The Intersection Of IP And Artificial Intelligence (AI)

By Susan Hallen, Kate Gaudry, Stephen Kabakoff and Sean O'Connor

In 1956 at the Dartmouth Summer Research Project on Artificial Intelligence, John McCarthy coined the term "Artificial Intelligence,"[1] which he would define as "the science and engineering of making intelligent machines." The usage/capabilities of artificial intelligence (AI) as a technology have been accelerating ever since, as computing power has grown exponentially, and more research has been devoted to artificial intelligence and its use cases. As a result, artificial intelligence has become an umbrella term that encompasses everything from Natural Language Processing/Understanding (NLP) to Machine Learning to Robotics to Deep Learning, and much more. Think of it as Cognitive Computing…a range of technologies that mimic the human brain and take actions.

While AI typically refers to "Artificial Intelligence," some entities have proposed alternative terminology. For example, IBM sometimes considers AI as representing "Augmented Intelligence," analogous to the smart co-worker that reads, understands and analyzes quickly, then helps you combine your knowledge and experience to leverage insights. Nearly all interpretations are based on an understanding that AI, as a whole, allows machines to work "intelligently," and that AI corresponds to a set of methods, systems or approaches for computers that try to mimic human judgments and human decision making. Machine Learning is a particular type of AI that provides computers with the ability to continue learning without being pre-programmed. Machine Learning comprises algorithms that learn from data and create foresights based on this data. Deep Learning is a particular type of Machine Learning that leverages neural networks and further automates learning without the requirement for massive data sets.

Thus, AI is a continuum that spans levels of adaptiveness, data requirements, and sophistication. The overall power of the techniques and utilization of the technology has been exploding. For example, DeepMind Technologies developed AlphaGo, which beat the top-ranked player in 2017 in the game Go. As another example, GPT-3 (which has been recently exclusively licensed to Microsoft) produces text that is – in many circumstances – nearly indistinguishable from that written by a human. Further, Narrow AI (focused on a single task or domain) is widespread today and growing; and Broad and General AI (which has the capability to perform across tasks and domains) is in progress under development.

As AI evolves, addressing particular questions about the development, control, and use of AI becomes increasingly important. Many of these questions boil down to trust…How can one trust a black box; and how can one trust that use of a powerful algorithm will be unobjectionable? What is the basis for decisions generated by AI? How was the system trained? Sustained and mainstream use of AI generally requires that users trust the system understands context, is free of bias, is trained by valid and relevant data sets, is secure, and is transparent. This is where ethics come into the discussion. Most large companies have some AI ethics policy today, and many are based on the European Commission, High-Level Expert Group on AI, "Ethics Guidelines for Trustworthy Artificial Intelligence."[2] Typical pillars include Accountability, Value Alignment, Explainability, Fairness and User Data rights.

Developing a system establishing performance and/or trust in these pillars may provide as much business advantage as would developing a system that establishes a target performance accuracy. As an example, various AI systems may be prone to bias, *i.e.*, prejudicing in favor or against something, which may lead an entity to behave unfairly toward a certain group compared to others. To illustrate, if an AI system is built to recognize whether a picture contains a person or not, the AI

■ Susan Hallen,
Associate Partner,
AI Workflows,
IBM Global Services,
Chicago, IL U.S.A.
E-mail: sahallen@us.ibm.com

■ Kate Gaudry,
Partner,
Kilpatrick Townsend & Stockton LLP,
Washington, D.C. U.S.A.
E-mail: kgaudry@kilpatrick-townsend.com

■ Stephen Kabakoff,
Partner,
Miller & Martin PLLC,
Atlanta, GA U.S.A.
E-mail: stephen.kabakoff@millermartin.com

■ Sean O'Connor,
Professor of Law,
George Mason University,
Founding Director, Innovation Law Clinic,
Faculty Director, Center for Intellectual Property x Innovation Policy (C-IP2),
Arlington, VA U.S.A.
E-mail: soconn3@gmu.edu

1. John McCarthy, *https://computerhistory.org/profile/john-mccarthy/?alias=bio&person=john-mccarthy*.

2. European Commission, High-Level Expert Group on AI, "Ethics Guidelines for Trustworthy Artificial Intelligence" - *https://ec.europa.eu/futurium/en/ai-alliance-consultation.*

system may be trained with a large number of pictures associated with labels indicating whether a person is depicted. However, if the examples are not diverse, representative, balanced, or inclusive of the human population, the AI could fail to generalize what it has learned to pictures never seen before. If all examples contain people with lighter skin tones, the system will have trouble recognizing people with darker skin tones. Embed this AI system into a decision-making process, and it is easy to see that this could lead to unfair treatment for certain groups. More than 180 human biases have been defined and classified, any one of which can affect how we make decisions. Many of these biases can find their way into AI systems that are used by governments and businesses to make decisions. IBM, for example, has published research on how to detect and reduce bias in a training dataset.

