VIA EMAIL TO AIPartnership@uspto.gov
Andrei Iancu
Under Secretary of Commerce for Intellectual Property and
Director of the U.S. Patent and Trademark Office
P.O. Box 1450
Alexandria VA 22313-1450

RE: NYIPLA Response to USPTO NOI on Artificial Intelligence Innovation

Dear Under Secretary Iancu:

The New York Intellectual Property Law Association ("**NYIPLA**") welcomes the opportunity to respond to the Notification of Inquiry on Artificial Intelligence Innovation (the "**NOI**"), which was published in the October 30, 2019 edition of the Federal Register, 84 FR 58141. The NYIPLA is a professional association comprised of over 1,000 lawyers interested in Intellectual Property law who live or work within the jurisdiction of the United States Court of Appeals for the Second Circuit, and members of the judiciary throughout the United States as ex Officio Honorary Members. The Association's mission is to promote the development and administration of intellectual property interests and educate the public and members of the bar on intellectual property issues. Its members work both in private practice and government, and in law firms as well as corporations, and they appear before the federal courts and the United States Patent and Trademark Office. The NYIPLA provides these comments on behalf of its members professionally and individually, and not on behalf of their employers. In regard to the NOI, the NYIPLA offers the following comments and observations with respect to certain of the questions specifically addressing copyright law.

1. **Should a work produced by an AI algorithm or process, without the involvement of a natural person contributing expression to the resulting work, qualify as a work of authorship protectable under U.S. copyright law? Why or why not?**

Under current U.S. law, AI algorithms and processes cannot themselves be considered "authors": works of authorship are only created through the involvement of human beings.[1]

Article I Section 8 Clause 8 of the U.S. Constitution grants Congress the power "to promote the progress of science and useful arts, by securing for limited times to authors and inventors the exclusive right to their respective writings and discoveries." Congress has interpreted "authors" here to mean only natural persons. Even when a limited exception applies that results in the legal fiction that a legal entity has authored a work, it is necessarily the case that the creator(s) of the work in question were human beings. *See, e.g.*, 17 U.S.C. § 101 ("The authors 'widow' or 'widower' is the author's surviving spouse…"; "a work is a 'United States work' only if--…all the authors of the work are nationals, domiciliaries, or habitual residents of, or in the case of an audiovisual work legal entities with headquarters in the United States…"). Likewise, the Copyright Office has concluded that "[t]o qualify as a work of 'authorship' a work must be created by a

---

[1] As discussed in response to Question 2 below, however, certain lines of existing case law provide some framework for assessing whether the works produced by such algorithms and processes may nevertheless be protectable.

human being." Copyright Compendium Vol.3 at 313.2. The courts concur. *See, e.g.*, *Burrow-Giles Lithographic Co. v. Sarony*, 111 U.S. 53, 58 (1884) (Author means "he to whom anything owes its origin; originator; maker"); *Naruto v. Slater*, 888 F.3d 418 (9th Cir. 2018) (monkey lacks standing to sue under Copyright Act for use of selfie photo).

Whether U.S. copyright law *should* be amended to permit work produced by an AI algorithm or process to qualify as works of authorship raises significant questions about the policy and purpose of U.S. copyright law, namely whether affording machines author status would undermine the objective of U.S. copyright law, which is to protect the creative efforts of natural persons, and such a change would serve the constitutional purpose of "promoting the progress of science and useful arts." On the one hand, adding more works to a field of endeavor traditionally has been thought to further the progress of science and useful arts, as the more works there are to explore on a topic, the more advancements can be made. However, in the case of artificial intelligence, where AI devices can run non-stop twenty-four hours a day and can generate works at a much faster pace than humans, there may be a risk that one or more AI devices could flood a particular field with works and, in effect, block or discourage others from producing works in that field for fear of copyright infringement. Additionally, U.S. copyright law has traditionally distinguished between creative content, which is protected, and non-creative work that constitutes "sweat of the brow," which is not. Granting copyright protection to AI-generated work may blur those distinctions, leading to uncertainty and confusion on the scope of copyright protection.

