



Before the
World Intellectual Property Organization
Geneva, Switzerland

In re

Impact of Artificial Intelligence on IP Policy: Call
for Comments

**COMMENTS OF
INTERNET ASSOCIATION, HIGH TECH INVENTORS ALLIANCE, TECHNET**

Internet Association (IA),¹ High Tech Inventors Alliance (HTIA),² and TechNet³ appreciate the World Intellectual Property Organization's (WIPO) interest in laying the groundwork for discussions on how artificial intelligence (AI), particularly machine learning, impacts intellectual property (IP) policy. As WIPO seeks to develop a list of relevant issues that might form the basis of future structured discussions, it is important that WIPO not outpace ongoing conversations in countries around the world.

AI and algorithms are used daily by every American with an internet connection or smartphone. AI helps consumers navigate the internet with ease and empowers internet companies to provide more useful and more personalized experiences online. In the last decade, we have made incredible strides in machine learning algorithms, which have powered artificial intelligence. As this technology has become more powerful and integrated into everyday life, there are various policy implications to consider, particularly in the realm of IP.

Many countries, including the U.S., are currently undertaking reviews of their own policies regarding AI and IP. The United States Patent and Trademark Office (USPTO) recently published several requests for comments on IP and AI but has not yet released its findings or recommendations. It would be premature for there to be any changes to the global IP system until such reviews are complete.

As WIPO continues to examine the policy implications of AI, WIPO should be prepared to consider several important AI-related issues concerning IP that Associations have already shared with USPTO in response to its requests for comments.

Patents and AI

AI is making an increasingly large impact on the patent ecosystem. The number of AI patents granted globally has increased rapidly over the last decade.⁴ The current ecosystem is encouraging and facilitating a boom in AI-related patents and no immediate changes to the law are necessary. Rather, IP law should continue to empower these advances and not limit future technologies.

¹ <https://www.internetassociation.org/our-members/>

² <https://www.hightechinventors.com/about>

³ <http://technet.org/our-story/about>

⁴More than half of all identified AI patents were published in the last five years. WIPO, "Technology Trends 2019: Artificial Intelligence." 2019. https://www.wipo.int/edocs/pubdocs/en/wipo_pub_1055.pdf

There are multiple, distinct types of inventions that could be termed “AI inventions.” Those various types of AI inventions can be loosely separated into three categories: (1) *AI innovations*—inventions that embody an advance in the field of AI; (2) *AI applications*—inventions that apply AI to another field; and (3) *AI-developed inventions*—inventions that are produced by AI.

An *AI innovation* is an advance in the field of AI technology itself and may include, for example, a new neural network structure of an improved machine learning model or algorithm. Despite their complexity, such inventions have existed for decades and can be described, claimed, and examined in the same way as other software inventions. As a result, there is no conflict with established claiming and disclosure practices, and these inventions are unlikely to present significant new challenges with respect to the application of substantive patentability requirements.

However, some AI innovations will present more complex examination difficulties that stem from the nature of machine learning—the dominant form of AI. Machine learning does not involve explicit instructions but instead, the computer is “trained” using statistical methods that produce an analytical or mathematical model based on data analysis. Unlike a series of algorithmic steps, machine learning training produces a mathematical model that is derived by a computer and is expressed in a form that may not be comprehensible to human experts. The inability to describe precisely how particular results are produced by a trained model is often referred to as the “black box” or “interpretability” problem in AI. But just as some pharmaceutical innovations, for example, may not operate in a way that is fully understood by their own inventors, inventors can and must still meet the disclosure and enablement requirements, to ensure that another skilled artisan can make and use the claimed invention and be assured that the inventors possessed the invention at the time of filing.

For example, even if the inventor cannot explain exactly how the AI works, she must still explain the various functions implemented in each element of the neural network along with the topology of the network and the type of data needed to build the model. In other words, existing patent principles that accommodate alternative ways to describe inventions can—and must—be applied even to complex AI innovations.

