Before the

World Intellectual Property Organization

Geneva, Switzerland

In re

Impact of Artificial Intelligence on IP Policy: Call for Comments

COMMENTS OF THE SOFTWARE AND INFORMATION INDUSTRY ASSOCIATION

The Software and Information Industry Association commends the World Intellectual Property Organization for examining the intellectual property (IP) issues created by advances in artificial intelligence (AI). SIIA is the principal U.S. trade association for the software and digital content industries. With over 800 member companies, SIIA is the largest association of software and content publishers in the country. Our members range from start-up firms to some of the largest and most recognizable corporations in the world. The innovative companies that make up SIIA's membership have embraced artificial intelligence

Our members actively use AI on many fronts—from journalism to fraud detection, money laundering investigations, and locating missing children. They use artificial intelligence to help people make use of an increasingly large pool of data sets and invest billions in its development, acquisition and use. At the same time, however, the use of AI must comply with existing statutory requirements and respect for established intellectual property rights.

Artificial intelligence has been with us for some time. Today, when people refer to AI, they are typically referring to areas like "machine learning" or "deep learning." "Machine learning" refers to a process by which the computer improves the exercise of particular functions by correcting its errors. The computer accomplishes its objective through a process of trial and error as it assigns different weights to particular inputs received by individual "nodes," which can be analogized to human neurons. These nodes are the basis of "neural networks," which are more concisely visualized than explained:

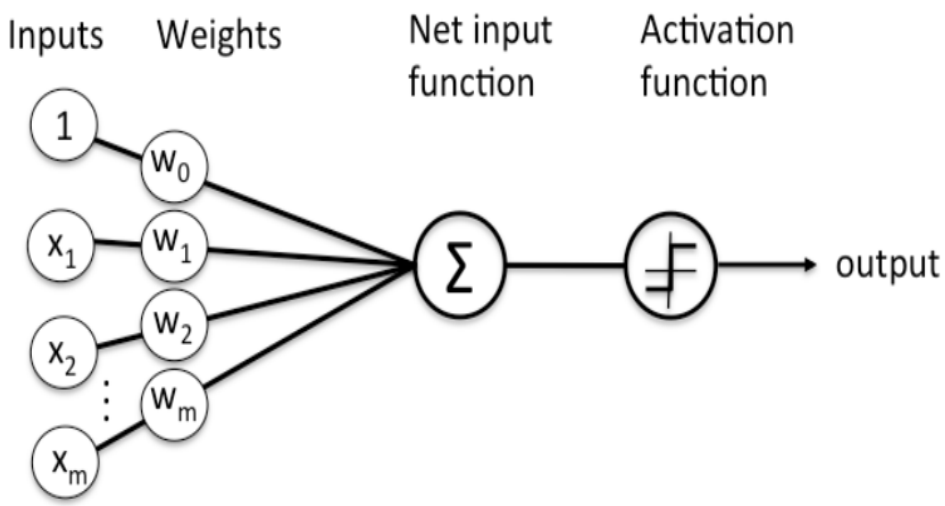

Source: skymind.ai

Suppose, for example, that a computer tried to determine whether an email was spam. Each input (a word or phrase) would be assigned different weights and combined into a net input function, e.g. "spam" or "not spam" output. If the output is incorrect, then the algorithm adjusts the weights and performs the function again. Humans speed the process along by training the machine so it learns. Each error results in a correction until the computer gets it "right."

"Deep learning" simply refers to the number of layers of nodes through which a particular input has to pass before identifying a pattern—whether identifying spam or a particular image. Given the realities of current technology, modern deep learning can consist of hundreds of thousands of these nodes or more, updating themselves multiple times per second. These kinds of algorithms can identify patterns and correlations in unstructured data such as photographs, newspaper and journal articles, sound recordings and video.

Applications for this technology abound, and we are facing a healthy environment for innovation, development and implementation. Venture capital investment in the software and internet industries has hit $45 billion, and our members invest billions in technological improvements.[1] That industry environment in general is even more robust when it comes to AI: between January 2015 and January 2018, the number of AI startups has increased by 113%, as compared to 28% for startups generally.[2] Similarly, venture capital funding for AI increased by 350% between 2013 to 2017, a rate over three times higher than that for venture capital

---

[1] National Venture Capital Association, Venture Monitor, 4Q 2018, https://files.pitchbook.com/website/files/pdf/4Q_2018_PitchBook_NVCA_Venture_Monitor.pdf.