Intellectual property systems provide avenues to incentivize discovery, *e.g.*, of new AI techniques, new uses of AI techniques, new approaches for establishing trust in AI techniques, and/or new uses of AI results, by virtue of providing limited monopoly. What kinds of discoveries might societies want to incentivize in the field of AI? Are our current intellectual property systems designed to provide these incentives?

Here, we focus on the intersection of AI and intellectual property (IP) in the form of patents and copyrights:[3]

- Part I focuses on patenting prospects pertaining to AI-related technologies, such as new types or uses of AI. Given the wide variety of technologies that may pertain to these innovations, are the prospects of securing patents relating to AI uniform across technologies? Is this uniformity even desirable?
- Part II then considers the powerful set of AI tools that can be used as part of the process of securing patents on other inventions. Such AI tools and related technology can themselves be patented, copyrighted (where applicable as source code), and protected as trade secrets.
- Part III discusses certain developments in the U.S. Patent & Trademark Office (USPTO) and elsewhere regarding IP protection of AI inventions, code, and data. Should patentability requirements be revised as AI becomes increasingly used throughout society?
- Finally, questions of AI inventorship and authorship are still live issues and will likely become more so over time. A basic analysis of these issues is presented in Part IV, while noting that, in practice, commercial entities will have little incentive to position their AI tools as independent inventors or authors; rather this is more relevant to those who want to show the power of the tools or challenge existing IP systems.

I. IP Protection of New Types or Uses of AI

While patent offices and stakeholders alike have been attuned to policy issues related to AI innovations, AI innovations is a broad and imprecise concept. AI innovations may include: (1) innovations where AI contributed to conception of an invention; (2) innovations focused on a new type of AI (*e.g.*, a new type of machine-learning model) or a new AI-based processing flow (*e.g.*, that uses a new type of pre- or post-processing and further uses a machine-learning or intelligent rules-based model); or (3) a new use of AI (*e.g.*, using an existing AI-based processing technique to process a type of data not previously processed by the technique and/or generate a type of result not previously generated by the technique).

Even when innovations where AI contributed are excluded, "AI innovations" can relate to many different technological fields, *e.g.*, natural language processing, robotics, computer vision, bioinformatics, etc. A USPTO study tracked recent AI filings, where the agency considered AI filings to be those relating to AI technologies for machine learning, AI hardware, natural language processing, evolutionary computation, speech, vision, knowledge processing, and planning/control, recognizing that some patent applications are AI-related because they include multiple AI technologies.[4] Like other patent offices, the USPTO has observed a substantial increase in the number of AI-related patent applications that have been filed. The USPTO, for example, reported that in the 16 years from 2002 to 2018, the USPTO's annual AI patent applications have increased by more than 100 percent, rising from 30,000 to more than 60,000 annually. The USPTO further reported that patent applications containing AI appeared in about 9 percent of all technology subclasses used by the USPTO in 1976 and increased to more than 42 percent by 2018.[5]

As discussed in Part III, the rapid growth of AI-related patent applications has raised interesting questions about whether conventional legal standards should apply to such AI applications. The World Intellectual Property Organization (WIPO) has conducted meetings among its members and identified many legal issues that arise from the intersection of AI and IP.[6] The USPTO has similarly

3. While there are also interesting AI and IP issues in the area of trade secrets and trademarks, they are outside the scope of this present article.

