Computer & Communications
Industry Association
Tech Advocacy Since 1972

In re: Artificial intelligence and intellectual property: call for views by the United Kingdom Intellectual Property Office

COMMENTS OF
COMPUTER & COMMUNICATIONS INDUSTRY ASSOCIATION

The Computer & Communications Industry Association (CCIA)[1] submits the following comments in response to selected questions from the United Kingdom Intellectual Property Organization's call for views on artificial intelligence (AI) and intellectual property.[2]

CCIA's members are leaders in AI innovation. Most modern AI technology utilizes hardware manufactured by CCIA member companies, and much of it relies on AI technology released as open source contributions by CCIA members. Because of this, CCIA members have a significant interest in ensuring that the various types of AI-related inventions are promoted, rather than suppressed, by the intellectual property system.

I. Call for Responses: Patents

A. Question 1: What role can/does the patent system play in encouraging the development and use of AI technologies?

Much like any area of technology, the patent system has a role to play in encouraging the development and use of AI technologies. And like any area of technology, a patent system that grants overly broad patents on AI, or grants patents on AI that are not necessary to generate the patented innovation, can slow or stall innovation in AI technologies.

A properly calibrated balance is necessary. In the case of AI inventions, that means treating inventions created by humans regarding AI technology, or inventions created by humans employing AI tools, similarly to inventions in other fields. In the case of autonomously generated AI inventions, that balance likely bars any reward of a patent.

B. Question 2: Can current AI systems devise inventions?

1. To what extent is AI a tool for human inventors to use?

2. Could the AI developer, the user of the AI, or the person who constructs the datasets on which AI is trained, claim inventorship?

3. Are there situations when a human inventor cannot be identified?

Currently, CCIA does not believe that AI systems can autonomously generate inventions. Instead, they act as tools for human operators to use.

---

[1] CCIA is an international nonprofit membership organization representing companies in the computer, Internet, information technology, and telecommunications industries. Together, CCIA's members employ nearly one million workers and generate approximately a quarter of a trillion dollars in annual revenue. CCIA promotes open markets, open systems, open networks, and full, fair, and open competition in the computer, telecommunications, and Internet industries. A complete list of CCIA members is available at http://www.ccianet.org/members.

[2] Available at https://www.gov.uk/government/consultations/artificial-intelligence-and-intellectual-property-call-for-views.

25 Massachusetts Avenue NW #300C
Washington,DC 20001
+1(202) 783-0070
https://www.ccianet.org

Rue de la Loi 227, First Floor
1040 Brussels, Belgium
+32-2-888-8462

The AI developer and the person who constructs the training dataset cannot claim inventorship in the AI's output, any more than the developer who writes electronic design automation (EDA) software can claim intellectual property in the output of that software without their own involvement in that output or the inventor of the Bunsen burner could claim patents in the chemicals developed using burners in the laboratory.

The user of the AI may be able to claim inventorship, to the extent that they independently contribute something of novelty, with inventive step, to the output of the AI. To the extent the user simply files a patent on the output of the AI without further involvement, that simply represents an attempt to claim an invention anyone of ordinary skill could have produced.

In circumstances in which no human contributed any novelty or inventive step to the application, then there is no proper inventor. This is a desirable situation—a patent is a government-granted monopoly and should be disfavored where that monopoly is not required to obtain the idea. An AI tool, not motivated by the economic reward of a patent, does not justify the economic harm of restraining others from using its output simply because one individual was the first to do the mechanical act of applying for a patent on its output.

C. Question 3: Should patent law allow AI to be identified as the sole or joint inventor?

CCIA opposes allowing an AI to be named as an inventor. The law does not presently permit such a situation, as evidenced by the European Patent Office's recent decision to refuse a patent application filed naming an AI as an inventor. Nor is there any policy justification for naming an AI as an inventor. However, neither should a human be named as an inventor of technology created by an AI—in such circumstances, there is no inventor who could be properly named on the patent.

This does not mean that there may not be a need to differentiate between AI-contributed and human-contributed aspect of an invention when a human utilizes autonomous AI as part of the process used by the human to create an invention—in such circumstances, it would be appropriate to name the human as the inventor and owner under the ordinary course of patent law. That ownership would be solely with respect to the contribution the human made to the invention—the portion contributed autonomously by AI would not itself be patentable, as it would represent nothing more than what is available to any artisan of ordinary skill using the available prior art. However, that human contribution must itself provide novelty and inventive step over the prior art and the skill in the art—including the output of autonomous AIs.