In our view, in light of the rapid rate at which AI technology is currently developing, it is simply too soon to consider any specific amendments to the Copyright Act to address AI. If and when Congress concludes that amending U.S. copyright law to create authorship in AI-generated works, however, there are at least two approaches that could be taken that are worthy of consideration and further study.

The first such approach would be to consider the example set by certain other countries to amend the definition of "authorship" expressly to include AI-generated works. For example, Section 9(3) of the U.K.'s Copyright, Designs and Patents Act (1988) ("**CDPA**") provides:

> In the case of a literary, dramatic, musical or artistic work which is computer-generated, the author shall be taken to be the person by whom the arrangements necessary for the creation of the work are undertaken.

The CDPA defines "computer-generated" in relation to a work, to mean a work is generated by computer in circumstances such that there is no human author of the work. *See* CDPA § 178. Under this definition, "author" would include, at a minimum either by the programmer(s) of the AI device or their employer if the coding was done in the context of a work-made-for-hire arrangement. We note, however, that this definition was created in 1988, when artificial intelligence was in its infancy. The closer technology develops to create a device that "thinks for itself," the more antiquated this definition becomes. Also, identifying an AI algorithm or process as an "author" will have implications throughout U.S. copyright law wherever the concept of authorship is discussed. For example, if an AI algorithm or process can be an author, should it also qualify as an author for purposes for "joint authorship," which, under current U.S. law, requires a mutual intention of both authors that the work be merged together in a single work?

The second approach would be to amend the Copyright Act's definition so that AI devices could be considered an "employee" or recognized as having equivalent status of an "employee" making the output "works made for hire." For a discussion of this approach, see Kalin Hristov, Artificial Intelligence and the Copyright Dilemma, 57 IDEA 431 (2017), available at https://ipmall.law.unh.edu/sites/default/files/hosted_resources/IDEA/hristov_formatted.pdf . This approach seems to be a more narrowly tailored fix to the problem, as having the AI

algorithm or process constitute an employee solely for purposes of "works made for hire" does not implicate all the issues raised by amending the bedrock definition of authorship, which impacts a wide range of provisions under U.S. copyright law.

2. **Assuming involvement by a natural person is or should be required, what kind of involvement would or should be sufficient so that the work qualifies for copyright protection? For example, should it be sufficient if a person (i) designed the AI algorithm or process that created the work; (ii) contributed to the design of the algorithm or process; (iii) chose data used by the algorithm for training or otherwise; (iv) caused the AI algorithm or process to be used to yield the work; or (v) engaged in some specific combination of the foregoing activities? Are there other contributions a person could make in a potentially copyrightable AI-generated work in order to be considered an "author"?**

As an initial matter, we note that these questions are not entirely new or unique to AI. While there is no settled body of law plainly on point, courts have considered similar inquiries in the context of software and output of software, and have articulated a possible framework. In *Design Data Corp. v. Unigate Enterprise, Inc.*, for example, the Court of Appeals for the Ninth Circuit considered whether copyright in software might extend to cover the output of that software. 847 F.3d 1169 (9th Cir. 2017). While not deciding the question, the Court considered the possibility it may be appropriate to extend the copyright protections afforded the computer program itself to the program's output where the program performs "the lion's share of the work" in creating the output and the contributions of any user of the software were "marginal." *See id.* at 1173 quoting *Torah Soft Ltd. v. Drosnin*, 136 F. Supp. 2d 276 (S.D.N.Y. 2001). More recently, in *Rearden LLC v. The Walt Disney Company*, the District Court for the Northern District of California dismissed a claim for copyright infringement in the output generated by a computer program that captures facial motion to create special effects in movies as plaintiff had not alleged that its program did the lion's share of the work, or that the contributions of the actors/directors were marginal (nor did the Court think this plausible). 293 F.Supp.3d 963, 970-71 (N.D. Cal. 2018). Thus, if the AI algorithm is embodied within copyrightable code[2] and is largely responsible for a particular output (which would otherwise qualify as a copyrightable expression), then it may be the case that the output is copyright authorship created by the author(s) of the code.

