

DEPARTMENT OF COMMERCE

Patent and Trademark Office [Docket No. PTO-C-2019-0038]

Comments on Intellectual Property Protection for Artificial Intelligence Innovation

AGENCY: United States Patent and Trademark Office, Department of Commerce.

Attention: Coke Stewart, Office of the Under Secretary and Director of the USPTO

About Obeebo, Inc.

Obeebo, Inc. (“Obeebo” or “we”) is a California based company developing AI technology for the composition and production of music. Obeebo’s technology can rapidly compose music in a range of styles, either “from scratch” or as variations of existing works.

Contact: info@obeebo.ai

...

Obeebo’s comments to the USPTO Questions

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?

A work that would be copyright-eligible if produced by a natural person should still be copyright-eligible if produced by an AI. It is the nature of the work and not the nature of the author that should determine copyright-eligibility, and a work that would otherwise be conventionally copyright-eligible should not have its copyright-eligibility negated based on the nature of its author.

If works created by an AI could not be copyrighted, then works created by an AI could be freely copied by third parties. This would be bad for businesses that are investing in AI, and bad for the burgeoning field of AI in general. IP protection mechanisms such as copyright are intended, at least in part, to motivate investment in creativity and innovation. Exempting AI-created works from copyright protection would create a disincentive to advance certain AI technologies, and IP protection systems should never create a disincentive to advance a technology.

The key question isn’t “should work produced by an AI qualify as a protectable work,” but rather “when an AI creates copyright-eligible work, who should own the copyright?” It simply does not make sense for an AI itself to own copyright because an AI cannot realize the benefits and responsibilities of being a copyright owner. Furthermore, though an AI may appear to be an author of a creative work we posit that it is more accurate and appropriate to regard an AI as a tool through which a work is created. All copyright-eligible works are generated using tools, such as a paintbrush, a piano, a typewriter, etc. Just as the user of a paintbrush, piano, typewriter, or other tool owns the copyright in what they produce with such tool, the user of an AI should own the copyright in the output of the AI. Now, because an AI can sometimes operate fairly autonomously, the concept of the user of an AI requires consideration. AI typically exists in the form of copyable and distributable software and therefore multiple instances of the same or similar AI are possible. The user of an AI is not necessarily the original creator of the AI or the owner of any copyright in the code that makes up the AI. For the purposes of copyright ownership, we posit that the user of an AI (and therefore the owner of any copyright in any work produced by the AI) should be defined as the owner or licensee of the specific instance of the AI by which the work was generated (“User”). This is analogous to a typical

employer/employee relationship (with the User being the employer and the AI being the employee), wherein an employer typically owns IP created by its employee as “work for hire.”

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”?

Firstly, we disagree with the premise of this question and repeat (from our response to question 1) that involvement by a natural person should not be required for an AI-created work to be protectable under copyright. Any AI will necessarily have natural persons associated therewith, such as (a) IP owners (of the AI algorithm and/or associated software implementation thereof) and (possibly) (b) licensed users. Copyright in the output of the AI should follow the rights bestowed by IP ownership and licenses. For example, suppose a company or individual X owns the AI algorithm/software but has not licensed it to others. Then the company would own the copyright to the output of the AI. However, if the company X that owns the AI algorithm has licensed use of an instance thereof to another, then the licensee should assume the copyright in the output of the licensed instance of the AI. This is analogous to a company owning the IP in, for example, a word processor the company provides as a product, but the licensee would own the rights to any works the licensee creates using the word processor.

However, if we do assume the premise of the question, then the question is what kind of involvement by a natural person is required for the output of the AI to be deemed copyright protectable. As we contend that AIs can generate copyright-eligible material with or without assistance by a natural person, we must take a very permissive stance on this question. Item by item, if a natural person:

- (i) designed the AI algorithm or process that created the work;
 - Yes the work is copyright eligible and the owner of the copyright is the licensee of the AI, or the owner of the AI if no licensee exists.
- (ii) contributed to the design of the algorithm or process;
 - Yes the work is copyright eligible and the owner of the copyright is the licensee of the AI, or the owner of the AI if no licensee exists
- (iii) chose data used by the algorithm for training or otherwise;
 - Yes the work is copyright eligible and the owner of the copyright is the licensee of the AI, or the owner of the AI if no licensee exists
- (iv) caused the AI algorithm or process to be used to yield the work; or
 - Yes the work is copyright eligible and the owner of the copyright is the licensee of the AI, or the owner of the AI if no licensee exists
- (v) engaged in some specific combination of the foregoing
 - Yes the work is copyright eligible and the owner of the copyright is the licensee of the AI, or the owner of the AI if no licensee exists

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?

All human creators/authors are influenced by the (copyrighted) works of those that came before. Provided that a new work is sufficiently distinctive from earlier works, no recognition is currently afforded to the influences of a human

creator/author. Eric Clapton is under no obligation to acknowledge the influence of Robert Johnson unless he is performing a Robert Johnson song. There is no “volume threshold” when it comes to the influences of a human creator/author and there is no basis or need to introduce such a threshold if the creator/author is an AI. As long as a new work created by an AI is sufficiently distinctive from earlier works (using the same bar that would be applied to a new work created by a human), then the sheer volume processing capability of the AI should not trigger any extra recognition of its influences.

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?

