AMERICAN BAR ASSOCIATION
Intellectual Property Law Section

321 N. Clark Street
Chicago, IL 60654-7598
T 312.988.6254
iplaw@americanbar.org
www.americanbar.org/iplaw

October 30, 2023

Submitted by Online Submission Procedure
Hon. Shira Perlmutter
Register of Copyrights and Director
U.S. Copyright Office
101 Independence Ave., S.E.
Washington, DC 20559

RE: Notice of Inquiry and Request for Comments: Artificial Intelligence and Copyright (Docket No. 2023-18624)

Dear Register Perlmutter:

We are writing to express the views of the American Bar Association's (the "Association") Section of Intellectual Property Law (the "Section") responding to the U.S. Copyright Office ("Office") August 30, 2023 Notice of Inquiry and Request for Comment ("NOI") on artificial intelligence ("AI") and copyright issues. The views expressed herein have not been reviewed or approved by the House of Delegates or the Board of Governors of the American Bar Association and, accordingly, should not be construed as representing the policy of the American Bar Association. These comments do not represent the policy or views of any government employee who is a member of the Section, its Council, or its Interest Groups.

Prior to the NOI, the Section participated in two prior requests for comments concerning AI that are relevant to the Office's inquiry: (1) the U.S. Patent & Trademark Office's requests for comments on Intellectual Property Protection for Artificial Intelligence Innovation, 84 Fed. Reg. 44889 (August 27, 2019) and 84 Fed. Reg. 58141 (October 30, 2019) (respectively, "USPTO AI Patent Comments" and "USPTO AI IP Comments"); and (2) the World Intellectual Property Organization's request for comments on its Draft Issues Paper on Intellectual Property Policy and Artificial Intelligence, WIPO/IP/AI/2/GE/20/1 (December 13, 2019) ("WIPO Comments"). Those prior comments are enclosed herewith for your review.

Before responding to the Office's specific questions, we share the Section's views on three critical issues. *First*, the Section opposes, in principle, recognizing an artificial intelligence as an "author" under US copyright law. An AI process or machine (*e.g.*, AI agent) does not, on its own, satisfy the

requirements to be an author set forth in the U.S. Supreme Court's opinions concerning what it means to be an author. Instead, an implicit component of the Court's decisions is that human creation is necessary for authorship. The core question, then, is when a human being uses a machine or device (e.g., a camera used to make a photograph, or an AI used to make an image) in developing the human's conception, what is the nature or degree of human participation required to treat the creation as the product of human authorship.[1] Our view is that the same reasoning that has been applied to other machines used to create works of authorship should be applied where a human being uses an AI agent or other instrumentality – e.g., through programming or input – to develop their conception and fix it in a tangible medium of expression. *Second*, the Section opposes, in principle, recognizing an artificial intelligence as an assignee, licensee, or other type of party having an ownership or possessory interest to a copyright recognized under Title 17. *Third*, the Section opposes, in principle, a new sui generis law to supplement US copyright, patent, trade secrets, data access (e.g., Computer Fraud and Abuse Act, 18 USC § 1030) or contract law to protect artificial intelligence data sets and databases.

We appreciate the Office's inquiry on these matters and support the overall goal of exploring complex questions and issues, including those highlighted by the Office such as the use of copyrighted works to train AI models, the appropriate levels of transparency and disclosure with respect to the use of copyrighted works, and the legal status of AI-generated outputs. We also appreciate the Office's inclusion of a glossary of key terms to ensure consistent terminology and reference to AI technologies are used by all stakeholders. Capitalized terms in this response have the meaning set out in the NOI's glossary unless otherwise noted. We have answered below ten (plus one subpart) of the questions that are presented in the NOI. Additionally, we recognize the need and potential benefit of having more time to consider and debate input collected from its members at a deeper level given the variations in, and still-developing circumstances surrounding, how AI systems are developed, used, and protected.

I. General Questions

Question 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?