To the extent that technology reaches a point at which it would be correct to say that an AI is self-motivated and contains sufficient indicia of consciousness to justify treatment as a person, such a change would implicate issues far beyond the expertise and scope of patent law and should be addressed at that time.

D. Question 4: If AI cannot be credited as inventor, will this discourage future inventions being protected by patents? Would this impact on innovation developed using AI? Would there be an impact if inventions were kept confidential rather than made public through the patent system?

The inability to credit AI as an inventor will not discourage future inventions from being protected by patents, nor would it have a negative impact on innovation developed using AI. In circumstances in which a human utilizes AI as a tool and independently provides inventive step, the human may be named as an inventor and the issue does not arise.

The sole circumstance in which a human should not be named as an inventor and where not naming an AI as inventor might block patents is when an AI solely and autonomously innovates. While CCIA does not believe that situation has occurred at present, in that circumstance, there is no reason to issue a patent.

The AI is not motivated by the prospect of a reward of exclusivity—it is not motivated at all. It would produce the idea upon being run regardless of the availability of a patent, suggesting that the

essential value of a patent as an "encouragement to [] ingenious discoveries"[3] is absent and suggesting that the "embarrasment [sic] of an exclusive patent"[4] is unnecessary to cause the idea to be created. In addition, the output of such a machine should never qualify as having inventive step for the reasons described below in response to Questions 13 and 14, *infra*, suggesting there is no circumstance in which such an AI generator of an idea would qualify as the inventor of a patent. The creation of an idea-creating machine would still be incentivized, as it would be patentable (unless it was itself the output of an AI), but there is no need for the economic incentive of a patent to sufficiently incentivize operation of such a machine.

Other economic incentives, such as first-mover advantage, network effects, and competitive necessity will be sufficient to justify any resources expended in operating such a machine. And there is no need to incentivize disclosure by the operator of such a machine as any other operator of a similar machine would be able to obtain the same invention, suggesting that the disclosure-based justification for the patent bargain is weak at best with respect to autonomous AI inventions.

Given the lack of necessity and the potential negative impacts of permitting the owner of an autonomous invention-generating AI to own any idea that comes from it, effectively preempting a number of fields without any of its own contribution, innovation is likely to be positively served by refusing to credit AI as an inventor.

E. Question 5: Is there a moral case for recognising AI as an inventor in a patent?

Moral rights are justified based on the concept of recognizing the dignity and personhood of a creator. An AI, being a non-sentient computer program, lacks personhood and dignity.

Thus, there are no moral rights considerations, as the only potential holder of such a right—the AI itself—lacks any sense that would justify a moral right in the invention. The AI is not motivated by the prospect of a reward of recognition—it is not motivated at all. It would produce the idea upon being run regardless of the availability of a patent or its recognition as an inventor, suggesting that any moral right to recognition for an AI is inappropriate and unnecessary.

As above, to the extent that technology reaches a point at which it would be correct to say that an AI is self-motivated and contains sufficient indicia of consciousness to justify treatment as a person and concomitant moral rights, such a change would implicate issues far beyond the expertise and scope of patent law and should be addressed at that time.

F. Question 6: If AI was named as sole or joint inventor of a patented invention, who or what should be entitled to own the patent?

While CCIA disagrees that AIs should be named as inventors on a patent, in the event they were, the invention should belong to no one.

This does not mean that there may not be a need to differentiate between AI-contributed and human-contributed aspect of an invention when a human utilizes autonomous AI as part of the process used by the human to create an invention—in such circumstances, it would be appropriate to name the human as the inventor and owner under the ordinary course of patent law. That ownership would be solely with respect to the contribution the human made to the invention—the portion contributed autonomously by AI would not itself be patentable, as it would represent nothing more than what is available to any artisan of ordinary skill using the available prior art. However, that human contribution

[3] Letter from James Madison to Thomas Jefferson (Oct. 17, 1788), *available at* https://founders.archives.gov/documents/Madison/01-11-02-0218.
[4] Letter from Thomas Jefferson to Isaac McPherson (Aug. 13, 1813), *available at* https://founders.archives.gov/documents/Jefferson/03-06-02-0322.

must itself provide novelty and inventive step over the prior art and the skill in the art—including the output of autonomous AIs.

