GENERATIVE AI AND COPYRIGHT PROTECTION: AI DREAMS OF ELECTRIC SHEEP

BY DR. CHRISTIAN E. MAMMEN

& DANIEL GRIGORE

Partner and associate, respectively, Womble Bond Dickinson, San Francisco.

TechREG CHRONICLE NOVEMBER 2023

Visit www.competitionpolicyinternational.com for access to these articles and more!

GENERATIVE AI AND COPYRIGHT PROTECTION: AI DREAMS OF ELECTRIC SHEEP

By Dr. Christian E. Mammen & Daniel Grigore

The IP and tech community has for most of the past decade been captivated by the emergence of newly capable artificial intelligence algorithms and whether their outputs should be protected by existing IP law paradigms (notably concerning patents and copyright). This has been to the forefront of the debate surrounding certain generative AI platforms, such as ChatGPT, Dall-E, and Stable Diffusion, among others. But even more latterly, the focus has shifted: the question has now turned to the IP protection of the inputs necessary for the functioning of such systems. This article draws out the consequences of this shift, with a focus on recent policy developments in the U.S.

01
INTRODUCTION

For much of the past four years, the IP and tech community has been captivated by the emergence of newly capable artificial intelligence algorithms – first DABUS[2] on patents, then DABUS on copyright, and then multiple generative AI platforms, such as ChatGPT, Dall-E, Stable Diffusion, and the like.

Initially, the conversation focused on whether *outputs* generated by an AI algorithm could be protected by patents or copyrights.

For example, last year, the U.S. Court of Appeals for the Federal Circuit affirmed a refusal by the U.S. Patent and Trademark Office ("PTO") to issue two patents to Dr. Stephen Thaler's DABUS, interpreting the Patent Act to require a "natural" person – that is, a human being – as an inventor.[3] More recently, and in the copyright sphere, a district court concluded that Dr. Thaler's AI system, the "Creativity Machine," could not be the copyright holder of a piece of artwork it generated because a piece of art, generated solely by AI with no human input, is not copyrightable.[4] This result should come as no surprise since the Copyright Office itself provided a range of scenarios that explain when an AI-generated work may or may not be copyrightable, and, almost always, requires some type of human involvement in the authorship for a chance at protection.[5]

Consequently, it appears that questions related to outputs have, for the time being, been definitively resolved in the U.S.: AI-generated works, operating without the requisite amount of human involvement, cannot be protected by patents or copyrights.

Globally, however, there are a few exceptions. South Africa has permitted patents to issue on AI-generated inventions.[6] The U.K. permits copyrights (of a shortened duration) for computer-generated works.[7]

More recently, largely since early 2023, the question has turned to IP protection of the *inputs.* It is now well-known that ChatGPT and other generative AI chatbots are trained on large language models ("LLMs") and image-generating AI platforms like Dall-E and Stable Diffusion are likewise trained on massive volumes of data.

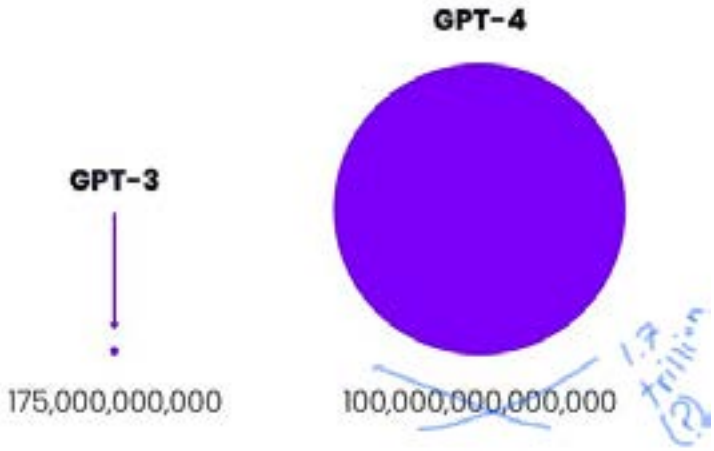

As rumors continue to swirl, it is assumed that GPT-4 runs on about 1.7 to 1.8 trillion parameters, an order of magni-

2 DABUS is an acronym for "device for the autonomous bootstrapping of unified sentience," and is the name for an AI algorithm developed by Dr. Stephen Thaler. Wen Xie, *AI Inventorship: Will Our Patent Laws Stand Up? My Conversation with Dr. Stephen Thaler*, IPWatchdog (May 14, 2023, 12:15 PM), https://ipwatchdog.com/2023/05/14/ai-inventorship-will-patent-laws-stand-conversation-dr-stephen-thaler/id=160832/.