4. USPTO report, "Inventing AI, Tracing the diffusion of artificial intelligence with U.S. patents," Office of the Chief Economist, *IP Data Highlights,* No. 5, Oct. 2020 - *https://www.uspto.gov/sites/default/files/documents/OCE-DH-AI.pdf.*

5. *Id.*

6. "WIPO Conversation on Intellectual Property (IP) and Artificial Intelligence (AI)," WIPO Doc. No. WIPO/IP/AI/2/GE/20/1 REV, dated May 21, 2020—*https://www.wipo.int/edocs/mdocs/mdocs/en/wipo_ip_ai_2_ge_20/wipo_ip_ai_2_ge_20_1_rev.pdf).*

solicited public comments on issues concerning IP rights in relation to AI technology.[7]

For example, consider how patent applications pertaining to AI should be examined by a patent system. Typically, each examiner at each patent office is designated to examine patent applications pertaining to a particular type of technology. In the U.S., the USPTO further assigns individual patent applications to a single "art unit" (that is associated with one or more supervisors and a given type of technology) and one or more classifications associated with a given type of technology. For example, class 706 (examined by examiners in art unit 2122) is associated with patent applications relating to AI model innovations; a different art unit 1631 examines patent applications in classes related to bioinformatics (nearly all of which relate to AI); and class 382 is associated with image-analysis innovations.

Table 1. Allowance Prospects

	% Actions that were allowances			
	Q4 2018	Q2 2019	Q2 2019	Q2 2020
Straight A (Class 706)	22.1	46.1	30.7	25.5
Bioinformatics (Art Unit 1631)	19.9	30.4	24.1	23.9
Image Analysis (Class 382)	53.5	51.1	53.0	46.9

Perhaps unsurprisingly, the prospects of securing a patent vary quite significantly across different parts of the patent office and over time. To illustrate this variability, data was collected from LexisNexis® PatentAdvisor℠ to determine the number of office actions (rejections) and the number of allowances issued by various classes and art units within each quarter. The fraction of allowances divided by sum of the office actions and allowances provides some approximation of allowance prevalence. (Notably, this ratio is generally lower than the "allowance rate," which indicates an overall likelihood of securing a patent, given that many allowances are issued after multiple office actions having already been issued and addressed.) Table 1 shows that the allowance prospects in these AI-related areas varies by approximately two-fold during most time periods. Further, the changes in allowance prospects across time periods can be substantial (*e.g.*, 22.1% to 46.1% for class 706 and 19.9% to 30.4% for art unit 1631 when comparing Q4 2018 to Q2 2019).

The high variability in the allowance prevalence across time and technology areas may be a result of, for example, differences in patent-office training and/or case-law applicability; differences in supervisors' interpretation of case law and agency guidance; differences in how much prior art exists in various areas; etc. Rejection-specific analyses indicated that most – but not all – of the variability across technology areas and times tracks (in an inverse manner) the prevalence of eligibility rejections, which asserts that a claimed innovation does not correspond to a type of subject matter that can be patented.

II. AI Tools Usable in the Patenting Process

AI tools are now also becoming increasingly used to improve applicants' odds of securing patents. For example, as evident by the above analysis, big data is now available to identify the allowance prospects for various art units, classifications, and time periods. Natural language processing tools now exist that allow entities to predict to which art unit or class an application will be assigned. Thus, an applicant can make informed decisions as to which invention aspects to emphasize and/or whether to file a patent application at all. As another example, not only do big data tools track high-level statistics, but they also track statistics as to how effective various prosecution strategies may be, *e.g.*, appealing rejections to the Patent Trial and Appeal Board rather than continuing to engage a particular examiner. Thus, applicants can dynamically shift prosecution strategies based on the relative allowance prospects associated with different prosecution strategies.

Further, AI may facilitate searches that inform whether an entity may be able to operate without infringing another entity's IP rights or may be able to secure their own IP rights. Think of what it takes to come up with your list of keywords and perform searches, then sort out the data and combine spreadsheets, iterate, read a ton of large documents, and much later you have your short list of results or potential answers to whatever question you had. Do you really have the most relevant matches? By leveraging big data ingestion/understanding, machine learning, NLP, and elastic search, you can get to that detailed analysis much, much faster. Then apply the NLP to assist with existing IP analysis (*e.g.*, claim analysis) and product documentation matching. Machine learning and training in specific technology domains also increases the accuracy and completeness of the results. So applying AI on the back end with training models and the front end with NLP, inference checking and other tools, creates a search and analysis tool that is faster, more accurate, and provides better results.

For example, IBM created a tool called IP Advisor with Watson to perform these kinds of searches based on business use cases: Evidence of Use, Prior Art, Landscaping. Tools like this lead the patent professional to the target and allow further investigation, leveraging

7. U.S Patent & Trademark Office, *Public Views on Artificial Intelligence and Intellectual Property Policy* 19-20 (2020); U.S. Copyright Office, *Compendium of U.S. Copyright Office Practices* §§ 306, 313.2 (3d ed. 2017).

the searcher's experience, to get to the bullseye. Patent interpretation is subjective in the eyes of the examiner, the patent creator, and the patent attorney. So, while AI can provide useful information, it is well recognized that human assessments remain an integral part of the analysis.