The current law concerning joint ownership of copyrights may provide another potential framework for this analysis. Section 101 of the Copyright Act defines "joint work" as "a work prepared by two or more *authors with the intention that their contributions be merged* into inseparable or interdependent parts of a unitary whole." 17 U.S.C. §101.

Using joint authorship as a framework, however, creates certain difficulties. First, there is the issue discussed above in response to question 1 that an AI machine cannot qualify as an author. Assuming legislation overcomes that issue, there remains a secondary issue, namely the nature of the author's contribution to the joint work. Courts disagree as to whether joint ownership requires each author's contribution to the work at issue to be separately and independently subject to copyright protection, or whether the joint author's contribution must simply be more than a *de minimis* contribution to the work as a whole. *Compare, e.g., Childress v. Taylor*, 945 F.2d 500 (2d Cir. 1991) (contributions and suggestions made to playwright did not create joint authorship works in play) *with Gaiman v. McFarlane*, 360 F.3d 644 (7th Cir. 2004) (writer's creation and description of characters for use in a comic book gave writer joint ownership in the comic book, even though the contribution made was not by itself subject to copyright protection). *See generally* M. Nimmer and D. Nimmer, *Nimmer on Copyright*, Vol. 1, § 6.07[A][3] (2019).

If the obstacles regarding authorship by AI devices and the split among courts

---

[2] An algorithm per se is not copyrightable authorship.

regarding nature of the contribution required for a joint work are resolved, the joint authorship framework could potentially support an analogous claim to copyright ownership by one who authors some portion of the relevant code, or perhaps someone who selects and arranges training material. It seems less likely that simply causing an algorithm to be used could result in copyright ownership.

To the extent amendment is desirable, however, we offer these suggestions regarding which parties would be proper authors and/or joint authors, using the principles outlined in the U.K.'s CDPA as a guideline.

Designers or contributing designers of the AI algorithm or process (collectively, "**Programmers**") (see (i) and (ii) above): Programmers are the parties whose copyrighted work perhaps most directly contributes to the AI device's creation of the work. Designating Programmers as authors of the AI device's output provides practical advantages, namely the available use of work-made-for-hire agreements between the Programmers and their employers so that their employers, who may have invested significant resources in developing the AI device, may reap the economic benefit of its investment. This result on its face appears to advance the Constitutional purposes of the U.S. copyright laws, which is to promote the progress of science and useful arts insofar as it rewards Programmers who create AI devices that produce economically significant works. Please note, however, that extending protections to AI-produced works may disincentivize others in the same field from producing their own works, whether from concerns about copyright infringement or competitive disadvantage. Another potential downside to granting Programmers the copyrights in the output of their AI devices would be that Programmers of the AI device may receive a copyright windfall by obtaining copyrights in both the underlying software and the work product created by their software. The extension of copyright protections to software output is the very issue identified, but not decided, by the Ninth Circuit in *Design Data Corp. v. Unigate Enterprise, Inc.* In the case of AI-generated works the windfall may be particularly problematic insofar as works generated by the AI program may be less directly related to the creative efforts of the Programmer than in the case of traditional software, given the relatively greater autonomy of the operations of AI programs.

Persons who chose data used by the algorithm for training or otherwise (collectively "**AI Trainers**") (see (iii) above). Like Programmers, AI Trainers make decisions regarding the materials to be taught to the AI device which may have a direct impact on the work created by the AI Device. Whether AI Trainers should be considered joint authors of the work created by the AI Device would depend on the nature of the specific contribution made by the AI Trainer to the AI Device, and on whether the AI's contribution is sufficient under the Copyright Act's definition of "joint work" as interpreted by the U.S. courts to qualify as a work of joint authorship. For example, assume the AI Device is programmed to generically replicate the painting of any artist for whom the AI Trainer provides input. The AI Trainer selects the artist, and selects the examples of such artist's works, and the number of artist's works to be submitted to the AI Device for learning purposes. The AI Trainer also instructs the Programmer as to which elements of the artist's work on which the AI Device should focus: e.g., brush technique, lighting, compositional elements. Under this circumstance, the AI Trainer may well qualify as a joint author under applicable law, based on the AI Trainer's selection and arrangement of the materials that the AI used to create the work. By contrast, if the AI Trainer's function requires only a routine aggregation of certain pre-selected data, then the AI Trainer is less likely to qualify as a joint author.