G. Question 7: Does current law or practice cause problems for the grant of patents for AI inventions in the UK?

Current law and practice are sufficient for protecting AI inventions in the UK. CCIA members regularly obtain AI-related patents without issue.

However, there is a potential that the current practice may over-grant AI patents. Much like software before it, AI presents the potential for unique issues with obtaining and searching prior art. While standard AI techniques are more likely to be described in the literature than is the case in software, there is still a significant proportion of AI technology that is undocumented except in source code. This source code may or may not be available and is generally considered difficult to search for. USPTO Director Kappos spoke about this issue with respect to software in 2012, noting difficulties with respect to "software, where much prior art is in the form of previously written software, which is difficult to find and more difficult to understand unless you wrote it."[5] An FTC report received similar comments from stakeholders, arguing that time constraints do not allow adequate searches for software prior art.[6] AI is likely to present the same search issues.

This is relevant because examiners are less likely to identify and cite sources of prior art that are more difficult to access and search. Consistent with Director Kappos's comments, a recent GAO study of U.S. examiner behavior, based in part on examiner interviews, found that software-related non-patent literature was the prior art source examiners were most likely to search "rarely" or "never."[7] Examination in AI technology is likely to proceed similarly, with AI-related non-patent literature infrequently searched or cited.

Searching patent prior art for AI inventions is likely to involve difficulties similar to those in the software space. Software patents have traditionally been classified according to the end use of the software, meaning that software for controlling the temperature of a pizza oven and software for controlling the temperature of a kiln would be classified differently even though the underlying software concepts are identical.[8] AI patents may be treated similarly. As a result, classification-based examiner search techniques are more difficult to employ with respect to AI and often miss relevant prior art.

The combination of these factors means that documentation of AI prior art may be difficult or impossible to obtain, leading examiners to be more likely to issue patents on AI inventions that are not actually novel or that lack inventive step.

H. Question 8: Could there be patentability issues in the future as AI technology develops?

CCIA does not foresee significant new issues, outside of those discussed in these questions, arising from the current trends in AI technology.

---

[5] USPTO Director David Kappos, *An Examination of Software Patents* (Nov. 20, 2012), *available at* https://www.uspto.gov/about-us/news-updates/examination-software-patents.

[6] *See* FTC, *To Promote Innovation: The Proper Balance of Competition and Patent Law and Policy* at 10 n. 35 (Oct. 2003), *available at* https://www.ftc.gov/sites/default/files/documents/reports/promote-innovation-proper-balance-competition-and-patent-law-and-policy/innovationrpt.pdf.

[7] *See* GAO, *Patent Office Has Opportunities to Further Improve Application Review and Patent Quality* 4-5 (Sept. 16, 2016), *available at* https://www.gao.gov/assets/680/679830.pdf; *cf.* Cohen, *Reverse Engineering and the Rise of Electronic Vigilantism: Intellectual Property Implications of "Lock-Out" Programs*, 68 S. Cal. L. Rev. 1091, 1178 (1995) ("In an area that relies so heavily on published, 'official' prior art, a rejection based on 'common industry knowledge' that does not appear in the scholarly literature is unlikely.").

[8] *See* Lemley et al., *Software and Internet Law* 332 (2000).

I. Question 9: How difficult is it to secure patent protection for AI inventions because of the list of excluded categories in UK law? Where should the line be drawn here to best stimulate AI innovation?

CCIA members have not experienced difficulty securing appropriate patent protection based on the list of excluded categories in UK law. Generally speaking, AI inventions should be treated similarly to other software inventions.

J. Question 11: Does the requirement for a patent to provide enough detail to allow a skilled person to perform an invention pose problems for AI inventions?

1. Could there be a need to disclose more than a basic trained AI model, for example training data or the coefficient or weight of the model? If yes, is it clear how much information would be sufficient for a skilled person to be able to work the invention? Are special provisions needed for this information to be filed and stored?