Generative AI Systems, as that term is defined by the Office in the NOI, have the potential to automate many tasks in "knowledge work." "Knowledge work" can be defined as workers who apply theoretical and analytical knowledge, often acquired

[1] By "conception" we refer not to a mere idea (which is not protected by copyright) but to the author's detailed plan for creating a work. As the U.S. Supreme Court recognized in the context of photographs, although they use a mechanical device (a camera), they are "representatives of original intellectual conceptions of the author." *Burrows-Giles Lithographic Co. v. Sarony*, 111 U.S. 53, 58 (1884).

through formal training, to develop products and services.[2] "Knowledge workers" include computer programmers, web designers, system analysts, technical writers, researchers, pharmacists, public accountants, engineers, architects, physicians and other healthcare workers, scientists, financial analysts, design thinkers, and, notably, lawyers. Although there has been a surge in the proliferation of commercially available Generative AI Systems over the past year, this technology is still in its earliest stages, so the impact on various stakeholders may be little more than speculation. However, credible studies suggest that Generative AI Systems will follow the pattern of similar disruptive technologies in that they will displace certain categories or levels of knowledge workers while creating opportunities for others.

One recent study suggests that Generative AI Systems may fill a critical gap in ensuring the future health of the U.S. economy, namely by offsetting recent declines in worker productivity and counterbalancing the impact of an aging worker population. [3] Moreover, we also note that although the Office is focused on the impact of Generative AI Systems, there is considerable entanglement between those systems and the "AI Model" (as defined by the Office) that enable the Generative AI Systems. Foundation AI Models, such as GPT-4 and LLaMA, by their definition, have a myriad of uses beyond consumer-facing chatbots and image-generation services. Alongside of the proliferation of ChatGPT, for example, is a proliferation of private instantiations of GPT-4 and other foundation AI models to enhance the operations of corporate stakeholders at multiple levels of the enterprise, from research and development, to human resources, to marketing. Large Language Models are also being trained for industry-specific applications e.g. law, medicine, insurance etc. The anticipated gains in efficiency and R&D insight from these enterprise instantiations are only beginning to be assessed.

In addition to its impact on knowledge work, Generative AI Systems may also have a significant impact on creative industries. Research from Goldman Sachs suggests that generative AI has the potential to automate 26% of work tasks in the arts, design, entertainment, media and sports sectors.[4] While the same study also predicted significant productivity gains and a potential increase in global gross domestic product by 7%, such gains may come at the cost of some displacement of pre-existing creative jobs.

There are also pre-existing and continuing developments of alternate technologies of machine learning, including more specialized use of neural networks, as well as of surrounding technologies such as filtering and adapting of raw data and application and user interfaces, much of which is being performed by smaller entities and "independent" developers, including authors and artists. Some of these, exemplified by apps used on smartphones or specialized medical image analytics, might be characterized as generative

---

[2] Peter Drucker, *The Landmarks of Tomorrow* (1959).
[3] McKinsey & Company, The Economic Potential of Generative AI, (June 14, 2023) https://www.mckinsey.com/capabilities/mckinsey-digital/our-insights/the-economic-potential-of-generative-ai-the-next-productivity-frontier
[4] Joseph Briggs and Devesh Kodnani, "The Potentially Large Effects of Artificial Intelligence on Economic Growth," March 23, 2023. Goldman Sachs. ,https://www.key4biz.it/wp-content/uploads/2023/03/Global-Economics-Analyst_-The-Potentially-Large-Effects-of-Artificial-Intelligence-on-Economic-Growth-Briggs_Kodnani.pdf .

AI but may warrant different legal treatment than the large Foundation AI Models that may cost upwards of a billion dollars to create, practically limiting new entrants.

The U.S. Department of Health and Human Services (HHS) published its Trustworthy AI (TAI) Playbook in September 2021.[5] The HHS TAI Playbook contains a succinct overview of how AI use cases are broken down by various AI solutions, which are comprised of one or more AI methods.[6] Generative AI solutions are but just one example of how several AI methods can work together to power solutions across countless use cases.

In our view, great care needs to be taken to ensure that the deleterious effects of Generative AI Systems are minimized, while not simultaneously thwarting the development of a technology that may yield crucial benefits to the U.S. economy and U.S. global competitiveness. A rush to provide support because of perceived benefits (or impose obligations because of perceived threats) from AI to one or more industries may have unintended consequences on other industries that are no less valuable to the U.S. economy and international competition. Aggressive throttling of AI could impede the development of foundation models, while aggressive support could make it even more difficult for individual authors and small creators to create new non-AI-assisted works. We applaud the Office's efforts to closely monitor the development of these technologies, and encourage further thoughtful monitoring and assessment of them before deployment of regulatory and other Office action. The purpose of the Copyright Clause of the U.S. Constitution in promoting the progress of science and the useful arts, is best achieved by the careful calibration of disincentivizing infringement and other risks posed by Generative AI Systems, while incentivizing the development of these powerful new tools. .[7]

Question 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?