In our responses above we argue that the User of an AI (i.e., the owner or licensee of the instance of the AI) should be the owner and copyright holder for any work produced by an AI. Continuing this reasoning, it follows that the User of an AI should be responsible/liable for the output of the AI. If the output of an AI infringes a 3rd party’s copyright, then the User of the AI should be liable, especially since it is presumably the User and not the AI that would enjoy any benefit from the infringing work. A clear analogy here is this: if the output of a typewriter infringes a 3rd party’s copyright then it is the author using the typewriter that is liable, not the typewriter itself or the manufacturer of the typewriter.

It follows that an AI itself cannot commit willful infringement of any form of IP protection; however, a User can be guilty of using an AI to commit willful infringement.

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?

Any User (i.e., owner/licensee) of an instance of an AI should be able to own copyright in works produced by the AI, provided such is consistent with the terms (e.g., licensing terms) under which the User engages with the AI. For example, a company exploiting an open source AI algorithm should not be able to own copyright in the output of the AI algorithm if such would be inconsistent with the terms of the open source license governing use by the company of the AI algorithm. In practice, works created by the AI should be treated as “works for hire” commissioned by the User and the term of any copyright should be set accordingly (e.g., the shorter of 120 years after creation or 95 years after publication, in the US).

6. Are there other copyright issues that need to be addressed to promote the goals of copyright law in connection with the use of AI?

We feel the simple rules described above readily accommodate existing copyright laws and see no outstanding issues to address.

7. Would the use of AI in trademark searching impact the registrability of trademarks? If so, how?

The use of AI in trademark (or patent) searching will lead to more sophisticated searching and more refined search results. This will improve the identification of cases of potential overlap/confusion (or issues of patentability). AI-based searching will raise the bar for trademark distinctiveness (or patentability), making it harder for applicants to get trademarks (or patents) registered but ultimately improving the quality of the trademark (or patent) register. The number of frivolous filings will likely be reduced, as will the amount of time that human Examiners need to devote to searching.

8. How, if at all, does AI impact trademark law? Is the existing statutory language in the Lanham Act adequate to address the use of AI in the marketplace?

Treating an AI as a tool used by a User and the output of an AI as a “work for hire,” as detailed in the responses above, should generally fall in line with existing trademark law.

9. How, if at all, does AI impact the need to protect databases and data sets? Are existing laws adequate to protect such data?

We do not see any new issues introduced by AI here. Protecting data is extremely important in our world today but AI itself should not necessitate any new laws around the protection of data.

10. How, if at all, does AI impact trade secret law? Is the Defend Trade Secrets Act (DTSA), 18 U.S.C. 1836 et seq., adequate to address the use of AI in the marketplace?

Treating an AI as a tool used by a User and the output of an AI as a “work for hire,” as detailed in the responses above, should generally fall in line with existing trade secret law.

11. Do any laws, policies, or practices need to change in order to ensure an appropriate balance between maintaining trade secrets on the one hand and obtaining patents, copyrights, or other forms of intellectual property protection related to AI on the other?

We do not see any changes needed here. There is currently no requirement to indicate that an AI was involved in a patent/copyright/other protection unless the AI itself is part of the subject of such patent/copyright/other protection (e.g., as would be the case in a patent on the AI algorithm or a copyright on the AI code), in which case there’s no new issue. In the case where an AI is used to generate a work that is protectable and owned by a User as described in the responses above, we do not see any reason to require that the role of the AI be revealed or declared to the extent that a trade secret might be compromised.

12. Are there any other AI-related issues pertinent to intellectual property rights (other than those related to patent rights) that the USPTO should examine?

The biggest issue we see that has not been addressed above is the potential for AI to rapidly generate huge volumes of IP. For example, if AIs became very good at preparing and filing patent/TM applications, a User could deploy an AI to churn out a huge number of filings in a very short period of time and pay only the bare filing fees to initiate protection. This could overwhelm Examiners and introduce serious hurdles for competitors. Similarly, in the case of copyright it is conceivable that an AI could rapidly produce all permutations in a given space (e.g., all melodies having N notes or all images having X pixels) and effectively secure copyright thereof without ever pursuing the vast majority of such permutations in any meaningful way. This sort of exploitation of IP protection systems by AI cannot be permitted, but it is not obvious how to prevent it because there may be no clear mechanism to distinguish works produced by AIs from works produced by natural persons. Our best proposals are: i) for patents/TM, introduce some sort of volume controls in IP protection systems (like an upper limit on the number of patent/TM applications that can be filed by a given applicant / customer number in a certain period of time, or tiered filing fees that increase with volume); and ii) for copyright, introduce a meaningful use requirement that limits enforceable protection to works that are genuinely pursued or exploited within some reasonable time from creation.

13. Are there any relevant policies or practices from intellectual property agencies or legal systems in other countries that may help inform USPTO’s policies and practices regarding intellectual property rights (other than those related to patent rights)?

Our experience indicates that policies relevant to the intersection of IP + AI are rapidly evolving in many countries. These are difficult issues to address and we commend the USPTO for engaging with the community while these issues are under consideration. The most relevant learning from other legal systems with which we are familiar is the importance of clarity. Europe, for example, has been grappling with a number of these issues but the results have

only fostered uncertainty and confusion. We view this as a cautionary tale and hope the USPTO can adopt clearer, simpler solutions.

info@obeebo.ai