Sufficient disclosure for an AI invention is dependent on how AI is integrated into the invention and the type of AI in question. Much of current AI is based on machine learning techniques, and these comments focus on the disclosures that are most relevant to machine learning. However, other forms of AI exist and will likely be developed, and distinct disclosure considerations may apply to other AI architectures.

Inventions that relate to improving the functionality of AI inventions (AI-improving inventions), such as new algorithms for machine learning or training, will generally implicate the same disclosure concerns as computer-implemented inventions. The written description must describe not just the desired aspect of AI, but how it is achieved—typically, by disclosing the algorithm or implementing structure and how to utilize it. For example, a patent application for a new deep learning structure could describe the number of layers, the number of units per layer, data flows between units or layers, and activation functions. A patent application for a new training algorithm would likely describe the structure to be trained, the starting point and data used in training, and other relevant information necessary to reproduce the invention.

In general, CCIA does not believe a system of deposit or other special storage would prove useful. The algorithms, training methodology, and model details useful to enabling and describing existing AI tools are generally amenable to textual or digital description and should be submitted as part of the patent application. Any additional data necessary for enabling and disclosing an invention that is more voluminous than is appropriate for inclusion in the patent document could be submitted as an appendix to the application, similar to the permitted appendices for computer program listings.

To the extent that disclosure of future AI architectures would be well-suited for a depository system, such a system could be created at that time.

2. What would be the effect if competitors could use data to quickly train a different AI model?

Assuming the different AI model is outside of the scope of the patent, this primarily impacts the ease of designing around the invention. Typically, design-around is viewed as a beneficial outcome. There may be other issues relating to data, such as proprietary rights to the data, that could complicate this concern.

3. How would the skilled person know whether the invention could be repeated across the breadth of the patent claims or whether a claimed result could be achieved?

AI systems can be unpredictable. Seemingly minor changes in training data or algorithm can produce wildly varying output models.

Given this, it is crucial—particularly for AI-application inventions—that the specification enable the full scope of the claim. This means that the disclosure of a single species would not generally be sufficient to support a generic claim—*e.g.*, disclosure of a deep learning model that recognizes images of dogs would not be sufficient to support a claim to "using a deep learning model to recognize images by type" without additional support in the specification to provide modeling for generic image recognition. Claims to "do a function with AI" should only be considered enabled if they teach how to do the function with respect to any form of AI, not just a particular model of AI. To treat such claims otherwise would effectively, and unjustifiably, imply that the patent could predictably enable the use of novel AI architectures to perform the claimed function.

K. Question 13: Does or will AI challenge the level of inventive step required to obtain a patent? If yes, can this challenge be accommodated by current patent law?

CCIA submits that the law should exclude from patent protection inventions that are generated autonomously by an AI application. An invention autonomously generated by an AI application is inherently lacking in inventive step.

An invention lacks inventive step when it would have been obvious to a person skilled in the art. Once an AI is made available, it becomes part of the ordinary skill in the art. While that inventing AI might itself be patentable, its output would not be patentable absent additional human intervention that goes beyond the ordinary skill. Ordinary creativity—the output of a skilled person—exceeds the creativity of an 'automaton'. But an automaton is exactly what an AI is. In the event that an AI is capable of the needed creativity to create an idea, it is definitionally within the creativity of a person skilled in the art.

Much like the availability of computation and computer-aided design tools has affected what is reasonable to treat as ordinary skill, the availability of AI tools will affect what is reasonable to treat as the ordinary skill in the art. An ordinary artisan, relying on the output of an AI tool, has not created anything beyond the ordinary skill unless they contribute something to the combination that rises above what any ordinary artisan could do with the same AI tool. We would not find patentable the set of logic gates created on a field-programmable gate array by a VHDL compiler. It is simply the output of an ordinary tool—though the tool itself might contain patentable inventions. Similarly, the output of an AI is within the scope of ordinary creativity and presumptively unpatentable, even though the AI itself might be patentable.

To the extent that technology reaches a point at which it would be correct to say that an AI is self-motivated and contains sufficient indicia of consciousness to justify treatment as a person, such a change would implicate issues far beyond the expertise and scope of patent law and should be addressed at that time.

The only necessity to accommodate this within current patent law is to explicitly define ordinary skill in the art as including the ability to use AI invention assistance tools. Once that definition is made, the existing definition of "person skilled in the art" incorporates the skill provided by AI.

