submit the following responses to the Office’s questions on this topic.

¹ See, e.g., Jonathan T. Morgan, *Designing ethically with AI: How Wikimedia can harness machine learning in a responsible and human-centered way*, Wikimedia Foundation, July 2019, <https://wikimediafoundation.org/news/2019/07/18/designing-ethically-with-ai-how-wikimedia-can-harness-machine-learning-in-a-responsible-and-human-centered-way/>; Miriam Redi et al., *Can machine learning uncover Wikipedia’s missing ‘citation needed’ tags?*, Wikimedia Foundation, April 2019, <https://wikimediafoundation.org/news/2019/04/03/can-machine-learning-uncover-wikipedias-missing-citation-needed-tags/>; Miriam Redi, *How we’re using machine learning to visually enrich Wikidata*, Wikimedia Foundation, March 2018, <https://wikimediafoundation.org/news/2018/03/14/machine-learning-visually-enriching-wikidata/>.

² See Selena Deckelmann, *Wikipedia’s value in the age of generative AI*, Wikimedia Foundation, 12 July 2023, <https://wikimediafoundation.org/news/2023/07/12/wikipedias-value-in-the-age-of-generative-ai/>.

General Questions:

1. As described above, generative AI systems have the ability to produce material that would be copyrightable if it were created by a human author. What are your views on the potential benefits and risks of this technology? How is the use of this technology currently affecting or likely to affect creators, copyright owners, technology developers, researchers, and the public?

Overall, the Foundation believes that generative AI tools offer benefits to help humans work more efficiently, but that there are risks of harms from abuse of these tools, particularly to generate large quantities of low-quality material.

Generative AI systems have only recently been made available for public use.³ The Wikimedia Foundation, like many others, has watched with interest as these tools have been deployed at scale, sparking public debate about a wide range of concerns and policy considerations. At the same time, the developers of these tools continue to refine and advance both the underlying technologies and their practical applications. Although the emergence of text and image generators has proven massively popular, it is not yet clear what applications or use cases may prevail in the long term. This makes it difficult to predict what sets of policy issues to address with government action.

As for the risks arising from the currently available generative AI systems, we see the large-scale generation of unverifiable, inaccurate information as a threat to the global information ecosystem. Projects like Wikipedia depend on the availability of accurate, reliable sources of information, including both verifiable written text and multimedia content that accurately depict noteworthy topics. These sources may become increasingly difficult to find and verify as accurate amid a sea of synthetic content. In this context, it is critical that humans such as the volunteer editors who contribute to Wikipedia remain engaged in the creation and curation of new content. We urge the Office as well as policymakers outside the realm of copyright to maintain a focus on the humans who are essential to and impacted by the development of technologies like generative AI.

We are also concerned that an overreliance on generative tools to create text, images, music, and audiovisual works could disrupt the “skills pipeline.” For example, this pipeline can begin with a desire to learn how to write, draw, or play an instrument and, in time, produce a population of experts. If enough people turn to generative AI systems instead of acquiring and developing skills in the arts, there may one day be a significant lack of experts to pass along their knowledge and mentor the next generation of artists.

³ See, e.g., DALL-E, <https://en.wikipedia.org/wiki/DALL-E>; see also, ChatGPT, <https://en.wikipedia.org/wiki/ChatGPT>.

As for the potential benefits of generative tools, we see promise in applications that support humans' ability to summarize and translate text. The Wikimedia Foundation has itself been using machine learning to aid humans with content translation for many years, well before the current release of popular generative AI tools.⁴ We believe that combining machine-generated automatic translation with a human review stage allows for more effective production of knowledge across languages and for different audiences.⁵ By combining machine learning and human review, this approach also helps to prevent mistakes and avoid some of the potential harms—e.g., misinformation.

2. Does the increasing use or distribution of AI-generated material raise any unique issues for your sector or industry as compared to other copyright stakeholders?

The Foundation is somewhat uniquely positioned as both the host of a primary source of training material for generative AI and a user of many AI and ML tools that aid human editors with the creation of free knowledge. As such, we support copyright policies and technical tools that help our human editors work effectively. We also want to ensure that those same editors receive credit for their work and that their contributions to free knowledge projects like Wikipedia are not inhibited by a world filled with generative AI tools.