III. IP Protections of AI Inventions, Code and Data

Even holding aside questions of AI inventorship and authorship discussed in Part IV below—and acknowledging that both the USPTO and the U.S. Copyright Office (USCO) have asserted that only natural persons can be inventors and authors—the use of AI tools to assist human inventorship and authorship still presents many issues. First, to obtain a patent in the United States, a claimed invention must be novel and non-obvious from the perspective of a hypothetical person of ordinary skill in the art (35 U.S.C. §§ 102, 103). European patents have similar requirements for novelty and inventive step determined from the perspective of a person skilled in the art. But if an invention is generated using AI, where the inner-workings of the AI would not (and likely could not) be obvious to a person with ordinary skill, then it may be easier to establish non-obviousness or inventive step for an AI-generated invention based on the very nature of the AI itself. This potentially could lead to different legal standards for judging non-obviousness and inventive step of AI-generated inventions, as compared with conventional inventions.

Second, under U.S. law, the written description of the invention must be sufficiently clear to enable a person skilled in the art to make and use the invention (35 U.S.C. § 112(a)). Because an invention incorporating AI technology may be highly dependent on particular training-data sets input into the AI algorithm, it is possible that a person skilled in the art would not be able to practice the disclosed invention without further knowledge of the appropriate training data. This leads to the question of whether patent applications including AI-related inventions should be subject to additional requirements for disclosing details related to training data, hyperparameters, learning algorithms, loss functions, etc., and if so, under what circumstances. Some commentators have suggested applying the legal test for undue experimentation (using "Wands factors") to determine when the disclosure of an AI-related invention is sufficient to enable a person skilled in the art to practice the invention.[8]

The legal issues surrounding AI and IP are still evolving as AI technologies mature and become more prevalent. Given these uncertainties, it is unclear if AI technologies eventually may lead to new forms of IP protections, or variations of conventional IP rights, or instead may be eligible for only certain existing IP rights.

IV. AI Inventorship/Authorship

As mentioned above, much popular commentary assumes that Artificial General Intelligence (AGI) machines exist—or will shortly—and that they operate as sentient beings that could exercise autonomous intent and judgment to create something. We have been conditioned to worry about this ever since some of the earliest science fiction posited human-created machines and beings. However philosophically and ethically rich such discussions are, they are still not ripe for where the technology is today. Accordingly, the sense of AI as equivalent to a human inventor or author is not what we will consider in this article.

Instead, the pressing issues today involve use of AI tools and engines to assist in human invention and authorship, or to be *directed* by humans in what to "invent" or "author" in qualified senses of those terms. Thus, the real questions turn on how to allocate invention and authorship, at least for the ultimate purposes of determining whether patents can issue or copyrights can be registered, and to whom.

Both the USPTO and the USCO have asserted that only humans can invent and author.[9] In both cases—invention and authorship—it is not clear that this is a constitutional requirement, however. The U.S. Supreme Court has never squarely ruled on either issue. Rather the USPTO and USCO cite *Trade-Mark Cases,*[10] *Burrow-Giles Lithographic Co. v Sarony,*[11] and *Bleistein v Donaldson Lithographing Co.*[12]—all of which simply discuss human agency for the facts in those cases and do not directly hold that *only* natural persons can be inventors and authors. The USPTO acknowledges that a counterexample to such a categorical prohibition on non-human authors would appear to be the statutory work made for hire (WMFH) provision of the Copyright Act.[13] Yet, it brushes this off as merely a "legal fiction for non-human employers to be authors under certain circumstances."[14] It then adds in a footnote, "The [work made for hire] doctrine applies whether the employer is a human being or a corporation but the *actual creator* of any work protected by copyright has always needed to be human. . .," without giving a citation for this assertion.[15]

8. *See supra* Note 4.

9. *See supra* Note 7.

10. 100 U.S. 82, 94 (1879).

11. 111 U.S. 53, 58 (1884).

12. 188 U.S. 239 (1903).

13. 17 U.S.C. § 201(b).

14. *Public Views on Artificial Intelligence and Intellectual Property Policy, supra* note 7.

15. *Id.* at footnote 108.