L. Question 14: Should we extend the concept of "the person skilled in the art" to "the machine trained in the art"?

While CCIA submits that this is unnecessary, as a "person skilled in the art" inherently includes that person's utilization of AI assistance in the practice of their craft, this would be a viable alternative route to achieving the same outcome.

II. Call for Responses: Copyright

A. Question 1: Do you agree with the above description of how AI may use copyright works and databases, when infringement takes place and which exceptions apply? Are there other technical and legal aspects that need to be considered?

B. Question 2: Is there a need for greater clarity about who is liable when an AI infringes copyright?

C. Question 3: Is there a need to clarify existing exceptions, to create new ones, or to promote licensing, in order to support the use of copyright works by AI systems? Please provide any evidence to justify this.

D. Question 4: Is there a need to provide additional protection for copyright or database owners whose works are used by AI systems? Please provide any evidence to justify this.

The use of the data subsisting in copyright works without authorization for machine learning should not constitute an infringement of copyright. Similarly, the database right should not preclude the use of data from a database for the purpose of machine learning. However, the existing copyright exception for text and data-mining does not apply to commercial uses. Likewise, the database right provides for an exception only for noncommercial research. The exclusion of commercial uses from the scope of these exceptions will frustrate the government's objective of making "the UK a global centre for AI and data-driven innovation."

In contrast to the UK law, in the United States, the existing statutory framework and related case law concerning the fair use right, 17 U.S.C. § 107, clearly permit the ingestion of large amounts of copyrightable material for the purpose of an AI algorithm or process learning its function for commercial as well as noncommercial purposes. Moreover, there is no database right in the United States to interfere with commercial uses of nonoriginal databases. To highlight the need to broaden the existing exceptions in the UK, we will explain in detail how fair use permits the ingestion of data in the United States.

AI algorithms and other processes often require the ingestion of large amounts of material. Assembling that material may entail converting it into a more usable format, e.g., translating image files into machine-readable files.[9] In addition, backup copies of the materials will be necessary to protect against loss of data in the event of system failure.[10] Temporary reproductions of portions of the material in a computer's random access memory are a normal part of the process of training and AI algorithm.[11] All these copies are not viewable or consumable by the outside world.[12] Because these non-expressive copies are not consumable by the public, they do not function as market substitutes for copies of the

[9] In the Google Library Project, Google made a digital scan of each book it borrowed from a research library, then used optical character recognition software to convert the scan machine readable text. Google retained both the scanned image and machine-readable text. *Authors Guild v. Google, Inc.*, 804 F.3d 202, 208 (2d Cir. 2015).

[10] For example, HathiTrust created and maintained four text-only copies of its entire database (one on the primary server at the University of Michigan, another at the mirror server at the University of Indiana, and two encrypted backup tapes at two secure locations on the University of Michigan campus) for the purpose of balancing the load of user web traffic and serving as back-up in the case of a disaster. *Authors Guild v. HathiTrust*, 755 F.3d 87 (2d Cir. 2014).

[11] These temporary reproductions may not constitute copies under the Copyright Act. *See Cartoon Networks LP v. CSC Holdings, Inc.*, 536 F.3d 121 (2d Cir. 2008).

[12] Professor Matthew Sag characterizes acts of copying which do not communicate the author's original expression to the public as "nonexpressive uses." *See, e.g.*, Matthew Sag, *Copyright and Copy-reliant Technology*, 103 NW. U. L. REV. 1607, 1624 (2009). Professor Edward Lee describes three kinds of uses: creational uses (uses of copyrighted works to create a technology); operational uses (uses that occur during the operation of the technology once it has been created); and output uses (the distribution or display of works as an output of the technology). Edward Lee, *Technological Fair Use*, 83 SO. CAL. L. REV. 797, 842-44 (2010).

ingested works.[13] Numerous U.S. appellate courts have correctly found the mass copying of raw material to build databases for uses by AI processes to be fair use under 17 U.S.C. § 107. *See, e.g.*, *Authors Guild v. Google, Inc.*, 804 F.3d 202 (2d Cir. 2015); *Authors Guild v. HathiTrust*, 755 F.3d 87 (2d Cir. 2014); *A.V. ex rel. Vanderhye v. iParadigms, LLC*, 562 F.3d 630, 640 (4th Cir. 2009); *Perfect 10 v. Amazon.com, Inc.*, 508 F.3d 1146, 1165 (9th Cir. 2007); *Kelly v. Arriba Soft Corp.*, 336 F.3d 811, 818 (9th Cir. 2003). Judge Leval's opinion in Google provides the clearest analysis of why the creation of an AI database, and its subsequent uses, are fair uses.