The Wikimedia Foundation provides legal and technical support for Wikipedia and other projects. These projects are widely used as training material for generative models.⁶ The Foundation's mission is to empower people to collect and develop educational content, but also to disseminate it effectively.

To achieve these goals, we want to ensure that Wikipedia and other projects remain high-quality and are supported by robust and accurate source materials. As mentioned in response to Question 1, we are concerned about the impacts generative AI outputs could have on this source material, in terms of both accuracy and verifiability. To the extent that generative AI systems are used to create large volumes of lower-quality content, this could impact volunteers' ability to find and trust materials to support their additions to the projects.

However, we also want volunteers to have access to tools powered by AI or ML, including those involving language models, which can help them manage content on the projects by enabling automated vandalism detection or assisting in translation efforts. Neither is it clear at this stage of development whether generative AI systems will be well suited to meet the needs of Wikipedia and other Wikimedia project volunteer editors for these purposes or not, nor is it clear how such systems will affect the state of the art in machine learning as a whole.

⁴ See, e.g., Extension:Translate, <https://www.mediawiki.org/wiki/Extension:Translate>, and Content translation, https://www.mediawiki.org/wiki/Content_translation.

⁵ Chris Stokel-Walker, *AI tidies up Wikipedia's references, and boosts reliability*, Nature, 19 October, 2023, <https://www.nature.com/articles/d41586-023-02894-x>.

⁶ See Jon Gartner, *Wikipedia's Moment of Truth*, New York Times, 18 July, 2023, <https://www.nytimes.com/2023/07/18/magazine/wikipedia-ai-chatgpt.html>

Finally, although we do not view currently available generative AI systems as sufficiently accurate, reliable, or verifiable to be used as tools to directly query or access knowledge, AI-powered search and other applications may improve over time to meet these and other criteria to become trusted sources of knowledge. We are currently working with companies on projects to make improvements in these areas, primarily by allowing existing generative AI tools to draw on external sources of information like Wikipedia, rather than only generate text based on the trained model itself.⁷ To the extent that generative AI systems may one day become important interfaces through which people access knowledge, we want to support their development and ensure that human-created resources like Wikipedia are integral to their accuracy and verifiability.

3. Please identify any papers or studies that you believe are relevant to this Notice. These may address, for example, the economic effects of generative AI on the creative industries or how different licensing regimes do or could operate to remunerate copyright owners and/or creators for the use of their works in training AI models. The Office requests that commenters provide a hyperlink to the identified papers.

We recommend a few items for review by the Office:

- [Open Future Foundation texts on AI and the Commons](#). The Open Future Foundation has done various researches on the potential impacts and harms of AI tools, particularly with regard to open source and open data topics.
- ["Talkin' 'Bout AI generation" by Grimmelmann et al.](#) This paper is an academic review of both the technology and potential copyright implications of generative AI. It aims to present the issues objectively and covers a range of possible copyright interpretations and their potential impacts. We believe the Office would find this valuable when considering any recommendations related to the technology.
- [Federal Trade Commission \(FTC\) roundtable from 4 October, 2023](#) As the FTC is also considering this topic, we encourage the Office to review their discussion and consider already available material to avoid any potential regulatory or legislative conflicts.

4. Are there any statutory or regulatory approaches that have been adopted or are under consideration in other countries that relate to copyright and AI that should be considered or avoided in the United States? How important a factor is international consistency in this area across borders?

The Wikimedia Foundation supports international consistency for regulatory approaches related to AI and copyright. Volunteers on Wikipedia and other Wikimedia projects usually collaborate based on common languages, which is why the projects are typically categorized by language version rather than national borders. As a result, conflicting regulations addressing the use of training datasets like Wikipedia, the use of generative AI tools by editors, and the use of

⁷ Maryana Pinchuk, *Exploring new paths for the future of free knowledge: New Wikipedia ChatGPT plugin, leveraging rich social media apps, and other experiments*, 13 July, 2023, <https://diff.wikimedia.org/2023/07/13/exploring-paths-for-the-future-of-free-knowledge-new-wikipedia-chat-gpt-plugin-leveraging-rich-media-social-apps-and-other-experiments/>.