3 *Thaler v. Vidal*, 43 F.4th 1207 (Fed. Cir. 2022), *cert. denied*, __ U.S. __, 143 S. Ct. 1783 (2023).

4 *Thaler v. Perlmutter*, __ F.Supp.3d __, Civil Action No. 22-1564 (BAH), 2023 WL 5333236 (D.D.C. Aug. 18, 2023).

5 U.S. Copyright Office, Copyright and Artificial Intelligence (2023), https://www.copyright.gov/ai; Federal Register, Copyright Registration Guidance: Works Containing Material Generated by Artificial Intelligence (March 16, 2023), https://www.federalregister.gov/documents/2023/03/16/2023-05321/copyright-registration-guidance-works-containing-material-generated-by-artificial-intelligence.

6 Meshandren Naidoo & Dr. Christian E. Mammen, *DABUS Gains Traction: South Africa Becomes First Country to Recognize AI-Invented Patent*, PatentlyO (Aug. 4, 2021), https://patentlyo.com/patent/2021/08/traction-recognize-invented.html.

7 Copyright, Designs and Patents Act of 1988, §§ 178 ("'computer-generated' in relation to a work, means that the work is generated by computer in circumstances such that there is no human author of the work"); 9(3) ("In the case of a literary, dramatic, musical or artistic work which is computer-generated, the author shall be taken to be the person by whom the arrangements necessary for the creation of the work are undertaken."); and 12(7) (allowing computer-generated works to be protected for fifty years from the end of the calendar year in which the work was made).

tude larger than GPT-3.5's already-large 175 billion parameters.[8] [9]

This, of course, has led to an obvious question and an equally obvious answer: Where did all of that training data come from? Essentially, from materials that we humans have put online.

It turns out that lots of that information is subject to copyrights held by humans and corporations, leading to a collective "Can they do that???" Put another way, if Human A can read something that Human B has posted online, and Human A even learns something by reading it, is that really different from Computer Z reading it, and maybe learning something? What if Computer Z reads 1000 texts in the time it takes Human B to read ten? Or *one trillion* texts?[10] This is the guts of the main question underlying a half-dozen cases currently making their way through the U.S. court system.[11]

When copyright disputes arise, a standard defense is that the alleged infringement falls under the fair use exception. A court will balance four factors to determine whether something is a fair use: (1) the purpose and character of the use, (2) the nature of the copyrighted work, (3) the amount or substantiality of the portion used, and (4) the effect of the use on the potential market for or value of the work.[12] The result of this analysis will ultimately be heavily fact-dependent, especially in the AI space, where the threshold question is not simply "does fair use apply" but rather "*should* fair use apply?"

For argument's sake, if the answer to the "should" question is yes, then the first factor will consider whether the use is commercial, or nonprofit, or educational in nature, and whether there is some "transformation" that adds something new to the copyrighted work. Noncommercial uses are more likely to weigh in favor of fair use than those done for profit, so maybe those AI-generated rap lyrics in the style of Danielle Steel are better protected as examples in English class than printed on a novelty t-shirt for sale.

Interestingly, many existing AI platforms are trained on copyrighted works for commercial purposes. Even OpenAI, which was originally founded as a nonprofit,[13] launched a subscription model for its AI product, ChatGPT. Going forward, AI developers should keep in mind that even if a system is not *currently* being used for commercial purposes, *future plans* to enter the market could be enough to undermine fair use protection.

The second factor considers the nature of the copyrighted work being infringed with an eye toward how creative it is: the more creative the art, music, or literature that makes up the training data, the more likely the generated work may be considered infringing. While AI need not be trained solely on dictionaries and encyclopedias, fair use may be more attainable that way.

"Interestingly, many existing AI platforms are trained on copyrighted works for commercial purposes

Substantiality serves as the core of the third factor and measures the portion copied in comparison to the copyrighted work as a whole. Something to keep in mind when weighing this factor: it's not how much ends up in the final, generated work that matters, it is what proportion of the copyrighted work was *ingested*. To argue that the copyrighted work only had a minimal effect on the ultimate outcome, as is sometimes asserted, flips the third factor on its head. If an AI system ingests *all* or *substantially all* of the copyrighted work, it does not matter how much of that copyrighted work appears in the final product – the AI trained on the *whole* or *almost whole* work, regardless of the output. In sum, to meet the fair use requirement, an alleged infringer may not

8 Parameters determine how a neural network turns input data into output data and are learned during the training process. The more parameters, the more complex and expressive the model can be. Vitalii Shevchuk, *GPT-4 Parameters Explained: Everything You Need to Know*, Level Up Coding (May 17, 2023), https://levelup.gitconnected.com/gpt-4-parameters-explained-everything-you-need-to-know-e210c-20576ca#:~:text=The%20more%20parameters%20a%20model,4%20has%201.7%20trillion%20parameters.