1. The Purpose and Character of the Use

The Google court stated that it "had no difficulty concluding that Google's making of a digital copy of Plaintiff's books for the purpose of enabling a search for identification of books containing a term of interest to the searcher involves a highly transformative purpose. . . ." *Google*, 804 F.3d at 216. In reaching this conclusion regarding these copies, the court relied on *HathiTrust*, where the court found that "both the making of the digital copies and the use of those copies to offer the search tool were fair uses." *HathiTrust*, 755 F.3d at 105. The *Google* court noted that the *HathiTrust* court had found that the downloading and storing of complete digital copies of books "was essential to permit searchers to identify and locate the books in which words or phrases of interest to them appeared." *Google*, 804 F.3d at 217. The *Google* court quoted the *HathiTrust* court's conclusion that "the creation of a full-text searchable database is a quintessentially transformative use . . . [a]s the result of a word search is different in purpose, character, expression, meaning, and message for the page (and the book) from which it is drawn." *Google*, 804 F.3d at 217 (quoting *HathiTrust*, 755 F.3d at 97).

The *Google* court then cited *A.V. ex rel. Vanderhye v. iParadigms, LLC*, 562 F.3d 630, 640 (4th Cir. 2009); *Perfect 10 v. Amazon.com, Inc.*, 508 F.3d 1146, 1165 (9th Cir. 2007); and *Kelly v. Arriba Soft Corp.*, 336 F.3d 811, 818 (9th Cir. 2003), as "examples of cases in which courts had similarly found the creation of complete digital copies of copyrighted works to be transformative uses when the copies served at different function from the original." *Google*, 804 F.3d at 217 (quotations omitted). All three of these cases involved the creation of a database to which AI algorithms were applied. *Kelly* and *Perfect 10* involved search engines designed to find images on the World Wide Web. *iParadigms* involved a plagiarism detection service that enabled an instructor to find works from which a student paper may have been copied. The *Google* court explained that "[a]s with *HathiTrust* (and *iParadigms*), the purpose of Google's copying of the original copyrighted books is to make available significant information about those books, permitting a searcher to identify those that contain a word or term of interest. . . ." *Google*, 804 F.3d at 217. *See also iParadigms*, 562 F.3d at 639-40 ("iParadigms' use of plaintiffs' works had an entirely different function and purpose than the original works. . . . iParadigms' use of these works was completely unrelated to the expressive content and instead aimed at detecting and discouraging plagiarism.").

Significantly, the *Google* court made clear that the commercial motivation of a provider of the AI database should not tilt the first factor against the provider: "[o]ur court has . . . repeatedly rejected the contention that commercial motivation should outweigh a convincing transformative purpose and absence of significant substitutive competition of the original." *Google*, 804 F.3d at 219. *See also iParadigms*, 562 F.3d at 639. In short, the first fair use factor should always weigh in favor of the creator of an AI database.

2. The Nature of the Copyrighted Work

Neither the *Google* nor *HathiTrust* courts found the second fair use factor, the nature of the copyrighted work, to be dispositive. This is "because the secondary use transformatively provides

---

[13] The non-expressive uses of works for the creation of AI algorithm are analogous to the "intermediate copies" made during the course of software reverse engineering. Courts have found that fair use permitted the translation of machine-readable object code into human readable source code as an essential step in the development of noninfringing interoperable computer programs. In these cases, the source code was used internally and was never distributed to the public. *See Sony Computer Entm't v. Connectix Corp.*, 203 F.3d 596 (9th Cir. 2000); *Sega Enters. Ltd. v. Accolade, Inc.*, 977 F.2d 1510 (9th Cir. 1992); *Atari Games Corp. v. Nintendo of Am., Inc.*, 975 F.2d 832 (Fed. Cir. 1992).

information about the original, rather than replicating protected expression in a manner that provides a meaningful substitute for the original." *Google*, 804 F.3d at 220. In other words, because the copying involved in the creation of a search database does not provide a meaningful substitute of the original, the nature of the original has little relevance. *See also iParadigms*, 562 F.3d at 641-42 (second factor does not weigh against iParadigms because its "use of the works in this case—as part of a digitized database from which to compare the similarity of typewritten characters used in other student works—is . . . unrelated to any creative component."). There could be cases where the nature of the works used, such as more factual works, might also support the finding of fair use. In any event, the second fair use factor does not tilt against nonexpressive uses by AI database creators.