Our members are engaged in law practice, commentary, and policy development in every industry imaginable. In addition to our representation of other stakeholders, we are both authors and users of copyrighted material, with such copyrights often uniquely modulated by public and governmental functions, and supplemented by plagiarism norms. (Many of the AI models have likely been trained on materials, such as contracts, briefs, and research memos, for which our members or their clients may own copyright.) Issues raised from the earliest days of automation in the legal field (such as automated search of legal material databases, automated document assembly, and predictive coding of document screening) have been magnified by the current availability of automated generation of much larger amounts of text with relatively little understanding of underlying models.

---

[5] U.S. Department of Health and Human Services, "Trustworthy AI (TAI) Playbook". https://www.hhs.gov/sites/default/files/hhs-trustworthy-ai-playbook.pdf .

[6] *Id*. *See* pages 11-13.

[7] U.S. Constitution, Article I, Section 8, Clause 8.

As it relates to the practice of law, Generative AI Systems may impact tasks that have been traditionally time-intensive such as e-discovery and document review, generation of documents such as contracts, letters, complaints, summaries of discovery documents and deposition transcripts, and legal research. [8] Generative AI Systems may put pressure on traditional hourly billing fee structures and speed to producing work product, and could drive client demand for deployment and use of these systems to achieve greater value.

Lawyers may be interested in adopting the use of Generative AI Systems in some form. In 2012, the American Bar Association amended Comment 8 to Model Rule of Professional Conduct 1.1 (Lawyers Duty of Competence) to address technology competency. The Association amended the comment to address changes in technology and its use in legal practice, stating:

> To maintain the requisite knowledge and skill, a lawyer should keep abreast of changes in the law and its practice, including the benefits and risks associated with relevant technology, engage in continuing study and education and comply with all continuing legal education requirements to which the lawyer is subject.

Other ABA Model Rules may also be implicated. For example, Rule 1.6 sets forth a duty of confidentiality that may be implicated in prompts to Generative AI Systems where the prompts including client confidences are made available to a platform operator or incorporated into an underlying model. Rules 5.1 and 5.3 require the supervision of both lawyers and non-lawyers of their legal work that may involve Generative AI Systems.

Rules 8.4(d) and 8.4(g) proscribe assistance in discrimination and harassment and have been argued to apply to the use of biased automated systems, including Generative AI Systems. Bias may be found in the raw training material, the pre-processing of that material supplied to the model-generating process, the platform/user interface, and the prompting; that bias may not be intentional but may reflect societal distributions of information and circumstances. Copyright and neighboring issues may arise in claims of cultural or style appropriation that smaller developers may be less able to defend than incumbents.

Early judicial concern about the uninformed use of popular "chatbots" to write briefs has led to some blanket prohibitions against any use, which could chill efficient and ethical use. Similarly broad law firm prohibitions may be ineffective (as it was with the early use of the Internet) or unduly favor competitively those who can field in-house systems.

Thus, in some ways our members represent a microcosm of the benefits and disruption that are facing other sectors of the U.S. economy. Many of the AI models have likely been trained on materials, such as contracts, briefs, and research memos, for which our

[8] Andrew Perlman, The Implications of ChatGPT for Legal Services and Society, The Practice, Harvard Law School Center on the Legal Profession, March 2023, https://clp.law.harvard.edu/knowledge-hub/magazine/issues/generative-ai-in-the-legal-profession/the-implications-of-chatgpt-for-legal-services-and-society/ .

members or their clients may own copyright. Those models when used in Generative AI Systems may now threaten to disrupt the ordinary course of business for legal authors, but in a fashion that may ultimately benefit clients and society at large by driving down costs. Over time, Generative AI Systems fine-tuned to the legal profession have the opportunity to improve access to justice and democratize access to law and legal services.[9]

Question 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.**

The Section will cite and link relevant and supportive papers and studies in footnotes throughout its comments.

Question 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?**

We support efforts by the Office to closely monitor the legislative and policy developments of other countries to regulate the development, deployment, and use of Generative AI Systems. Consistency of operations across jurisdictions is generally desirable for most copyright stakeholders, and inconsistent rules and regulations may hamper the efforts of the U.S. to achieve the carefully calibrated approach suggested in response to Question 1 above.

