COPYRIGHT INFRINGEMENT OF AI-GENERATED OUTPUT

Danny Friedmann[*]

ABSTRACT

Where a generative AI memorizes copyrighted content from its training data and generates output substantially similar to these copyrighted works, this 'memorization' can lead to copyright infringement. Developers of large language models (LLMs) and deployers of generative AI (gAI) services could be held liable for this symmetry between input and output, if there is no fair use defense, or exception or limitation for unauthorized but permissible use. *Ultraman*, was the first case in China where AI-generated output was qualified as copyright infringement, and ongoing cases such as *New York Times v OpenAI* in the U.S., are slowly beginning to demonstrate the contours of these copyright infringement cases. The doctrine of copyright infringement seems ill-equipped to address the challenge of protecting the style of copyrighted works unless there is substantial similarity. The same can be said about deepfakes and voice cloning. As a result, there appears to be thin protection for copyrighted styles and likeness in voice or image under current law. These legal challenges highlight the need for more robust frameworks to address the complexities introduced by advanced AI technologies in case of AI-generated output substantially similar to copyrighted works.

KEY WORDS

memorization, copyright infringement, New York Times v OpenAI, style, Snoopy problem, Italian plumber problem

[*] Associate Professor of Law, Dr., Peking University School of Transnational Law, Shenzhen, China.

Introduction

In the U.S.,[1] and in the Czech Republic,[2] AI-generated images could, so far, not be protected, in contrast to China.[3] This does not mean that AI-generated output cannot constitute copyright infringement. If the output is substantially similar to a copyrighted work, the deployer of the gAI can be held liable for copyright infringement, in the absence of fair use in the U.S., or any exceptions or limitations in the EU, or any of both in China. Lemley argued that there can be copyright (and thus copyright infringement) in the prompt itself.[4]

One can argue that AI is an algorithmic mimicry of human capacities.[5] The idea is that during the training process, the LLM makes statistical inferences of the data in order to generate new and never seen or heard before results. But when there is mimicry in the last part of the 'AI-supply chain,'[6] so that the generated output is too similar to the input, this can lead to copyright infringement.

Italian plumber problem

Generative AI does not generate content *ex nihilo*. It learned about copyrighted works in the training data and sometimes memorized it a bit too much, so that the output of generative AI (gAI) is substantially similar to that copyrighted work in the training data. This 'memorization,' can cause copyright infringement. According to LLM developers and AI-service deployers, 'regurgitation' is unintentional, a bug and not a feature.[7] Memorization in LLMs and the generation of content that closely resembles the expressive works in their training data can be attributed to several factors. By generating outputs that mimic the input in its training data, the LLM reflects the machine-learning phenomenon known as 'overfitting,' which happens when it is too closely adapted to the data on which it was trained, making it difficult for the model to generalize to new data sets. LLMs with billions of parameters have a greater capacity to store information. While this enables them to learn complex patterns and generate high-quality content, it also increases the risk of memorizing and generating specific content from the training data in the output. The chance this happens increases if these works appear frequently across the dataset. This is what Sag identified as the 'Snoopy problem,'[8] and what Lee and Grimmelmann called 'the Italian plumber problem,'[9] referring to users of gAI that use this prompt for generating images of the game character 'Super Mario.' This problem can be avoided or at least mitigated by de-duplicating the training data, and increasing the diversity of the training data.

The attention mechanism in transformer models allows the model to focus on specific parts of the input data. While this is useful for capturing context and generating coherent content, it can also contribute to memorization if the model learns to pay special attention to particular sequences it encountered frequently during training.

Substantial similarity in times of gAI

The substantial similarity analysis can be determined by identifying the protected elements through a process of filtration, which involves identifying original works, separating ideas from the expression of ideas, and excluding *scènes à faire*. Following this, the selection and arrangement of protected elements that are substantially similar due to access to the original work or works remains relevant.

So far, most AI-service deployers have been unwilling to provide transparency about the training data used. Dividing the training process of the data among different companies, can complicate both the transparency in the training data and any questions about liability regarding copyright infringement. When the relics of the 'Getty' trademark were generated together with images, as was the case in *Getty Images (US), Inc. v. Stability AI Ltd*,[10] or producer tags were generated together with music files ('CashMoneyAP') as in *UMG Recordings v. Suno*,[11] access seems irrefutable. In *Skidmore v. Zeppelin*,[12] the Ninth Circuit abandoned the inverse ratio rule between the degree of access and similarity, but some courts might still apply it.