[2] AI Index 2018 Annual Report, at 31, available at http://cdn.aiindex.org/2018/AI%20Index%202018%20Annual%20Report.pdf.

investing generally.[3] Job growth in the field is large and accelerating.[4] And the technology itself has achieved remarkable milestones. For example:

- An AI system achieved humanlike translation quality when translating Chinese news stories into English.[5]
- An AI system examined 29,450 clinical images of 2,032 different diseases and achieved diagnostic parity with board-certified dermatologists.[6]
- An AI system generated a 70% success rate in detecting prostate cancer by examining specimens—exceeding the 61% rate of board-certified pathologists.[7]

Most businesses lack the ability to design and implement custom AI solutions, but the demand for this technology is vast. Access to these tools is quickly becoming democratized, and that democratization has been advanced by three factors. First, the advent of cloud computing has enabled "off the shelf" open source solutions that can be trained and implemented by corporations and small businesses.[8] Second, solutions are advancing such that a customer can simply use their own data and ask an AI service to use machine learning to create a custom-trained model.[9] Third, hardware is improving through the development of specialized chips that can supply the computing power necessary for deep learning analysis.[10] And finally, in many instances, a user will not train an AI network at all—they can incorporate tools like voice recognition or language translation into an existing product via pre-existing APIs. [11]

---

3 *Id.* at 32.

4 *See id.* at 33.

5 Microsoft reaches a historic milestone, using AI to match human performance in translating news from Chinese to English (Mar. 14, 2018) available at https://blogs.microsoft.com/ai/chinese-to-english-translator-milestone/.

6 Dermatologist-level classification of skin cancer with deep neural networks, Nature (Jan 2017), available at https://www.nature.com/articles/nature21056.

7 Improved Grading of Prostate Cancer with Deep AI Learning, Google AI Blog, https://ai.googleblog.com/2018/11/improved-grading-of-prostate-cancer.html.

8 *See,e.g.,* https://www.tensorflow.org/tutorials/quickstart/beginner.

9 *See, e.g.*, https://azure.microsoft.com/en-us/services/cognitive-services/custom-vision-service/. Microsoft is not alone in providing these services. E.g., https://www.clarifai.com; see also https://cloud.google.com/automl/ (requiring minimal machine learning knowledge).

10 See, e.g., https://www.intel.ai/#gs.dvcap3, https://www.theverge.com/2018/7/26/17616140/google-edge-tpu-on-device-ai-machine-learning-devkit.

11 Microsoft, Google, and other companies have created these interfaces. See https://nordicapis.com/5-best-speech-to-text-apis/.

The combination of growth and democratization means that before long, ordinary businesses will be training these networks by feeding them unstructured data and using that output to create new products. The balance of these comments addresses SIIA's views with respect to both patent and copyright and AI, which we define as a neural network capable of mimicking human thinking.

Patents.

As applied to patents, issues will arise in three basic AI use cases: (1) inventions that embody an advance in the field of AI; (2) inventions that apply AI to another field; and, arguably, (3) inventions that are produced by AI. We refer to the first group as *AI innovations* and the second group as *AI applications*.

An *AI innovation* is an advance in the field of AI technology itself and might include, for example, a new neural network structure of an improved machine learning model or algorithm. Aside from their complexity, such inventions could be described, claimed, and examined in the same way other software inventions have been. 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.

There will be, of course, AI innovations that present more complex examination difficulties that flow from the nature of machine learning, which is the dominant form of AI. 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. In other words, existing patent principles that accommodate alternative ways to describe inventions can—and must—be applied even to complex AI innovations. For example, if the inventor cannot explain exactly how the AI works, she may 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.

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. There may be technical difficulties to particular applications that amount to patentable inventions, but just as we have come to understand with computers or the internet, it is the technical solution to a technical problem that forms the core of patentability and not simply the notion, or recitation in a claim, that AI be applied.

Because the assessment of inventorship is fact dependent and because the potentially-qualifying contributions are numerous and varied, it is unlikely that all the ways in which a natural person might qualify as a named inventor can be

anticipated *ex ante*, making it impossible to compile an exhaustive list. As a general matter, however, AI should not be an "inventor" for patent purposes. Conception of an AI invention should be assessed like any other invention and will thus depend on the characteristics of the invention claimed, as well as the development process that results in the invention. For inventions that are advances in the field of AI, AI innovations, the nature and development of the invention are likely to be very similar to those of existing software-related inventions, and the assessment of inventorship will generally be based on the same considerations and factors that are already employed with respect to software.