An *AI application* is just that—the application of AI to a particular field or problem. Just as the invention of computers naturally led to their use in conventional problem solving, and just as the internet led to its use in communications and commerce, AI will have natural benefits in existing fields. However, as we have come to understand with computers or the internet, it is the technical advance that must be considered for patentability purposes and not simply the notion, or recitation in a claim, that AI be applied to a new use. As with AI innovations, and as discussed below, the Associations believe that existing laws adequately address patent applications directed to AI applications.

The Associations acknowledge the possibility of a third type of invention that one might call an *AI-developed invention*. Specifically, an AI-developed invention is an invention in any art that is developed by a machine rather than a human. As a threshold matter, the Associations do not believe that machines can be inventors. On closer inspection, depending on what is being claimed, it is likely the machine programmers or technologists who are evaluating the output of a machine learning model who are the actual inventors. AI is used as a tool to identify solutions to a problem already defined by human operators or to model alternative potential solutions to evaluate the suitability of each solution and identify the best candidates. Generally, these candidates will have to undergo further human evaluation and adjustment to demonstrate that they are a suitable solution. As a result, consideration of a separate “developed by” category is not analytically useful. At least for now, the focus should be on the more meaningful distinction between AI innovations and AI applications rather than on the inventor.

The AI invention category of AI applications presents particular challenges to examiners. One concern is that an examiner who is unfamiliar with AI may incorrectly believe that claims are adequately described or enabled. Similarly, an examiner unfamiliar with AI may not readily understand the true breadth of the claims and inadvertently allow a patent to issue with claims that cause downstream hindrance of

innovation. When a technology like AI expands rapidly, applicants may engage in what amounts to a land grab—purposefully attempting to tie up as much scope as possible. In order to police this potential behavior, it is critical that patent examiners receive regular technical training on the subject matter they examine. As part of this training, it is imperative that patent examiners are aware of and have access to sources of non-patent literature when searching for prior art, as that is often where the most relevant information is found. Training patent examiners, as opposed to establishing new IP policy, should be the primary response to the emergence of AI.

Appropriate disclosure is essential to narrowly drawn, well written patents. While AI inventions as a class do not create unique considerations relating to disclosure, the same rules and principles that apply to all other types of inventions are appropriate in the AI context. Nevertheless, there are considerations relating to the fact that AI is a rapidly evolving technology likely to play a critical role in the global economy. Given that, it is critical to avoid misallocation of rights, and to ensure that a company's development or adoption of AI technology does not subject it to unreasonable litigation or business risks as a result of overly broad or invalid patents.

Regardless of the specific nature of the AI invention, the structure of the machine learning model, system, or software algorithm should be described with enough specificity to show possession of the model or improvement as claimed. If the claims are directed to a class of AI innovations, the specification should include language showing examples of or guidance for achieving the result using the class of algorithms, and not just a single example.

In the case of an AI application, a skilled artisan is already aware of an existing machine learning model or technique. Such a description should therefore include identification of an algorithm by name with additional descriptions of what portions of the algorithm have been modified to meet the application, and, if applicable, how different algorithms are connected to each other. If the purported AI invention is described as working with numerous types of artificial neural network algorithms, the application should describe whether and how the invention works with each of the type of artificial neural networks, such as what particular changes would be required to apply the invention as claimed.

Examiners should use care in the case of claims to an AI application to confirm that the inventor has sufficiently described possession of the application itself. Thus, depending on where the innovation lies, one might expect to see detailed descriptions of the data inputs, training data, or the coefficients that resulted from the training data and the outputs to the existing machine learning model needed to adapt it to the new application. The description would likely also need to describe the particular type and scope of data in addition to its structure because certain training algorithms may be dependent on the data's characteristics. To the extent that specific details on hardware (e.g., use of AI accelerator hardware) are needed to actually implement the invention or achieve the results claimed, the specification should include details on the hardware capabilities required.