We are paying particular attention to the Artificial Intelligence Act (AI Act) proposed by the European Commission in the European Union (EU), which currently stands out as the most comprehensive attempt to address a number of perceived threats and misuses of Generative AI Systems. Especially relevant to the Office's NOI is Article 28b(4)(c), which states that the "provider" of a foundation model enabling a Generative AI System must "document and make publicly available a sufficiently detailed summary of the use of training data protected under copyright law,"[10] which must be done before the model is placed on the market or put into service in the EU (Article 28b(1)). It is not currently clear how this provision will be implemented, and whether the "sufficiently detailed summary" is contemplated on a per work basis, an identification of websites where training material may have been scraped, or some broader summary of AI model training.

---

[9] American Bar Association, Big problems (and benefits) of generative AI are here, August 2023. https://www.americanbar.org/news/abanews/aba-news-archives/2023/08/problems-and-benefits-of-ai/ .

[10] Note that the use of "training data" in this quote would most likely be equivalent to the definition "training material" supplied by the Office in the NOI.

This provision is likely to have considerable impact on the U.S., as many companies developing and licensing AI Models are headquartered in the U.S., but their platforms are available for use worldwide.[11]

Moreover, it will be important for the Office and U.S. policymakers to monitor approaches to platform liability. The EU AI Act appears to contemplate governance of U.S.-based companies.[12] The proliferation of customized instantiations and licensed uses of AI Models create a complex ecosystem in which the lines between who is a Generative AI developer, who is a Generative AI user, and who is the owner of training datasets, may be blurred. Article 3 of the EU AI Act attempts to address these differing roles, but the distinctions between "provider," "user," and "distributor" under the Article 3 Definitions of the AI Act are not entirely clear. Similarly, the EU AI Act's approach to regulating AI products and services according to categories of risk is not fully developed, giving rise to ambiguities that could disincentivize the development of certain AI technologies.[13]

Finally, we note that attention should be given to the global variations on the recognition of authorship of outputs created by Generative AI Systems—referred to as "non-human authorship" by the Office.[14] Article 6 of the Berne Convention provides that "Authors shall enjoy, in respect of works for which they are protected under this Convention, in countries of the Union other than the country of origin, the rights which their respective laws do now or may hereafter grant to their nationals, as well as the rights specially granted by this Convention." Many countries are still developing their approaches to copyright protection for AI-generated works, but some countries, such as the UK, do recognize at least some limited protection for such works, and others, including the US, are grappling with guidelines to identify when work that is created with or through use of Generative AI Systems constitutes copyrightable subject matter.[15] Failing to harmonize the U.S.'s position on the protection of AI-generated works may result in tensions, if not violations, of the U.S.'s obligations under the Berne Convention.[16] Moreover, inconsistency may threaten increased litigation costs and protracted discovery related to disputes over the enforceability of foreign works that implicate AI. Lastly, divergent developments of policy and regulation surrounding the recognition of non-human

---

[11] Lie, April, Europe's new AI law threatens US firms, https://thehill.com/opinion/technology/4014747-europes-new-ai-law-threatens-u-s-firms/.

[12] Article 3 of the EU AI Act includes definitions for "placing on the market", "making available on the market", and "putting into service," and "importer" that together suggests it is meant to reach beyond EU-based companies only.

[13] U.S. Chamber of Commerce, Navigating the EU AI Act: Striking a Responsible Balance, https://www.uschamber.com/international/navigating-the-eu-ai-act-striking-a-responsible-balance

[14] U.S. Copyright Office Review Board. Second Request for Reconsideration for Refusal to Register Théâtre D'opéra Spatial (SR # 1-11743923581), Sept. 5, 2023.

[15] "Computer-generated" is defined as "generated by computer in circumstances such that there is no human author of the work" (Section 178, Copyright, Designs and Patents Act (CDPA)). Section 9(3) of the CDPA provides that the author of a computer-generated work is deemed to be the person "by whom the arrangements necessary for the creation of the work are undertaken."

[16] Ginsburg, J.C. People Not Machines: Authorship and What It Means in the Berne Convention. IIC 49, 131–135 (2018). https://doi.org/10.1007/s40319-018-0670-x

authorship could lead to forum shopping for companies and organizations seeking to innovate and develop AI technologies, including generative AI.

Question 5. Is new legislation warranted to address copyright or related issues with generative AI? If so, what should it entail? Specific proposals and legislative text are not necessary, but the Office welcomes any proposals or text for review.