In short, existing frameworks are capable of addressing this and other problems. We generally agree with the conclusion reached in the study on this topic that was commissioned by the EPO and authored by Dr. Noam Shemtov that "at present there are no particular difficulties associated with ownership enquiries relating to inventions involving AI activities" and that the economic arguments for extending inventorship to AI machines, systems, or software are "not supported by empirical data or a credible economic model."[12]

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. 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, such as claims to achieving a result where the claims cover potential use of multiple different algorithms and not just a single algorithm, 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. With these kinds of difficulties in mind, the focus of enablement and written description should always remain on whether a person of ordinary skill in the art could make and use the claimed technology and confirm that the applicant was in possession of what she claimed. Maintaining that focus, however, does not require a re-evaluation of fundamental patent doctrine.

Copyright.

In SIIA's view, it is premature to advocate changes to copyright laws based on the state of AI development, as many of our members are deploying it both robustly and responsibly. With that said, it is not difficult to envision circumstances in which AI could be misused. Thus, for example, the translation and publication of entire

[12] Noam Shemtov, *A Study on Inventorship in Inventions Involving AI Activity*, 24, 34, (Feb. 2019); Study Commissioned by the European Patent Office, available at http://documents.epo.org/projects/babylon/eponet.nsf/0/3918F57B010A3540C125841900280653/$File/Concept_of_Inventorship_in_Inventions_involving_AI_Activity_en.pdf.

news stories to a different language would require a license of the right to create derivative works, as they could easily supplant the market for the copyrighted originals.

These principles of incentives and harms of substitution are not technology-dependent—quite the opposite. Article 9(2) of the Berne Convention, for example, instructs that limitations on reproduction not interfere with normal exploitation of copyrighted works. And they require a common understanding of what the term "AI" means. As we discuss copyright and AI, SIIA assumes for the purposes of argument that AI consists of a neural network capable of producing material equivalent to that produced by a human author. With that general framework in mind, we urge the following points for WIPO's consideration.

First, a work produced by an AI algorithm or process, without the involvement of a natural person contributing expression to the resulting work should not qualify as a work of authorship. The text of the Berne convention refers to the rights of authors authors. It does not (and should not) be read to encompass works in which a machine supplies creative expression in the absence of human originality. Thus, existing copyright instruments do not as a general matter provide for such rights, and such a construct exists well outside both the U.S. and the international understanding of copyright law.

With that said, authors may of course use automated tools to create works, from word processors and increasingly robust editing software to robust image editors and other kinds of creative assistance. At some point, the link between a human's creativity and the software becomes too attenuated, and the requisite link disappears and the work is not subject to protection. And at this point in time, it is premature to guess as to the effect that an abundance of AI-created and public domain works would have on different industries and different economies.

Second, to the extent an AI algorithm or process learns its function(s) by ingesting large volumes of copyrighted material, existing frameworks already address how this issue ought to be handled: on a fact-specific basis. Certainly, intonation of the phrase "AI" should not inculcate acts of unauthorized copying from examination. In particular, where information is made available by license, such licenses ought to be respected, and many of our members, especially those who publish journals, make their works available for exactly this purpose.

At the same time, in the absence of privity, existing law does permit the use of materials to train AI. Many of our members obtain publicly available information and use it to train their AI engines, or link to it. The interference of that activity with established exclusive rights must be determined on a case by case basis, balancing the author's interest in the preservation of her incentive against the nature and purpose of the use, the substantiality of the copying, the nature of the work copied, and the effect of the use on the actual or potential market for the

underlying work. By definition, the legality of these kinds of uses will be a fact-specific decision that augurs against the development of bright-line rules. Development of an international standard, therefore, seems to us to be premature.

Third, copyright principles should (and generally do) ensure that the user of the AI should be legally responsible for the consequences of the output as well as the input. Established principles of corporate and individual liability would seem in most cases to be readily transferable to the AI context. Thus, if a person directs an AI engine to copy and otherwise use a particular work, liability would fall on the person who so directed it. Cases in which an AI engine – on its own and without human intervention – copied a protected work cannot be addressed in the abstract, but it is fair to posit that liability in such a circumstance, if it exists, at a minimum ought to flow to the person that benefitted from the activity.

Conclusion

Many governments, including the U.S. government, are examining he implications of AI's growth for intellectual property. The United States Patent and Trademark Office (USPTO) recently received comments on issues raised by AI with respect to patents, copyright and trademark. Proposed changes to the global IP system should not emerge until these kinds of reviews are complete.

With that said, this is an important and rapidly developing field, and we commend WIPO for its examination. Thank you for considering our views.

Respectfully submitted,

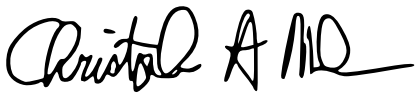

Christopher A. Mohr

Vice President for Public Policy and General Counsel

February 14, 2020