Fair use or Exceptions and Limitations

Not all unauthorized outputs that are substantially similar, are impermissible. It may still be legally permissible if these acts meet the criteria for fair use,[13] or exceptions and limitations. However, *Warhol Foundation v. Goldsmith*[14] appears to have shifted the Supreme Court's emphasis of the first fair use factor analysis to exclude uses that substantially share the same purpose as the original work, and therefore cannot be qualified as transformative. Content that is generated for the same purpose as the original expressive works, can also substitute for these original works. Thus, the fourth fair use factor that focuses on the market effect will likely weigh in favour of the plaintiff.

New York Times v. OpenAI,[15] is deemed to become an influential case about copyright infringement by AI-generated content in the output. OpenAI's ChatGPT allegedly generated responses to users' prompts that contained verbatim excerpts and detailed summaries of NYT articles. OpenAI made

allegations that the NYT prompt hacked the results that were nearly identical to the original texts of the NYT. OpenAI filed a motion to dismiss parts of the NYT's lawsuit, suggesting that the examples provided in "Exhibit J" were the result of extensive efforts to elicit extraordinary specific responses from ChatGPT,[16] namely providing short snippets of the beginning of NYT articles as prompts.[17] OpenAI demanded that NYT makes clear how it generated the results of Exhibit J.[18]

In another ongoing case, *Authors Guild v. OpenAI*,[19] the plaintiffs alleged that OpenAI used their copyrighted books to train its LLMs, resulting in AI-generated texts that may include substantial portions of the copyrighted material. They argued that the AI-generated output can be seen as infringements of the right to prepare derivative works. In *Thomson Reuters v. ROSS Intelligence Inc.*,[20] the plaintiffs alleged that ROSS Intelligence used content from the Westlaw database to train an AI-based legal research tool. The plaintiff claimed that both the use of its copyrighted materials for training and the resulting AI-generated outputs, which might contain substantial portions of its copyrighted content, including "headnotes," constitute copyright infringement.

"Ultraman"

Most notable, in China, the Guangzhou Internet Court decided the first case on copyright infringement of AI-generated output on February 8, 2024:[21] the Japanese copyright holder Tsuburaya Productions Co., Ltd., of the fictional character 'Ultraman,' granted Shanghai Character License Administrative Co., Ltd. (SCLA) an exclusive license for these works in China. This license included the rights of reproduction, preparation of derivative works, and enforcement. Additionally, Tsuburaya has registered the 'Ultraman' series images with the copyright office of the National Copyright Administration of China. SCLA sued a provider of gAI services in China for copyright infringement (reproduction, adaptation, and information network dissemination) by training on 'Ultraman' works and generating substantial similar works. After the defendant received the complaint, he implemented prompt filtering. However, when 'Diga' was used as a prompt, substantial similar images of 'Ultraman' were still being generated. The Guangzhou Internet Court held that these were infringements of the rights of reproduction, adaptation and dissemination, and the defendant had to immediately cease its infringing activities and implement corresponding technical measures to prevent the generation of infringing images. The defendant was also ordered to compensate RMB 10,000 to the plaintiff.

Direct, vicarious and contributory liability of copyright infringement

In the litigation against AI service deployers, next to claims of direct infringements, claims of vicarious and contributory infringements are made.

Vicarious infringement examines whether the AI service provider profited from and had control over the infringing activities. In *Getty Images v Stability AI*,[22] the plaintiff argued that the defendants profited from infringing activities (subscription fees of users that can create Getty-like stock photos) and had the ability to control the use of its platform but did not take adequate measures to prevent the infringement. In *Richard Kadrey v Meta Platforms, Inc.*,[23] the claim was dismissed because the plaintiffs did not adequately allege that any outputs generated by Meta's LLM named 'LLaMA' contained protectable expression from their books; and thus there was no vicarious copyright infringement.

Contributory infringement considers whether the AI service provider knowingly facilitated or contributed to the infringement by users. In *UMG v Anthropic*,[24] the plaintiffs alleged that Anthropic's AI model 'Claude,' facilitated the generation of lyrics and musical compositions that were substantially similar to copyrighted songs owned by UMG *et al*. The plaintiffs alleged that Anthropic was aware of these infringing activities and materially contributed to them by providing the AI tool and failing to implement sufficient safeguards.