Persons who caused the AI algorithm or process to be used to yield the work (collectively, "**Operators**") (see (iv) above). Operators, without more, do not appear to be authors under even the most expansive definition of the term. The decision to activate the AI device, while necessary to cause the AI device to generate a work, does not involve the creative processes and effort which copyrights are designed to reward.

Persons engaged in some combination of the foregoing activities (see (v) above). Whether persons who engage in one or more of the activities described above would qualify as authors will depend on the specific circumstances surrounding their contribution. An Operator who is also an AI Trainer or Programmer, or both, may qualify as an author of the AI Device work, depending on the nature of that person's contributions, and further provided that the contributions are not *de minimis*. The best current guidelines currently available in this regard are those for "joint authorship" and "joint works" under the Copyright Act and applicable interpretive law. The ultimate decision may depend on whether the view among U.S. courts continues to prevail that a contribution does not become a joint work unless the contribution is independently copyrightable.

Finally, although we cannot rule out the possibility that there may be others involved in creation of the AI device or the works the AI device generates that would qualify as joint authors of the AI-generated work, we would expect that category of others to be limited. It is well established that, in traditional analog copyrighted works, editors of the works do not gain copyrights in the books they edit, and publishers do not become authors of those books simply because they arranged for the physical publication of the books. These contributions, while important to the creation, distribution and success of the work, are not the sort of creative efforts traditional copyright law rewards via the grant of copyrights. Persons working on or with AI devices should similarly not be entitled to recognition as joint authors of the AI-generated work if their contributions are equivalent to those of an editor, publisher, or analogous function which traditionally does not qualify for recognition as authorship.

3. **To the extent an AI algorithm or process learns its function(s) by ingesting large volumes of copyrighted material, does the existing statutory language (e.g., the fair use doctrine) and related case law adequately address the legality of making such use? Should authors be recognized for this type of use of their works? If so, how?**

The use of copyrighted material to "train" AI processes (so-called "machine learning") may well violate the reproduction right of a copyright owner under 17 U.S.C. § 106(1). However, under certain circumstances, some instances of machine learning may constitute non-infringing fair use. The current fair use doctrine, codified at 17 U.S.C. § 107, is flexible and, and is capable of adapting to the use of copyright works in connection with AI without additional modification.

The challenge in applying the fair use doctrine lies in striking a balance between protecting copyright owners from a loss of value in their copyrights, while not placing undue restrictions on AI's ability to learn that will hinder its technological development.

Under the fair use doctrine, 17 U.S.C. § 107, courts weigh factors including the following:

1) the purpose and character of the use, including whether such use is of a commercial nature or is for nonprofit educational purposes; 2) the nature of the copyrighted work; 3) the amount and substantiality of the portion used in relation to the copyrighted work as a whole; and 4) the effect of the use upon the potential market for or value of the copyrighted work.

The first and fourth factors are of particular relevance in considering whether the use of copyrighted materials to train AI constitutes fair use. The first factor concerns the "transformative" nature of the use – specifically, "whether the new work merely supersedes the objects of the original creation . . . , or instead adds something new, with further purpose or different character, altering the first with new expression, meaning or message." *Campbell v. Acuff-Rose Music, Inc.*, 510 U.S. 569, 579

(1994). The fourth factor focuses on "whether unrestricted and widespread conduct of the sort engaged in by defendant would result in a substantially adverse effect on the potential market for the original." *Id.* at 590. The second factor is solely concerned with the nature of the underlying copyrighted work, e.g., whether it is primarily a factual work or a fanciful work, or whether it is published or unpublished, which is unaffected by the involvement of AI. The third factor concerns how much of the underlying work has been copied. For the purposes of machine learning, it seems likely that the work will be copied in its entirety.