We believe that it is premature to develop specific legislation to address the issues raised in the Office's NOI, but that continued rapid development and deployment of Generative AI Systems, and the concomitant proliferation of legal issues attending such proliferation, warrant constant monitoring and assessment. Consideration of new legislation might be prompted, for example, by the EU's adoption and further refinement to the AI Act, either in adopting certain approaches or distinguishing treatment of these issues in the U.S. Similarly, potential domestic legislative efforts in legal areas impacted by AI technologies other than copyright law, such as privacy, right of publicity, consumer protection, and content moderation, may inform copyright policy and should similarly be monitored.

With regard to AI, we believe that a balance of factors must be considered when proposing a legislative solution, and we oppose any legislative change that would shift the "conception" aspect found in copyright law away from the human author to an AI agent. Thus, it is our position that the current legal frameworks together with patent, trade secret, and contractual arrangements provide the required balance between the various competing factors with respect to authorship and ownership.[17]

Additionally, we point to the development of responsible AI frameworks and standards as an early manner of self-governance and self-regulation by various stakeholders across industries. The previously mentioned HHS TAI Playbook serves as a robust example of a responsible AI framework for the adoption of AI technologies.[18] Content authenticity and provenance are the goals of the Content Authenticity Initiative, which offers open-source tools and resources in support of its focus on the creation, distribution, and consumption of both traditional and synthetic media.[19] Another example focused on the output is the Partnership on AI's Responsible Practices for Synthetic Media.[20]

The Center for Security and Emerging Technology, a policy research organization within Georgetown University's Walsh School of Foreign Service, published a paper examining process frameworks for organizations implementing responsible AI. However, their research notes that process frameworks can complicate the ability to implement responsible AI frameworks that have been broadly drafted from a high-level governance

[17] https://www.americanbar.org/content/dam/aba/administrative/intellectual_property_law/advocacy/aba-ipl-comment-letter-on-ec-ai-white-paper.pdf
[18] HHS at 15-23.
[19] Content Authenticity Initiative, https://contentauthenticity.org/ .
[20] Partnership on AI, "Responsible Practices for Synthetic Media". https://syntheticmedia.partnershiponai.org/ .

viewpoint to the development and production groups applying and utilizing those frameworks in practice.[21]

It is noted that as of the date of this letter 44 bills have been introduced in Congress that primarily concern AI, or would institute significant regulations, policies, or programs concerning AI. The Section regularly monitors these developments, and is committed to continuing dialog with the Office to offer the expertise of its members as needed and appropriate.

II. Training

The Section has no comment on Questions 6 and 7.

Question 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.

Opinions differ on whether current applications of the fair use doctrine, and current theories of infringement (particularly with regard to distinguishing between an author's "ideas" or "style" and the copying the author's expression) are sufficient to address concerns regarding whether use of a copyrighted work as training data constitutes infringement. Moreover, Generative AI Systems vary in how they use copyrighted works, particularly in the ways that works are ingested, retained, and/or reused. The combination of limited information regarding the precise training methods used to develop currently deployed AI tools, and the rapidly evolving and dynamic nature of the way that various Generative AI Systems and AI Models are adopted, adapted, and customized across myriad industries, make it difficult to issue a one-size-fits-all proclamation that AI training is or is not fair use. Put differently, the fair use assessment is highly fact dependent, and the facts regarding AI training are oftentimes either unknown or in flux.

The American Bar Association, however, encourages a consistent approach to the application of copyright's fair use doctrine in the context of the digital environment.[22] Likewise, we support the principles that the judicially-created concept of "transformative use" should be applied in the fair use analysis with due consideration for the derivative work right and that courts should consider each of the statutory factors enumerated in 17 U.S.C. § 107, and that the judicially-created concept of "transformative use" should be considered as only part of the analysis under the "purpose and character of the use" factor. We also support the principle that, when considering the "effect of the use upon the potential market for or value of the copyrighted work" in the fair use analysis, courts should consider whether a "traditional, reasonable, or likely to be developed" market does or could exist for the secondary use at issue and whether unrestricted and

---

[21] Mina Narayanan and Christian Schoeberl, "A Matrix for Selecting Responsible AI Frameworks" (Center for Security and Emerging Technology, June 2023). https://doi.org/10.51593/20220029 .

[22] https://www.americanbar.org/content/dam/aba/administrative/board_of_governors/greenbook/greenbook.pdf

widespread conduct of the sort engaged in by the proponent of the fair use defense would impact that market, which ought not be limited to sheer quantification of the harm to that market such as might occur in an assessment of damages.