Copyright and AI

There are several ways in which an expansion of copyright law in the area of AI could actually frustrate the purpose of copyright by impeding the progress of science and learning. For example, copyright law should not be expanded to grant copyrights to works created solely by AI algorithms. Algorithms are not human authors. Copyright was originally conceived as protection for “original intellectual conceptions of the author”⁵ and the “fruits of intellectual labor,”⁶ which are, by definition, “founded in the creative powers of the mind.”⁷ Under existing international copyright law standards, if a work does not contain any human expression, copyright authorities in the United States and elsewhere properly will refuse to register copyright in it. Further, AI is capable of producing an enormous volume of works in a very short

⁵ See *Burrow-Giles Lithographic Co. v. Sarony*, 111 U.S. 53, 58 (1884).

⁶ See *Trade-Mark Cases*, 100 U.S. 82, 94 (1879).

⁷ See *id.*

time. Granting copyright to this type of machine output would very quickly create an unprecedented copyright thicket and a minefield of legal issues, which necessarily would lead to a dramatic uptick in copyright litigation.

The second way in which lawmakers could actually frustrate the purpose of copyright by expanding it would be to require owners of AI algorithms to secure licenses for source material that algorithms use to “learn.” Machine learning requires the ingestion of large amounts of source material in a form that machines can read. This, in turn, usually entails converting that material into a more usable format, which requires the owner of the algorithm to make copies,⁸ not only to facilitate the process but also to protect against loss of data in the event of system failure.⁹

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.¹⁰ All these copies are not viewable or consumable by the outside world.¹¹ Because these non-expressive copies are not consumable by the public, they do not function as market substitutes for copies of the ingested works.¹² Numerous appellate courts have undertaken a fact-intensive inquiry and correctly found the mass copying of raw material to build databases for uses by AI processes to be fair use under U.S. law in those cases.¹³ Numerous legislatures have implemented (or are in the process of implementing) legal exceptions that permit the use of works for such computational analysis.¹⁴ Such “non-expressive uses” are not unimportant uses; to the contrary, having access to a large volume of readable source material is critical to an AI algorithm’s ability to do what humans cannot—*i.e.*, process large volumes of information and make the connections that humans could not make in a lifetime of study, and such uses should fall under exceptions to copyright law such as fair use. There simply is no other way to put these algorithms to work, and there is no way humans could do this work without AI.

⁸ In the Google Library Project, for example, Google made a digital scan of each book it borrowed from a research library, then used optical-character-recognition (OCR) software to convert the scanned image into machine-readable text. Google retained both the scanned image and the machine-readable text. *Authors Guild v. Google, Inc.*, 804 F.3d 202, 208 (2d Cir. 2015).

⁹ 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, Inc. v. HathiTrust*, 755 F.3d 87 (2d Cir. 2014).

¹⁰ 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).

¹¹ 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).

¹² 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).

¹³ See *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).

¹⁴ See e.g. Europe: Articles 3 and 4, <https://ec.europa.eu/digital-single-market/en/modernisation-eu-copyright-rules>; Japan: 2018 revisions to Japan’s Copyright law, adding new Articles 30-4 (non-consumptive uses), 47-4 (incidental copying for computer uses), and 47-5 (allowing creation of searchable databases of copyrighted works for analytics) at http://www.mext.go.jp/b_menu/houan/kakutei/1405195.htm; Singapore: Public Consultation on Proposed Changes to Copyright Regime in Singapore at <https://www.https://www.mlaw.gov.sg/news/public-consultations/public-consultation-on-proposed-changes-to-copyright-regime-in-s>; Canada: <https://www.ourcommons.ca/Content/Committee/421/INDU/Reports/RP10537003/indurp16/indurp16-e.pdf>

Conclusion

Our Associations appreciate WIPO's interest in this important topic and the opportunity to submit these comments. Issues related to IP and AI are certainly worthy of discussion, and many of those discussions are currently underway in individual countries. WIPO should allow those analyses to produce meaningful results before attempting to tackle the same challenging questions in a different forum.