AI-generated outputs could create negative impacts for both volunteers to the Wikimedia projects (who may find their work used inconsistently across borders) and for the Foundation as a developer of AI and ML tools (which would be more difficult to do if the potential use of data depended on the geography from which it was contributed). On the use of data, specifically, we encourage regulators and legislators to align their approaches with existing models, such as the European Union's inclusion of an exemption for text and data mining in the Directive on Copyright in the Digital Single Market, that enable public interest research and other beneficial uses of protected works.⁸

For example, conflicting or inconsistent regulations on the use of generative AI tools or their generated outputs could complicate the deployment of such tools, especially those used by the Foundation, across the more than 300 language versions of Wikipedia that span international borders. The Foundation strives to promote equity among volunteers around the world, but diverse approaches to regulation could make this goal more difficult to achieve if volunteers in one country are prohibited from using tools that empower their counterparts in another country.

In addition to international consistency, we encourage the Office to consider the potential impacts that changes to copyright law could have on competition among AI developers. If copyright law changes are enacted such that the acquisition and use of training materials becomes more expensive or difficult, there is a risk that dominant firms with greater resources will become further entrenched while smaller companies, including nonprofit organizations, struggle to keep up with mounting development costs. We do not suggest that regulators face a binary choice between allowing only a small number of highly regulated firms to develop AI applications and an open field of unchecked harms to some aspects of society. Rather, to the extent that the Office does identify harms that require changes in law or regulations, we encourage looking for approaches that address the harms without increasing costs to the point of causing industry monopoly or oligopoly.

Training

6. What kinds of copyright-protected training materials are used to train AI models, and how are those materials collected and curated?

While we do not have good information about the full range of proprietary training methods and datasets, we are aware from both formal review of some AI datasets and informal conversations with experts that Wikipedia forms an important part of generative AI training data.

Wikipedia content is regularly used to train AI models for several reasons: it is free to use without prior authorization for any purpose (typically requiring only attribution and a share-alike provision for derivative works under its Creative Commons license); it is a large corpus of reliable information; and, it offers more internal structure than many other sources of training

⁸ Directive (EU) 2019/790 of the European Parliament and of the Council of 17 April 2019 on copyright and related rights in the Digital Single Market and amending Directives 96/9/EC and 2001/29/EC, L 130/92, <https://eur-lex.europa.eu/eli/dir/2019/790/oj>.

material.⁹ Although the Wikimedia Foundation does not track who accesses, scrapes, or downloads content from Wikipedia, we are aware of some companies' uses of Wikipedia content for modeling purposes due to existing partnerships established for this and other purposes.¹⁰

6.1. How or where do developers of AI models acquire the materials or datasets that their models are trained on? To what extent is training material first collected by third-party entities (such as academic researchers or private companies)?

We note that because of Wikipedia's broad accessibility and inclusion in a variety of datasets, our working assumption is that Wikipedia content appears in both academic datasets as well as training materials compiled by developers of AI models.

6.2. To what extent are copyrighted works licensed from copyright owners for use as training materials? To your knowledge, what licensing models are currently being offered and used?

Although AI model developers offer varying degrees of transparency as to the sources of materials used for training, which may include private collections, we would categorize the corpus of publicly available works as follows:

Public, but not licensed. This includes materials collected from the open internet without prior authorization and whose use is not subject to prior authorization. In other contexts, courts have determined such uses of copyrighted works to be non-infringing fair uses,¹¹ although the full context of generative AI has not yet been ruled on. Given the large volume of content available on publicly accessible websites, including previously compiled and configured datasets like those available from the Large-scale Artificial Intelligence Open Network (LAION),¹² and the relatively low upfront costs of acquisition, unlicensed public content may represent a significant portion of the training materials used to develop and refine models.

Public and licensed. Compilations of copyrighted material like Wikipedia are licensed under some form of license, usable by any member of the public, which explicitly or implicitly permit use of works for purposes such as model training. Creative Commons licenses like CC0 are unconditional and clearly permit any form of use, including for commercial purposes.¹³ Use of

⁹ Use of Wikipedia content is primarily governed by the Creative Commons BY-SA 4.0 license. Some individual material may be public domain or accessed and shared under other licenses.

¹⁰ See https://en.wikipedia.org/wiki/Wikipedia:Database_download for more information on how to download Wikipedia content. More information about commercial partnerships is available at <https://enterprise.wikimedia.com/>.