Cases in which fair use is at issue are rarely ever "open and shut" cases. For example, at one extreme, a fair use argument would be at it weakest in a circumstance where the purpose of the creation of the copyrighted material was to train an AI Model (or where the copyrighted material is itself a Training Dataset) and a market for that purpose already exists. A fair use assertion would be much stronger when the purpose of the material does not include use to train an AI and no market does or could reasonably exist for that purpose. For situations between those two extremes, courts are likely to be guided by existing fair use caselaw.

Question 8.1. **In light of the Supreme Court's recent decisions in Google v. Oracle America and Andy Warhol Foundation v. Goldsmith, how should the "purpose and character" of the use of copyrighted works to train an AI model be evaluated? What is the relevant use to be analyzed? Do different stages of training, such as pre-training and fine-tuning, raise different considerations under the first fair use factor?**

See response to 8 above.

The Section has no comment on Questions 8.2 through 8.5 and 9 through 14.

III. Transparency & Recordkeeping

The Section has no comment on the questions in Section IV (Transparency & Reporting).

IV. Generative AI Outputs

a. Copyrightability

Question 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?**

The Section references its response to the USPTO AI IP Comments, particularly with regard to questions 1 and 2. In particular, it is our position that a Generative AI System, standing on its own, cannot be considered an author. That is because a work produced by an AI algorithm or process, without the involvement of a human, does not and should not qualify as a work of authorship protectable under U.S. copyright law.

It, however, is our view that there are certain instances where users of generative AI systems could be considered authors. For example, a person who caused AI to yield a new work might be an author, provided he or she contributes sufficient human authorship. Likewise, a person using AI to create a new song or image, or to enhance or create a mashup of existing images, absent a contract to the contrary, would own copyright in the output, also provided he or she contributes sufficient human authorship. The copyright in the new work would not extend to the code that helped generate the new work or any existing work ingested to train the AI. In sum, a person could make contributions to AI that would be copyrightable, and authorship would reside in the persons who created an identifiable copyrightable work. If there are multiple contributors that are not all employees of the same work for hire employer, then they will be joint authors. By way of illustrative example, the popular video game MineCraft, which allows its players to create their own worlds, does not vest copyright in the programmers of the game unless re-assignment back to the game owners is stipulated in the licensing.

As noted by former Register Barbara Ringer, and as highlighted in the NOI, the Office ought not "take the categorical position that registration will be denied merely because a computer may have been used in some manner in creating the work." NOI at 3-4. Further, as was recently noted in a United States District Court decision addressing the registrability of works created using AI tools, Congress's power to protect "writings" is not "limited to script or printed material," but rather encompasses "any physical rendering of the fruits of creative intellectual or aesthetic labor" *Thaler v. Perlmutter*, No. 22-cv-1564, 2023 WL 5333236, at *4 (D.D.C. Aug. 18, 2023), quoting *Goldstein v. California*, 412 U.S. 546, 561 (1973). Thus, although authorship ought not be granted to a Generative AI System as a theoretical stand-alone entity, there are many instances where a human's interaction with, and direction of, AI processes may result in copyrightable expression.

The principle that U.S. copyright law protects works beyond traditional forms of expression was recently reaffirmed by the U.S. District Court for the District of Columbia, in *Thaler v. Perlmutter*, No. 22-cv-1564, 2023 WL 5333236, at *4 (D.D.C. Aug. 18, 2023). In affirming the Office's denial of registration of a work purportedly created "autonomously" by a Generative AI System, the court emphasized the fact that "human creativity is the *sine qua non* at the core of copyrightability, even as that human creativity is channeled through new tools or into new media." Id at *3. The court analyzed the creative interaction between human and machine, and identified the kinds of human activity that might yield protectable expression emanating from a machine, by looking to the process used in the creation of a photograph as an illustrative example:

> A camera may generate only a "mechanical reproduction" of a scene, but does so only after the photographer develops a "mental conception" of the photograph, which is given its final form by that photographer's decisions like "posing the [subject] in front of the camera, selecting and arranging the costume, draperies, and other various accessories in said photograph, arranging the subject so as to present graceful outlines, arranging and disposing the light and shade, suggesting

> and evoking the desired expression, and from such disposition, arrangement, or representation" crafting the overall image. [*Burrow-Giles Lithographic Co. v. Sarony*, 111 U.S. 53, 59-60 (1884)]. Human involvement in, and ultimate creative control over, the work at issue was key to the conclusion that the new type of work fell within the bounds of copyright.