9 Maximilian Schreiner, *GPT-4 architecture, datasets, costs and more leaked*, The Decoder (July 11, 2023), https://the-decoder.com/gpt-4-architecture-datasets-costs-and-more-leaked/.

10 It is estimated that GPT-4 has about 1.8 trillion parameters across 120 layers. *Id*.; Matt Popovic, *ChatGPT Parameters Explained: A Deep Dive into the World of NLP*, ecoAGI (June 4, 2023), https://ecoagi.ai/articles/chatgpt-parameters.

11 See e.g. *Getty Images (US), Inc. v. Stability AI, Inc.*, No. 1:23-cv-00135-GBW (D. Del. filed Feb. 3, 2023); *Tremblay et al v. OpenAI, Inc., et al*, No. 3:23-cv-03223-AMO (N.D. Cal. filed June 28, 2023*); Silverman et al v. OpenAI, Inc. et al*, No. 3:23-cv-03416-AMO (N.D. Cal. filed July 7, 2023).

12 U.S. Copyright Office, U.S Copyright Office Fair Use Index (2023), https://www.copyright.gov/fair-use/.

13 Chinecherem Nduka, *How OpenAI Transitioned form a Nonprofit to a $29B For-Profit Company*, Hackernoon (March 27, 2023), https://hackernoon.com/how-openai-transitioned-from-a-nonprofit-to-a-$29b-for-profit-company.

take more than what is necessary to achieve the transformative purpose, and cannot gobble whole the copyrighted work. This factor's premise serves as a major portion of the lawsuit Getty Images brought against Stability AI, alleging infringement in the training by using millions of Getty pictures, complete with Getty's watermark.[14] Is that amount of training data "only what is necessary?" We shall have to wait and see.

Lastly, the fourth factor takes into account the allegedly infringing work's effect on the market value of the copyrighted work and argues against fair use in cases where the infringing work may act as a stand-in or substitute for the original work. In the AI sphere, generative AI systems may very well destroy the market (e.g. sale or license) for copyrighted works, first, by flooding the market with the sheer volume of outputs, resulting in oversupply and a price collapse, and second, by not compensating copyright owners for the already-available works used as training data. Even though many copyright owners offer AI training licenses, few AI developers, if any, pay copyright owners for their works. This has the potential to severely curtail the quality of training input and, if recent media coverage is an indication, can make a lot of people just plain angry.

It is this last factor that suggests the use of some sort of buffer between the copyright owner and AI dataset. With such a proliferation of AI-infused web spaces, it may make sense for copyright holders, who have works on the internet that they want to protect, to receive an opt in/opt out opportunity. In theory, it sounds good, but in practice, how does one effectively patrol the internet on a global scale?

Consider this as a possible analogy – your data online is like photos of you taken by street cameras (and other information about you) as you move around a city. An opt in system would say that people collecting those photos can only collect if you opt in. But who would you tell? How would the many surveillance cameras on the street be able to determine who had opted in and who hadn't? Who gets to control the surveillance cameras and what happens to that information after an AI system has trained on it?

Similar practical problem arises for opting out: how would the camera operators be able to determine which individuals' images should be masked out from the street cameras? Moreover, how would one determine the logistics for opting out? At what stage of AI training should someone have the choice: after collection but before the AI model is trained? After training? If after training, how can opted-out material, and its influence, be excised from a trained data model when it has all been mixed together like a box of melted crayons? And to whom does this opt-out request get forwarded – does it go platform by platform, or to some central repository?[15]

"Consider this as a possible analogy – your data online is like photos of you taken by street cameras (and other information about you) as you move around a city.

Weighing the two, is opting in better than opting out? What type of person would opt in? Dr. Thaler, perhaps. Yet, if one million Dr. Thalers opted into data collection, this could lead to remarkably skewed training datasets that lacked the diversity provided by all the non-Dr. Thaler types who choose to remain on the sidelines.