¹¹ See, e.g., *Authors Guild, Inc. v. Google, Inc.*, 721 F.3d 132 (2nd Cir. 2015), *Authors Guild, Inc. v. HathiTrust*, 755 F.3d 87 (2nd Cir. 2014).

¹² For more on LAION, see <https://laion.ai/>.

¹³ CC0, "No Rights Reserved," <https://creativecommons.org/public-domain/cc0/>.

Wikipedia content is subject to the Creative Commons BY-SA license, which requires attribution and further sharing according to the same licensing terms.¹⁴

Given that Wikipedia content is both freely available and relatively valuable from a training perspective, yet many large language models fail to disclose the sources of their training data, we infer that some model developers may be out of compliance with the attribution clause of the CC BY-SA license. We also recognize that the need for compliance with licensing agreements depends on whether courts rule that copying such data for training is or is not a fair use. Likewise, if models trained on Wikipedia content are considered to be derivative works of their training data and they are not made available under a CC BY-SA license, the owners of those models may also be out of compliance with the “share-alike” clause. The Foundation generally supports uses of Wikipedia content for purposes including AI model development. In line with the free and open culture values of our movement, we encourage reusers of Wikipedia content to consider ways to provide attribution and to release models they develop under free and open licenses, even if court rulings ultimately determine that this is not legally required.

Works licensed explicitly for AI model training. Some developers and researchers have compiled datasets made of works for which they either own or have obtained the rights to use as training material. As we understand them, applications such as Adobe’s Firefly are built on models trained from exclusively licensed materials and materials in the public domain.¹⁵ This is a relatively new trend in AI model development and the size and contours of the market for these datasets are not well established. It is very likely that copyright case rulings or regulation could substantially influence this potential market as well.

8. Under what circumstances would the unauthorized use of copyrighted works to train AI models constitute fair use? Please discuss any case law you believe relevant to this question.

While the Wikimedia projects primarily host freely licensed or public domain content, some fair use material appears on the projects. The editors of Wikipedia decide on these uses on a per-language basis, usually for situations where no freely licensed image is available, such as for illustrating individuals who have died, groups that cannot be reassembled, or recent notable works where the copyrighted work is the topic of educational discussion. Based on this experience, we believe that Wikipedia provides a strong understanding of fair use in practice and we offer some thoughts on how fair use may apply in the context of generative AI.

We believe the fair use question should be answered holistically, considering both how AI tools are *trained* and *used*. Fair use analysis always depends on the specific facts and context of each situation.¹⁶ Thus, training a model for research purposes might be different than training a model to launch a commercial service. Similarly, an entity may not be able to simply license a research model for the purpose of building a commercial product, since the researchers may

¹⁴ CC BY-SA 4.0 Deed, <https://creativecommons.org/licenses/by-sa/4.0/>.

¹⁵ For more on Adobe Firefly, see <https://www.adobe.com/sensei/generative-ai/firefly.html>

¹⁶ See, e.g., *Andy Warhol Foundation for the Visual Arts, Inc. v. Goldsmith*, 598 US (2023).

lack the rights to offer such a license even if their own use is fair. Reusers of content on Wikipedia may be familiar with this already. For example, while a Wikipedia article about a popular movie may display an image of the movie poster to illustrate the free, educational article that references and comments on that poster, that does not give for-profit publications the right to freely copy the movie poster.

Key questions that should be considered in a fair use analysis of training include:

Considering factor one in determining fair use as established in [Section 107 of the US Copyright Act](#), what purpose should courts consider when assessing whether a use was transformative? If courts consider the most immediate purpose—typically, copying a work for the purpose of extracting data about that work to inform the mathematical weights or parameters of a particular model—they might arrive at a very different conclusion than if they consider the purposes for which that model might be used. Courts adopting the former approach may be more likely to find that the first factor weighs in favor of fair use, since developing a statistical model is almost certainly a different purpose than the original work was made to serve.

Courts adopting the latter approach are likely to be less consistent in their findings because of the wide variety of purposes a model might serve, but these findings may be more likely to address social and commercial harms as a result. In theory, a model developed using the works of one artist exclusively could be used in a system that mimics that artist's style to generate new works. In contrast, the same model could be used in a system that identifies possible forgeries by predicting whether a given work was produced by the artist whose works were used to train the model. Additionally, the way courts approach the interplay between the first and fourth factors could impact the consistency of fair use jurisprudence.