The court noted that, "Copyright has never stretched so far, however, as to protect works generated by new forms of technology operating absent any guiding human hand…" The question then becomes, where in the various potential uses of Generative AI Systems is the "guiding human hand"?

As part of the AI Initiative, experts, stakeholders, and other industry leaders shared their varying viewpoints on this subject. They encouraged the Office to consider both the details of human interaction with Generative AI Systems, and various theories of "authorship" that may sustain a copyright across industries. For example, the "iterative series of text commands or prompts" identified in this question number 18 could possibly be sufficient to claim authorship when it reflects the Generative AI System's user's "mental conception" and creative expression through iteration or compilations.

One possible example might be rejecting, editing, or refining the initial outputs of a Generative AI System and guiding and iterating with the system to produce text or an image that meets with the user's "mental conception". The preceding example demonstrates a measure of control and expressive vision, and may indeed be sufficiently creative to warrant ownership of a copyright in the output. Generative AI Systems operate much like a tool, that through continuous iteration and editing, the human author can develop a creative work that is akin to the author's "mental conception" or creative vision. The field of "prompt engineering," in which individuals are trained to produce desirable results from Generative AI Systems, and creatively iterate and dialogue with the system until their "mental conception" or creative vision is achieved, is in its infancy and ought to be incentivized to grow.

Moreover, U.S. copyright law recognizes types of authorship beyond the "classic" authorship of a lone artist creating a work in a studio or writer sitting with a typewriter. In finding copyrightable authorship in photographs, the Court in *Burrow-Giles* relied in part on an 1883 decision from the U.K., *Nottage v. Jackson*, where in upholding copyright for photographs under U.K, law, the court "said, in regard to who was the author, 'The nearest I can come to, is that it is the person who effectively is as near as he can be, the cause of the picture which is produced, that is, the person who has superintended the arrangement, who has actually formed the picture by putting the persons in position, and arranging the place where the people are to be – the man who is the effective cause of that.'"

As noted by the Author's Guild in its comments to the U.S.P.T.O. in 2019:

> Acts of authorship can be attributed to an agent or amanuensis worker in the manner that large sculptures or art installations, for example, often are actually

> "made" by workers under creator's direction. As Professor Ginsberg explains, "[t]he law [following agency rules] attributes authorship to the "mastermind," whose detailed conception so controls [the work's] subsequent execution that the individuals carrying out the embodiment exercise no creative autonomy." An AI system following the detailed guidance of its users and/or programmers and under their authority would be an amanuensis, and authorship in the work generated by the AI would be attributed to the human masterminds.[23]

The notion of human beings exercising control over the ultimate work created in part with a Generative AI System by acting as the "mastermind" finds some support in other "mastermind" authorship contexts, such as film production. *See Aalmuhammed v. Lee*, 202 F.3d 1227 (9th Cir. 2000) ("Burrow-Giles defines author as the person to whom the work owes its origin and who superintended the whole work, the 'mastermind.' In a movie, this definition, in the absence of a contract to the contrary, would generally limit authorship to someone at the top of the screen credits, sometimes the producer, sometimes the director, possibly the star or the screenwriter – someone who has artistic control."); *16 Casa Duse, LLC v. Merkin*, 791 F.3d 247 (2d Cir. 2015) (granting authorship to the production company as dominant author). A user of a generative AI tool might similarly exercise the kind of decision-making authority that exists in the film production example, exercising ultimate control of what goes in, and what stays out, of the final work to best comport with the user's mental conception of the final work. [24]

The Office regularly issues registrations to claimants that assert authorship over works whose authorship is not described in detail on accompanying registration applications. For example, the Office registers large groups of photographs without inquiring whether a particular photograph was created by posing the subject, selecting costumery, adjusting lighting, etc. Similarly, registrations in films and source code are made without specific inquiry into a claimant's contributions. Details regarding exactly what the asserted author did in constructing the work are not typically addressed at the registration phase. Rather, presumptions of validity and assertions of authorship might not be raised at all except when the copyright is enforced in litigation. Similarly, authorship and control issues may also be at issue in licensing negotiations, such that the commercial viability of a work may act as a reasonable proxy of proof of authorship.