The Purpose and Character of the Use

The purpose and character of an AI device's use of copyrighted materials is difficult to generalize, and may change over time. For example, two of the leading "transformative uses" of a work for fair use purposes are parody and criticism. Neither of these uses, however, appear to be available to AI technology at this time. By contrast, current AI technology is now targeted at being able to accomplish only the more "simple" tasks of creating works. *See, e.g.*, http://www.talktotransformer.com (providing an online tool demonstrating how a modern neural network completes text) (last visited Jan. 8, 2020). The capabilities of AI devices are, of course, expected to improve over time, and at some future point AI technology may be able to create parodies or criticism that would qualify for fair use.

AI devices today, however, may use copyrighted works in a non-expressive, transformative manner, which would support a finding of fair use. Courts have recognized a transformative, "non-expressive" fair use where computers have reproduced copyrighted works for the purposes of analyzing or storing those works, as opposed to reproducing them for human consumption. In *Sega Enters. Ltd. v. Accolade, Inc.*, 977 F.2d 1510 (9th Cir. 1992), the Ninth Circuit held that copying of defendant's video game software was non-infringing fair use because it was incidental to a non-expressive purpose – namely, copying was necessary to access the non-copyrightable functional elements of the video games, which were necessary to make games compatible with plaintiff's game consoles.

In *Kelly v. Arriba Soft Corp.*, 336 F.3d 811 (9th Cir. 2003) and *Perfect 10 v. Amazon.com, Inc.*, 508 F.3d 1146 (9th Cir. 2007), the defendants operated internet search engines and utilized "crawlers," software that traversed the internet, indexing web pages and downloading full size images to the defendant's servers, where defendant's software reduced the images to thumbnails and displaying them on its search engines. In each case the court held that, notwithstanding that defendants had reproduced expressive aspects of the copyrighted images, the use was transformative because the reproduced images had been assembled as "tools" to index and improve access to the images, and not as vehicles for conveying expression. To the extent that AI is using copyrighted materials for a purpose unrelated to their expressive elements, such "non-expressive" use may be found to be transformative. As in the *Kelly* and *Perfect 10* cases, these "non-expressive" uses are less likely to implicate the fourth factor, since non-expressive use (such as sorting and indexing) is more likely to enhance the value of a work rather than harm it.

Machine learning involving creative works gives AI the ability to derive information about the ways that authors express ideas, and may derive value from these expressive aspects of prior works. The transformative nature of such uses may be less apparent than with non-expressive uses. For example, an AI process may ingest hundreds of paintings by a single artist and output its own image based on all the data it has been provided concerning the artist's means of expression, *e.g.*, subjects, brush strokes, coloring, and shading. A court may have to consider whether, through the AI's processing and output of a new work, it has "added something new, with further purpose or different character, altering the first with new

expression, meaning or usage," or whether it has simply regurgitated different elements from the artist's paintings. In instances of "expressive use," where AI is analyzing copyrighted works to obtain data on how authors expressed themselves, for the purpose of using that data to produce its own works, courts may look to the fourth fair use factor to determine the potential for such uses to harm the potential market for the copyrighted work.

<u>The Effect of the Use Upon the Potential Market for or Value of the Copyrighted Work</u>

The fourth fair use factor evaluates a potentially infringing use by considering its adverse effect on the potential market for the copyrighted work. This inquiry is critical in the AI context because AI has the potential to replace human authors as a cheaper, more efficient option. For example, a company uses AI to compose music based on certain parameters set by its users (*e.g.*, mood, genre and instrumentation), and licenses the resulting songs. The company trains its AI by inputting a wide variety of musical works, many of them copyrighted. Applying the fourth factor, it is certainly likely that the AI-generated recordings will harm the market for traditionally recorded music. Moreover, courts applying the fourth factor will look beyond the defendant's use to "whether unrestricted and widespread conduct of the sort engaged in by the defendant (whether in fact engaged in by defendant or by others) would result in a substantially adverse impact on the potential market for, or a value of, the plaintiff's present work." *Campbell*, 510 U.S. at 590. If the use of copyrighted recordings as input data for AI composers becomes a widespread practice, it will almost certainly result in the displacement of human composers.