In factor two, how much will courts care about whether the underlying works used in training are more factual or more expressive? Courts assessing the first factor based on the purpose of a *system* in which a model is used may see the second factor as having some importance. The importance of factor two may be further enhanced if courts find that the use of a work primarily produces factual content or primarily produces more creative content. For example, using a work to develop a model that helps to generate lookalikes for an expressive work feels very different than using a work to develop a model that generates factual summaries of news reports.

In factor three, will courts adapt their analysis to match typical industry practices? The current understanding of the industry is that the purpose of training is best served by always copying an entire work. But conversely, the vastness of the datasets used in training mean that any single copy is barely a drop in the ocean of the whole. If factor three is to remain relevant, courts may need to assume that the entire work was copied and focus, instead, on the extent to which a work is weighted in the development of a model.¹⁷

¹⁷ We note that our understanding is that the topic of measuring the weight of a given work in training a generative AI model is an area of active technical research.

We understand that some methods for developing or refining models include enhancing the value of some parts of the training set relative to others. This can be done in a variety of ways and results in those works having a greater impact on the way a model performs, compared to other works in the training materials. Hypothetically, if a copyright protected work was manually weighted to have an outsized impact in model development, then one could argue that although the uses of other full works may be fair, the amplification of one particular work in the training set is not.

In factor four, how will courts define market substitution? A model that can generate all sorts of images or written works and that is employed in lieu of hiring any human artist would likely be too generic to actually substitute for any given work. On the other hand, some models can be more specifically trained to substitute for existing works, or can be trained poorly in ways that consistently produce works that would directly infringe the training data. It may be worth considering, in this respect, the level of sophistication and effort required to cause a given implementation of a model to produce works that are market substitutes for the training data. It is also worth considering efforts made outside of the training process itself, such as safeguards implemented in the interface to access the model.

Based on our analysis, we do not believe that training AI models should either be categorically fair use or categorically not fair use. Rather, the particulars of the training process and the way courts view the purposes of a use should inform whether a particular training process is fair or not.

Transparency & Recordkeeping

15. In order to allow copyright owners to determine whether their works have been used, should developers of AI models be required to collect, retain, and disclose records regarding the materials used to train their models? Should creators of training datasets have a similar obligation?

The Wikimedia Foundation supports transparency where possible as a part of our broader commitment to the free culture and open source movements. In the context of AI, we believe that transparency would provide value to people who want to study and understand these works for various purposes. We note, however, that the purposes of transparency extend beyond the interests of copyright law. We encourage the Office to consider whether a more holistic, non-copyright approach to transparency disclosures would protect interests within and outside of copyright. Likewise, we note that many purposes of transparency disclosures need not make a distinction between copyright-protected works and works in the public domain.

At the outset, we want to clarify a distinction implied in the question above: Works may be included in one or more datasets, but the mere inclusion of a work in a dataset raises different legal questions than the use of a work to train a model.

Transparency as to which copyrighted works are included in a dataset, as well as which datasets were actually used in AI model development, could provide value for various stakeholders, including creators of copyrighted works and researchers. The utility of this knowledge to those with copyrights, however, will depend on what means of legal recourse are viable in practice and whether disclosure of this information is made before or after using the works for training purposes.¹⁸

If disclosure of this information occurs *after* using the works, then it could be valuable as a way to provide clear credit to people whose works contributed to training AI and to researchers who want to understand what works were actually used to train AI. In some cases where there are problems with the permissions for the training process or the outputs, this disclosure could serve as evidence to support claims of infringement. These claims could stem either from the unauthorized use of a work in model development or from a generative AI system's output, where a work's inclusion in training materials could demonstrate the system's access to the work as part of a "substantial similarity" test. As discussed in response to Question 23, courts could also establish a presumption of access in lieu of requiring proof that a work was part of a model's training data.

If disclosure of a work's inclusion in a dataset occurs *before* that dataset is used to train an AI model, this would serve a different purpose. It would also help credit people whose work is likely contributing to many AI training processes (particularly in large repeatedly used datasets). It could also, in theory, allow people to object to the use of their work or request that their work be removed from the dataset before it's used. In practice, disclosures made prior to use of a dataset may also need methods to make them easy to find and a mechanism for granting or denying consent for the use of works.