3. The Amount and Substantiality of the Portion Used

Uses in the creation of AI algorithms often require the copying of entire works. The *HathiTrust* court concluded its discussion of the third fair use factor by noting that "[b]ecause it was reasonably necessary for the [HathiTrust Digital Library] to make use of the entirety of works in order to enable the full-text search function, we do not believe the copying was excessive." *HathiTrust*, 755 F.3d at 98. Likewise, the *Google* court found that "not only is the copying of the totality of the original reasonably appropriate to Google's transformative purpose, it is literally necessary to achieve that purpose. If Google copied less than the totality of the originals, its search function could not advise searchers reliably whether the searched term appears in a book (or how many times)." *Google*, 804 F.3d at 221. *See also iParadigms*, 562 F.3d at 642 (endorsing the district court's conclusion that iParadigms' use of the entirety of original works was limited in purpose and scope as a digitized record for electronic comparison purposes only). In other cases, an AI database might be made up of individual words or inconsequential portions of larger copyrighted works and thus the third factor would be expected to weigh in favor of such uses as well. In general, this factor and the case law interpreting it have well-served innovation in AI.

4. The Effect of the Use on the Market for the Original

The *HathiTrust* court found that the fourth fair use factor supported a finding of fair use because the ability to search the text of a book to determine whether it includes a search term "does not serve as a substitute for the books that are being searched." *HathiTrust*, 755 F.3d at 100. The *HathiTrust* court rejected the plaintiffs' suggestion that HathiTrust impaired the emergence of a market for licensing books for digital search: "Lost licensing revenue counts under Factor Four only when the use serves as a substitute for the original and full-text-search does not." *Id*. The *Google* court cited with approval *HathiTrust*'s conclusion that the search function does not substitute for the books being searched, *Google*, 804 F.3d at 223, and devoted the rest of its discussion of the fourth factor to snippet view. *See also iParadigms*, 562 F.3d at 644 ("Clearly no market substitute was created by iParadigms, whose archived student works do not supplant the plaintiffs' work in the 'paper mill' market so much as merely suppress demand for them, by keeping record of the fact that such works have previously been submitted.").

The *Google* court dismissed plaintiffs' assertion that the search function usurped their market for derivative uses. The court explained that "the copyright that protects Plaintiffs' works does not include an exclusive derivative right to supply . . . information [contained in the works] through query of a digitized copy." *Google*, 804 F.3d at 225. The court underscored this point: "Nothing in the statutory definition of a derivative work, or of the logic that underlies it, suggests that the author of an original work enjoys an exclusive right to supply information about that work of the sort communicated by Google's search functions." *Google*, 804 F.3d at 226. Given the absence of harm the Copyright Act is intended to prevent, the fourth fair use factor favors uses by AI database providers when those uses cause little or no harm to the legitimate markets of the originals.

5. The Four Factors Weighed Together

At the end of its fair use analysis, the Google court stated that "considering the four fair use factors in light of the goals of copyright, we conclude that Google's making of a complete digital copy of Plaintiffs' works for the purpose of providing the public with its search . . . functions . . . is a fair use and does not infringe Plaintiff's copyrights in their books." *Google*, 804 F.3d at 225. The reasoning of *HathiTrust*, *Google*, and *iParadigms* compels the conclusion that the uses necessary to make any AI database are fair use. Regardless of the nature of the content, providing AI functionality always has a different purpose and function from the content itself; copying entire works is often necessary to provide accurate operation of AI algorithms and processes; and such copying never substitutes for the original works. To help prevent this issue from being relitigated in every case involving an AI database, there should be a bright line permitting uses related to the creation and operation of an AI database. Such clear guidance not only would conserve judicial resources, it would prevent erroneous decisions. Drawing a bright line permitting the copying necessary to enable an AI process would benefit innovators and the courts.