Our position is that works created solely or "autonomously" by Generative AI Systems should not be afforded copyright protection. We therefore support the Office's position that such stand-alone content should be disclaimed in the registration process. However, there is real potential that a copyright claimant will seek registration of a work with a good faith belief that the author has used a generative AI tool in a fashion that the author had final decision-making authority over the work, and that work represented the author's

---

[23] https://www.uspto.gov/sites/default/files/documents/The%20Authors%20Guild_RFC-84-FR-58141.pdf

[24] It is also noted that the U.K., while taking a different view on AI authorship than the U.S., grants limited rights to AI-generated works under a theory similar to mastermind authorship. "Computer-generated" is defined as "generated by computer in circumstances such that there is no human author of the work" (Section 178, Copyright, Designs and Patents Act (CDPA)). Section 9(3) of the CDPA provides that the author of a computer-generated work is deemed to be the person "by whom the arrangements necessary for the creation of the work are undertaken."

"mental conception," similar to registrations of photographs described above. The mere fact that an AI tool was used in the creation of a work ought not disqualify the work from registration. The Section suggests that the Office consider registration guidance that defers to the claimant's judgment regarding the authorship brought to bear, and allows details regarding such authorship to be validated in enforcement actions and licensing deals in the ordinary course.

Question 19. Are any revisions to the Copyright Act necessary to clarify the human authorship requirement or to provide additional standards to determine when content including AI-generated material is subject to copyright protection?

The Section does not believe that an amendment is needed to the Copyright Act to address authorship questions. Common law doctrines of authorship should be sufficient to address authorship standards as discussed in the Section's response to 18 above.

We, however, note that consideration of the complex question of human authorship may not be necessary to determine copyrightability in many circumstances. For example, the Copyright Act already provides that "protection for a work employing preexisting material in which copyright subsists does not extend to any part of the work in which such material has been used unlawfully."[25] Likewise, the copyright in a "derivative work extends only to the material contributed by the author of such work, as distinguished from the preexisting material employed in the work, and does not imply any exclusive right in the preexisting material."[26] Thus, if a Generative AI Systems yields a work that unlawfully creates a derivative of a copyrighted work, that derivative work would not be protected. If the Generative AI System does so lawfully, the protection in such output would be limited to additional material contributed by the human author. [27] Accordingly, the existing Copyright Act may provide an alternative to the human authorship requirement that will delineate copyrighted material generated by an AI and uncopyrightable material.

Question 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?

AI-generated material that meets the standards discussed in response to Questions 18 and 19 above, and represents the mental conception of the author as produced by the author's final decision-making authority may in the right circumstances warrant copyright protection. Creating hurdles to obtaining copyright protection in AI-generated material may disincentivize adoption of generative AI technologies. As noted in response to Question 1 above, the use of both Generative AI Systems, and use by a myriad of

[25] 17 U.S.C. § 103(a).
[26] *Id.* § 103(b).
[27] *Id*.

industries of the AI models that enable such systems, has the potential to drive needed efficiencies and help ensure the health and global competitiveness of the U.S. economy. Placing the lion's share of AI-generated outputs into the public domain is likely to hinder adoption of those technologies.

However, encouragement of the development of Generative AI Systems needs to be measured against the impact on human authors and artists in creative industries as well. The ability of Generative AI Systems to create numerous works for its users without the involvement of human artists could substantially harm the creation of new works by such artists and the market for their works, and could result in significant displacement of creative industry jobs *See, supra* p. 7. Thus, some have argued that it is preferable to require AI outputs to be expressly labeled as such and those outputs not to be protected, so that their utility will be limited to situations where users of Generative AI Systems are unconcerned about asserting copyright in such outputs.

Question 20.1. If you believe protection is desirable, should it be a form of copyright or a separate sui generis right? If the latter, in what respects should protection for AI-generated material differ from copyright?

We do not believe that a *sui generis* right is necessary at this time, and believe that current U.S. copyright doctrines and policy are sufficient to extend protection of sufficient scope.

Question 21. Does the Copyright Clause in the U.S. Constitution permit copyright protection for AI-generated material? Would such protection "promote the progress of science and useful arts"? If so, how?

As noted in response to Questions 18, 19, and 20 above, the Copyright Clause should allow for protection assuming sufficient human intervention and involvement is demonstrated.

The Section has no comment on Questions 22 through 34.

* * *

Should you have additional questions, we would welcome further discussion with the Office.

Sincerely,

Steven P. Caltrider, Chair
ABA Section of Intellectual Property Law