15.1. What level of specificity should be required?

From a copyright perspective, if the goal is to provide sufficient information to enable people to determine whether their works appear in AI model training materials, we suggest the following for recordkeeping and disclosures:

Records should describe the primary sources of data and the nature of those sources, including indications of whether the dataset includes personal information collected in a structured manner (such as intentional collection of names, addresses, or similar details) or not.

If training material is drawn from an existing dataset or datasets, records should identify and refer back to them. When new collections of training materials are compiled, records should disclose the collection methodologies. When using automated collection methods to compile training corpora, records should include the source code for any scraping programs that dataset builders use. Ideally, those automated collection systems will create logs, including timestamps,

¹⁸ Here we mean actually having the work be rated in some manner that adjusts the parameters of an AI model, as opposed to merely being in a dataset but not actually having been used in a training process.

of the sources they accessed. At a minimum, disclosures should provide enough information to enable other parties to replicate the collection.

15.2. To whom should disclosures be made?

If the purpose of disclosing these records is to provide notice to people whose works may have been used in the training process, then they should be disclosed to the general public to maximize their availability to all those whose works may be included. Public disclosure also promotes greater accountability, empowers researchers, and mitigates opportunities for abuse through underinclusive or selective disclosures made in bad faith.

Generative AI Outputs

Copyrightability

18. Under copyright law, are there circumstances when a human using a generative AI system should be considered the “author” of material produced by the system? If so, what factors are relevant to that determination? For example, is selecting what material an AI model is trained on and/or providing an iterative series of text commands or prompts sufficient to claim authorship of the resulting output?

No, the human or humans providing inputs to generative AI systems should not be considered authors of the outputs. We believe that this produces the best result for society and most effectively serves the purposes of copyright law.

Generative AI systems cannot be said to produce the expressions of a human operator because they do not produce consistent, predictable outputs for any given input. Generative outputs reflect a series of predictions made using the system’s statistical models such that the selection and arrangement of each word or pixel is probabilistic and cannot be predetermined by the human operator. However, as the Office has noted, humans may become authors for some aspects of a work involving generative AI outputs by modifying or adding new elements that are not generated by an AI system.¹⁹ Some prompts, such as human drawings or sufficiently long and creatively written text, may also themselves be copyrightable as a human work even if the output that results is not. We support the Office’s case-by-case approach to determining which elements of a work are eligible for copyright protection.

We encourage the Office to avoid venturing into more complex line drawing exercises, such as attempting to quantify some threshold level of human inputs to a generative AI system sufficient to confer authorship of the output to the human. At the same time, we acknowledge that human-created inputs to a generative AI system may meet the necessary requirements to be

¹⁹ See letter from US Copyright Office Review Board to Ms. Tamara Pester, 5 September, 2023: <https://www.copyright.gov/rulings-filings/review-board/docs/Theatre-Dopera-Spatial.pdf>.

eligible for copyright protection and may be carried over into part of the output (such as instructing a system to retain a drawing provided as input and add new material around it).

20. Is legal protection for AI-generated material desirable as a policy matter? Is legal protection for AI-generated material necessary to encourage development of generative AI technologies and systems? Does existing copyright protection for computer code that operates a generative AI system provide sufficient incentives?

Legal protection for AI-generated material is unnecessary and undesirable as a policy matter.

First, it is not clear whether increasing the production of generative AI content is desirable. Although these systems can improve the efficiency of content production and offer potentially large cost savings in certain applications, content production may not be the most socially beneficial use case for systems built on large language models. Further, we urge the Office to preserve copyright law’s human-centric approach in order to encourage the creation of new expressions by humans—a feature we see as beneficial to society as well as to the generative AI systems derived from works created by people.

Second, the efficiency gains offered by generative AI systems, especially those that generate text, appear to already offer sufficient incentive to motivate many people and companies to adopt and use these applications. Similarly, companies continue to invest heavily in the development of generative AI systems in the absence of copyright protection for the outputs of those systems. Given the high levels of investment in and use of generative AI systems, no additional incentives appear necessary at this time.