E. Question 5: Should content generated by AI be eligible for protection by copyright or related rights?

A work produced by an AI algorithm or process, without the involvement of a natural person contributing to the resulting work, should not qualify for protection by copyright or related rights. Such a work would not be treated as a work of authorship protectable under the U.S. Copyright Act. The U.S. Copyright Office currently refuses to register a work that was not created by a human being. It explains:

> The copyright law only protects "the fruits of intellectual labor" that "are founded in the creative powers of the mind." *Trade-Mark Cases*, 100 U.S. 82, 94 (1879). Because copyright law is limited to "original intellectual conceptions of the author," the Office will refuse to register a claim if it determines that a human being did not create the work. *Burrow-Giles Lithographic Co. v. Sarony*, 111 U.S. 53, 58 (1884).[14]

The U.S. Copyright Office adds that it "will not register works produced by a machine or mere mechanical process that operates randomly or automatically without any creative input or intervention from a human author."[15]

Accordingly, if an AI algorithm or process creates a work that lacks expression from a natural person in the resulting work, the U.S. Copyright Office correctly would refuse to register the work. A court would uphold this refusal under the U.S. Supreme Court precedent cited in the Compendium.

As a matter of policy, the result should be the same in the UK. Withholding copyright protection from a work resulting from an AI process for which there was no expressive contribution by a natural person is justifiable from a policy perspective; the AI algorithm, and the computer that runs it, does not require the economic incentive provided by copyright in order to create works. Indeed, AI is capable of quickly producing an enormous array of works. Recognizing copyright in such output could quickly create a minefield of legal issues, leading to litigation and uncertainty.[16]

To be sure, the human creator of the software that runs the AI algorithm or process would receive a copyright in the expressive aspects of the AI software (and perhaps a patent for inventions in the AI software). We do not need copyright to incentivize the creation of AI generated works by giving copyright to each generated work because copyright already incentivizes the creation of AI software that generates all manner of content.

---

[14] Compendium of U.S. Copyright Office Practices III, § 306 (2017).
[15] *Id.* at § 313.2.
[16] *See* Daniel Gervais, *The Machine As Author*, 105 IOWA L. REV. (forthcoming 2020), *available at* https://ssrn.com/abstract=3359524.

F. Question 6: If so, what form should this protection take, who should benefit from it, and how long should it last?

As noted in response to Question 5, AI-generated works should not receive copyright protection. If a work resulting from an AI process is modified or enhanced by a natural person, that modification or enhancement may reflect sufficient originality to rise to the level of authorship. *See Feist v. Rural Telephone*, 499 U.S. 340 (1991). In the United States, a court applying traditional copyright principles to the facts of such a case would be able to determine whether there was enough expressive contribution by a human to rise to the level of authorship and whether the portions of a work that are reproduced in a particular case are protectable human creativity or unprotectable, AI-generated output. Of course, any copyright would extend only to the original elements contributed by the natural person.

CCIA opposes the concept of according a legal personality to an AI application that creates original works autonomously.

G. Question 8: Does copyright provide adequate protection for software which implements AI?

The UK provides adequate protection for software, including software which implements AI.

H. Question 9: Does copyright or copyright licensing create any unreasonable obstacles to the use of AI software?

We are not aware of any such obstacles.

25 Massachusetts Avenue NW #300C
Washington,DC 20001
+1(202) 783-0070
https://www.ccianet.org

Rue de la Loi 227, First Floor
1040 Brussels, Belgium
+32-2-888-8462

Respectfully submitted,

Joshua Landau
Patent Counsel
Computer & Communications Industry Association
25 Massachusetts Ave NW
Suite 300C
Washington, DC 20001
jlandau@ccianet.org

Ali Sternburg
Senior Policy Counsel
Computer & Communications Industry Association
25 Massachusetts Ave NW
Suite 300C
Washington, DC 20001
asternburg@ccianet.org

Victoria de Posson
Senior Manager, Public Policy
Computer & Communications Industry Association
1st Floor
Rue de la Loi 227
1040 Brussels Belgium
vdeposson@ccianet.org