Alignment to Copyright Law

To further align the AI, a filter can be added to the input field of the AI, so that some copyright inducing prompts can be blocked. In addition, a filter can be placed to the output of the AI. But these output filters need to take fair use or exceptions and limitations into account, such as caricatures, parody, pastiche, quotation and *de minimis* use. The number of characters that can be used in quotations, could be limited, as for example the 160 characters threshold used in Section 10 of the German Copyright Service Provider Act,[25] and also provides thresholds of 15 seconds for video and audio, and 125 kilobytes of an image. Alternatively, the LLM developers and AI deployers could license the training data to avoid any potential copyright infringement.

A DMCA-inspired notice-and-takedown regime cannot remove copyright infringing works from the LLM, since these are not stored as such in the LLM, as it only has the statistical patterns derived from these works in its weights. Instead of retraining the LLM every time a memorization is found, the output filter could be fine-tuned, which is much more cost and time efficient.

Personality rights

A special variation on the theme of copyright infringement is when the AI generates a person's unique voice, image or video without permission. This could lead to a violation of the publicity or personality rights (whether this was the question was in dispute between Scarlett Johansson and OpenAI).

On April 23, 2024, the Beijing Internet Court determined the first AI-generated voice personality rights infringement case in China.[26] One of the defendants, a Beijing cultural media company owned the copyright of audio works spoken by the plaintiff, the voice actor Yin. However, the court held that the copyright licensing arrangements did not include the rights to use the claimant's voice for training the AI and redeployment by AI without the individual's informed consent.

Signature style

The current doctrine of copyright infringement appears inadequate for addressing the challenge of protecting the signature style of copyrighted works unless there is a demonstrable substantial similarity between the original and the alleged infringing work. This inadequacy stems from the doctrine's focus on literal copying or near-identical reproduction, which leaves little room for safeguarding the distinct stylistic elements that characterize an author's work. Consequently, the protection afforded to styles under existing law is minimal, exposing such creative expressions to potential misappropriation without sufficient legal recourse. This gap underscores the necessity for evolving legal frameworks that can more effectively protect the nuanced aspects of artistic and literary styles in the face of modern technological advancements.