Some groups have pushed test cases at patent offices and in the courts regarding AI inventorship. In the United States, for example, Dr. Stephen Thaler of the Artificial Inventor Project filed for two patent applications each naming an AI engine "DABUS" as the sole inventor. The USPTO rejected these claims of inventorship, relying on its rule that only natural persons can be inventors. Dr. Thaler has appealed the USPTO's decisions.[16] He has similarly filed patent applications naming DABUS as the inventor in the United Kingdom and several other jurisdictions.[17] While the United Kingdom Intellectual Property Office also rejected inventorship claims for DABUS, finding a patent inventor must be a person—and this decision was upheld on appeal[18]—other countries have not yet decided the issue.[19]

Even if such groups establish some level of inventorship or authorship for AI, that does not settle the allocation of ownership. Instead, copyright law's WMFH doctrine and patent law's cousin—common law hired to invent doctrine—result in ownership vesting in the employer, whether natural or corporate person. Under copyright WMFH, the employer is the legal author *ab initio*—no transfer, assignment, or conveyance need occur. By contrast, the common law hired to invent doctrine effects an equitable transfer of ownership of the invention to the employer. Any issued patents or pending applications must still be assigned from the natural person inventor to the employer.[20] Accordingly, even if AI has the capacity—and is deemed to have invented or authored something, then provided this was still at the direction of the AI's owner—as "employer"—existing legal rules would almost certainly vest title in that owner-employer.

Only once some AGI is truly operating on its own with no human direction or input will we need to worry about authorship or inventorship. To be clear, this would require: (1) independently deciding there is a problem that needs to be solved, determining how best to solve it, and then solving it (patents); or (2), independently "desiring" to create something of its own "vision," determining how best to realize that vision, and then creating the work (copyright).

However, some argue that an implication of asserted rules against non-human inventorship and authorship is that there is simply neither patentable invention nor a copyrightable work. In such cases, there would then be no patent or copyright and the invention or work, respectively, would be in the public domain. This has primarily been of interest to, and promoted by, IP skeptics who would like to see IP systems weakened as a general matter.

Innovative firms and creative persons will likely accommodate any such concerns by ensuring that they direct whatever AI tools they employ in their work such that the results are covered by WMFH or hired to invent rules. They will have little incentive to push the issue of whether AI can, or has in any particular instance, invented or created, as this could complicate their own inventorship, authorship, or ownership. The only exception will be AI evangelists like Dr. Thaler, or firms who believe they can succeed without IP. It is, however, also possible that defendants in IP infringement cases might seek to show through discovery that the plaintiff firm or human persons were not the actual inventors or authors—either because the AI was or no one was because inventorship and authorship can only vest initially in natural persons. Careful inventors and creators should be able to establish a sufficient record of inventorship and authorship, as applicable, under advice of counsel.

Conclusion

Intellectual property systems around the globe have begun to recognize that it is essential to consider how to address the increasing power and utilization of AI. The number of patent filings pertaining to AI discoveries is skyrocketing across technology areas, as AI is increasingly used as a tool that can facilitate evaluating applications for IP protection and AI technology is increasingly used to contribute to new discoveries and works. To date, AI is a powerful tool that can expand the reach of innovators' discoveries and society's deliberation. Tracking the utilization, capabilities, and asserted capabilities (which may, or may not, be exaggerated) is most advisable for innovative companies and entities engaged with IP systems. The capabilities and utilization of this technology are evolving quickly, and the conversation is often imprecise. Staying informed as to the subtleties and recent advancements in the field of AI can provide much value. ■

Available at Social Science Research Network (SSRN): *https://ssrn.com/abstract=3897892.*

16. *Thaler v. Iancu,* Case No. 1:20-cv-00903-LMB-TCB (E.D. Va., August 6, 2020).

17. *See* Artificial Inventor Project's pending patent applications available at *https://artificialinventor.com/patent-applications/.*

18. [2020] EWHC 2412 (Pat) (21 September 2020).

19. *See supra* note 17.

20. This leads to the odd dichotomy that natural-person creators need not be listed in copyright registrations—a corporate employer can list only itself—while natural-person inventors must always be listed in issued patents—even where they assigned any rights in the patents to the employer in advance. For more detail on this topic, *see* Sean M. O'Connor, "Hired to Invent vs. Work Made For Hire: Resolving the Inconsistency Among Rights of Corporate Personhood, Authorship, and Inventorship," 35 *SEATTLE L. REV.* 1227 (2012).