In its current form, the fair use doctrine is flexible and readily applicable to acts of infringement arising from machine learning. The question remains as to who may be liable for such acts of infringement, which is discussed in the section below.

4. Are current laws for assigning liability for copyright infringement adequate to address a situation in which an AI process creates a work that infringes a copyrighted work?

For the time being, secondary liability may suffice to protect copyright owners from infringement caused by AI. That said, enforcement under theories of contributory and vicarious infringement will require the courts to consider novel issues regarding, among other things, agent, control, and foreseeability of the AI device's acts. And, as AI becomes increasingly autonomous, changes to the law may prove necessary.

Although an AI process is capable of infringing through reproduction or creation of a derivative version of a copyrighted work, a copyright owner cannot hold the AI process itself liable, as it is not a natural or juristic person capable of being sued. Instead, the copyright owner must establish secondary liability by some human or corporate agent or affiliate. Copyright law recognizes two forms of secondary liability: contributory infringement liability, and vicarious infringement liability. A party is liable for contributory infringement of a copyright where he (1) had knowledge of a direct infringement; and (2) induced, caused or materially contributed to that infringement. *Matthew Bender & Co. v. W. Pub. Co.*, 158 F.3d 693, 706 (2d Cir. 1998). A party is vicariously liable for copyright infringement if he or she (1) has the right and ability to control the infringing activity and (2) has a direct financial interest in such activity. *A&M Records, Inc. v. Napster, Inc.*, 239 F.3d 1004, 1022 (9th Cir. 2001).

To obtain redress for infringement of his copyrights by an AI process, a copyright owner must establish secondary liability of a party with a connection to the AI's actions – for example, under a contributory liability theory, the person(s) responsible for programming the AI, *e.g.*, a software designer, and the person(s) responsible for inputting

material to "train" the AI and their employers might be liable. Under a vicarious liability theory, the entity at whose direction and expense the AI has been programmed and/or "trained" may be liable. However, application of contributory infringement or vicarious liability claims may be more factually challenging as AI technology evolves. One of the primary features of AI is its autonomy, which leads to unpredictability and unforeseeable actions that AI may take when put into operation. As an AI process receives more autonomy, fewer parties may have the ability to control it or remain aware of what actions it is taking. Liability standards founded on agency, control and foreseeability become increasingly tenuous as the AI's behavior becomes increasingly distanced from its initial programming.

When AI commits copyright infringement, there must be a responsible party; otherwise the purposes of U.S. copyright law would be undermined. If designers and programmers and the companies that employ them are absolved from liability because AI's actions are unforeseeable and uncontrolled, then copyright owners will be left without a remedy for AI infringement of their copyrights. This result will encourage copyright infringement on a massive scale and harm the potential market for copyrighted works generally. However, as AI becomes increasingly autonomous, unless courts can keep pace with developing reliable standards for addressing issues concerning agency, control and foreseeability, changes in the law may become necessary. If the U.S. Copyright Act were amended to designate an AI machine as an "employee" for purposes of the U.S. Copyright Act's work-made-for-hire provisions, then it would be logical and natural to extend liability for infringement by that machine to its owner.

5. **Should an entity or entities other than a natural person, or company to which a natural person assigns a copyrighted work, be able to own the copyright on the AI work? For example: Should a company who trains the artificial intelligence process that creates the work be able to be an owner?**

We refer you to our discussion in response to Question 2, in which we address this issue.

***

The NYIPLA appreciates the opportunity to provide feedback to the USPTO on these important intellectual property law and policy issues and welcomes the opportunity to further participate in discussions regarding the effect of developments in innovation on the law and the need for changes to address those effects.