[1] Friedmann, Danny, 'Creation and Generation Copyright Standards' (2024). NYU JIPEL (forthcoming) https://ssrn.com/abstract=4770924.
[2] Municipal Court of Prague, 11 October 2023, AI-generated image "shaking hands over a contract". https://justice.cz/documents/14569/1865919/10C_13_2023_10/108cad3e-d9e8-454f-bfac-d58e1253c83a. Cerri, Alessandro, 'Czech court finds that AI tool DALL-E cannot be the author of a copyright work,' IP Kat, April 15, 2024, https://ipkitten.blogspot.com/2024/04/czech-court-finds-that-ai-tool-dall-e.html.
[3] *Supra* n. 2.
[4] Lemley, Mark A., 'How Generative AI Turns Copyright Upside Down', 25(2) COLUM. SCI. & TECH. L. REV. 25.
[5] Jaeger, Johannes, 'Artificial intelligence is algorithmic mimicry: why artificial "agents" are not (and won't be) proper agents,' arXiv:2307.07515 [cs.AI] 22 Feb 2024.
[6] Lee, Katherine, A. Feder Cooper, James Grimmelmann, 'Talkin' 'Bout AI Generation: Copyright and the Generative-AI Supply Chain,' (2023) J COPYRIGHT SOC'Y USA (forthcoming), https://arxiv.org/abs/2309.08133.
[7] 'OpenAI and journalism,' OpenAI Blog, Jan. 8, 2024, https://openai.com/index/openai-and-journalism/.
[8] Sag, Matthew, 'Copyright Safety for Generative AI,' (2023) 61 HOUS. L. REV. 295, https://houstonlawreview.org/article/92126-copyright-safety-for-generative-ai.
[9] Lee, Timothy B., and James Grimmelmann, 'Why The New York Times might win its copyright lawsuit against OpenAI,' Ars Technica, Feb. 20, 2024, https://arstechnica.com/tech-policy/2024/02/why-the-new-york-times-might-win-its-copyright-lawsuit-against-openai/3/.
[10] *Getty Images (US), Inc. v. Stability AI Ltd*, 1:23-cv-00135-UNA (D. Del., Feb. 3, 2023) https://cases.justia.com/federal/district-courts/delaware/dedce/1:2023cv00135/81407/1/0.pdf?ts=1682222407.
[11] *UMG Recordings v. Suno*, *Inc. et al*, 1:2024-cv-11611 (D. Mass., June 24, 2024) https://www.riaa.com/wp-content/uploads/2024/06/Suno-complaint-file-stamped20.pdf.
[12] *Skidmore v. Zeppelin*, 952 F.3d 1051 (9th Cir. 2020).
[13] 17 U.S.C. § 107.
[14] *Andy Warhol Found. for the Visual Arts, Inc. v. Goldsmith*, 598 U.S. 508, 528 (2023).
[15] *New York Times v Microsoft and OpenAI*, 1:23-cv-11195 (S.D.N.Y., Dec. 27, 2023) https://nytco-assets.nytimes.com/2023/12/NYT_Complaint_Dec2023.pdf.
[16] Edward Lee, 'Exhibit J to New York Times Complaint provides "One Hundred Examples of GPT-4 Memorizing Content from the New York Times,"' ChatGPT is Eating The World, Dec. 27, 2023, https://chatgptiseatingtheworld.com/2023/12/27/exhibit-j-to-new-york-times-complaint-provides-one-hundred-examples-of-gpt-4-memorizing-content-from-the-new-york-times/.
[17] Edward Lee, 'How did the New York Times figure out the glitch in GPT-4 for Exhibit J in its lawsuit v. OpenAI?' ChatGPT is Eating The World,' Dec. 28, 2003, https://chatgptiseatingtheworld.com/2023/12/28/how-did-the-new-york-times-figure-out-the-glitch-in-gpt-4-for-exhibit-j/.
[18] OpenAI Defendants' Notice of Supplemental Authority in Support of Motion to Compel, Case *NYT v Microsoft*, No. 1:23-cv-11195-SHS-OTW (S.D.N.Y., June 25, 2024) https://storage.courtlistener.com/recap/gov.uscourts.nysd.612697/gov.uscourts.nysd.612697.149.0.pdf.
[19] *Authors Guild v. OpenAI*, Class Action Complaint, Jury Trial Demand, 1:23-cv-8282-SHS (S.D.N.Y., Dec. 4, 2023) https://authorsguild.org/app/uploads/2023/12/Authors-Guild-OpenAI-Microsoft-Class-Action-Complaint-Dec-2023.pdf.
[20] *Thomson Reuters Enterprise Centre GmbH et al v. ROSS Intelligence Inc.*, No. 1:2020cv00613 - Document 547 (D. Del. Sept, 25, 2023) https://law.justia.com/cases/federal/district-courts/delaware/dedce/1:2020cv00613/72109/547/.
[21] (2024) 粤 0192 民初 113 号 – Guangzhou Internet Court. English translation by Jiaying Zhang and Yuqian Wang under the supervision of Robert Brauneis, https://ia600409.us.archive.org/3/items/scla-v-ai-company-guangzhou-internet-court-02082024-with-english-translation/SCLA%20v%20AI%20Company%20Guangzhou%20Internet%20Court%2002082024%20with%20English%20Translation.pdf.
[22] *Supra* n. 9.
[23] *Richard Kadrey v Meta Platforms, Inc.*, 3:23-cv-03417 (N.D. Cal., July 7, 2023) https://llmlitigation.com/pdf/03417/kadrey-meta-complaint.pdf.

[24] *Concord Music Group, Inc. v. Anthropic PBC*, No. 3:23-cv-01092, (Oct. 18, 2023, M.D. Tenn.) https://storage.courtlistener.com/recap/gov.uscourts.tnmd.96652/gov.uscourts.tnmd.96652.1.0.pdf.
[25] Act on the Copyright Liability of Online Content Sharing Service Providers of 31 May 2021 (Federal Law Gazette I, p. 1204, 1215), https://www.gesetze-im-internet.de/englisch_urhdag/englisch_urhdag.html.
[26] Tian Han, The first instance verdict of the country's first AI-generated voice personality rights infringement case, Beijing Internet Court, April 23, 2024 (Chinese), https://mp.weixin.qq.com/s/_GxGaG6Q2NYHJWQuOtMyrQ?from=industrynews&version=4.1.20.6024&platform=win.