To add another, legal, layer, does the mere fact that we are having a conversation about opting in or opting out an admission that fair use *doesn't* apply to AI-generated works? Is it an admission that it *shouldn't* apply? Should individuals be the ones making that call? Legislators? Big tech?

The reader may notice way more questions in the preceding paragraphs than answers. That is because, at this point, there is truly no "right" answer or easy solution. On this precipice of AI/copyright relations, it's truly anyone's best guess as to what will work or fail. To be safe, rather than sorry, a number of major brands, like the New York Times and Reddit, have made it clear that they do not want Google, Microsoft, and OpenAI to profit off of their content.

14 Blake Brittain, *Getty Images lawsuit says Stability AI misused photos to train AI*, Reuters (Feb. 6, 2023, 9:23 AM), https://www.reuters.com/legal/getty-images-lawsuit-says-stability-ai-misused-photos-train-ai-2023-02-06/.

15 While few AI developers have entered either field of opting in or opting out, trial runs of the concept may yield informative results: OpenAI's latest version of DALL-E allegedly allows artists to opt *out* of their content being used to train future generations of AI models and will also reject requests to make an image in the style of a living artist or portray public figures. *See* Kyle Wiggers, *OpenAI unveils DALL-E 3, allows artists to opt out of training*, TechCrunch (Sept. 20, 2023, 10:57 am), https://techcrunch.com/2023/09/20/openai-unveils-dall-e-3-allows-artists-to-opt-out-of-training/.

Both have updated their terms of service/use to include AI-specific language.[16]

Hoping to offer some guidance for AI use of the future, the White House published a "blueprint" for an AI bill of rights with the hopes of creating a safe and equitable user experience.[17] The White House based this AI blueprint on five principles, and the fifth, "Human Alternatives, Consideration, and Fallback," is most relevant here.

This principle recommends a human alternative to AI systems so that users may have an accessible, and human-based, remedy for AI system failures, errors, or appeals of the impact AI has on the user. In this way, and going beyond copyright, all users may completely forego an automated system in order to access a human alternative – in essence, a *complete* opt out of AI interaction, at least "where appropriate."[18]

"Where appropriate" is a concept that remains a little vague, but the White House makes a point of suggesting that having a human alternative makes a lot of sense in sensitive domains like the criminal justice system, employment, education, and healthcare, where, "absent appropriate safeguards," technology may lead to unfair, inaccurate, or dangerous outcomes.[19] Furthermore, a human alternative would be important in situations where automated systems fail during time-critical situations (think an automated door-opening system that fails during a building fire).

This all comes full circle: the White House's human alternative principle promotes human/AI teamwork that both the PTO and Copyright Office require for IP protection. Which makes sense, because, if androids dream of electric sheep, AI functions best when its dreams are inspired and guided by the incredible sheep/oeuvre of human creativity and ingenuity. ■

"This all comes full circle: the White House's human alternative principle promotes human/AI teamwork that both the PTO and Copyright Office require for IP protection

16 Jesse Weatherbed, *The New York Times prohibits using its content to train AI models*, The Verge (Aug. 14, 2023, 3:26 AM), https://www.theverge.com/2023/8/14/23831109/the-new-york-times-ai-web-scraping-rules-terms-of-service; Danny Goodwin, *New York Times: Don't use our content to train AI systems*, Search Engine Land (Aug. 10, 2023, 9:47 AM), https://searchengineland.com/new-york-times-content-train-ai-systems-430556; akhudek (u/ akhudek), Reddit (April 18, 2023, 3:56 PM), https://www.reddit.com/r/MachineLearning/comments/12r7qi7/d_new_reddit_api_terms_effectively_bans_all_use/.

17 The White House Office of Science and Technology, Blueprint for an AI Bill of Rights (Oct. 2022), https://www.whitehouse.gov/ostp/ai-bill-of-rights/.

18 The White House Office of Science and Technology, Human Alternatives, Consideration, and Fallback (Oct. 2022), https://www.whitehouse.gov/ostp/ai-bill-of-rights/human-alternatives-consideration-and-fallback/.

19 *Id*.

CPI SUBSCRIPTIONS

CPI reaches more than **35,000 readers** in over **150 countries** every day. Our online library houses over **23,000 papers**, articles and interviews.

Visit **competitionpolicyinternational.com** today to see our available plans and join CPI's global community of antitrust experts.

CPI COMPETITION POLICY® INTERNATIONAL