Finally, the long-term ramifications of extending copyright protection to AI-generated outputs could be catastrophic to the public domain and the future of creativity. For example, the potential scale at which AI systems could generate new works and the possibility to generate countless variants of expressions could effectively “lock out” future works by human authors.

Infringement

22. Can AI-generated outputs implicate the exclusive rights of preexisting copyrighted works, such as the right of reproduction or the derivative work right? If so, in what circumstances?

It is possible, if not likely, that some AI-generated outputs could infringe rights associated with one or more existing works. An AI-generated output could so closely resemble a protected work as to infringe the right of reproduction. Likewise, an AI-generated output could infringe on an author’s right to produce derivative works. However, the fact that a generative AI system was used to produce a potentially infringing work does not seem to significantly change the legal analysis necessary to determine infringement. Rather, its primary impact is adding complications to determine the identity of the infringer. This is because there may be separation in intent between the trainer of the system, the deployer of the system, and the user of the system. As a

simple example, a user who takes great care and time to prompt a system into generating an infringement would likely be an infringer, whereas a user who entered “cat” and happened to get an output that infringed the derivative rights of an existing painter of cats likely would not be an infringer, raising a question as to whether the model trainer or the model deployer ought to bear liability in that case.

23. Is the substantial similarity test adequate to address claims of infringement based on outputs from a generative AI system, or is some other standard appropriate or necessary?

The substantial similarity test remains an adequate means of addressing infringement claims, even those based on generative AI outputs. In cases where there is insufficient evidence to ascertain whether an author’s works were included in the training materials for one or more models used by a generative AI system, courts may need to establish a presumption that generative AI systems capable of producing potentially infringing works had “access” to the work prior to producing the allegedly infringing output. The remainder of the analysis need not change.

Additional Questions About Issues Related to Copyright

31. Should Congress establish a new federal right, similar to state law rights of publicity, that would apply to AI-generated material? If so, should it preempt state laws or set a ceiling or floor for state law protections? What should be the contours of such a right?

Although replicating a person’s likeness has long been possible even without technology, the ease with which accurate digital replications can be made using AI systems may require new policy considerations. We take no position on whether a new federal right is necessary at this time, but we offer the following observations:

We appreciate that the Office seeks to consider a wide range of policy implications in the context of generative AI, and we agree that copyright and a right of publicity could intersect in any number of fact patterns. However, if the Office plans to make recommendations to US Congress, we urge the Office to clearly delineate these two legal constructs and to avoid creating a perception that a publicity right should be modeled on or based in copyright law.

Second, recognizing and establishing new legal protections for such a right requires careful balancing against the First Amendment’s speech protections. In particular, a broad publicity right that allows a person to block or remove uses of their likeness must not limit the ability of others to document or report on items in the public interest, including the use of a person’s likeness where relevant. Existing causes of action like libel and defamation already offer pathways for abusive strategic lawsuits against public participation (SLAPP). The Wikimedia Foundation as well as volunteers who edit Wikipedia and other projects have been named as defendants in such SLAPP cases, often by people attempting to remove accurate documentation of their prior

actions.²⁰ Any new right or cause of action must be sufficiently narrow to mitigate this form of abuse.

32. Are there or should there be protections against an AI system generating outputs that imitate the artistic style of a human creator (such as an AI system producing visual works “in the style of” a specific artist)? Who should be eligible for such protection? What form should it take?

No. Like ideas and facts, styles and genres are not copyrightable. Artists have always mimicked the styles of others—it is through such iteration, with variations, that new styles and new forms of expression emerge. The use of a new tool in the iterative creative process does not change the foundational bases for and limitations to the scope of copyright protection. Additionally, such a prohibition would likely not survive any level of legal scrutiny under the First Amendment, especially not the strict scrutiny that courts typically apply to content-based restrictions on speech as proposed in the question.

²⁰ For example, the Wikimedia Foundation is currently defending itself and the privacy of several volunteers against a lawsuit based on well-sourced public information on the Portuguese language version of the Wikipedia project. See Jacob Rogers and Kabir Choudhary, *High stakes for the Wikimedia projects in Portugal: Fighting a strategic lawsuit against public participation (SLAPP)*, 27 July, 2023, <https://diff.wikimedia.org/2023/07/27/high-stakes-for-the-wikimedia-projects-in-portugal-fighting-a-strategic-lawsuit-against-public-participation-slapp/>.