U.S. Department of Justice

Criminal Division

Computer Crime and Intellectual Property Section *Washington, D.C. 20530*

April 15, 2024

Suzanne V. Wilson
General Counsel and Associate Register of Copyrights
Library of Congress
Copyright Office
101 Independence Avenue, SE
Washington, DC 20559-6000

Dear Ms. Wilson:

The Computer Crime and Intellectual Property Section (CCIPS) of the U.S. Department of Justice's Criminal Division appreciates the opportunity to comment as part of the Copyright Office's ninth triennial rulemaking proceeding under the Digital Millennium Copyright Act (DMCA).[1]

As CCIPS has noted in prior filings in the triennial rulemaking processes, the Department of Justice is responsible for criminal enforcement of a range of statutes protecting intellectual property, including enforcement of the DMCA's criminal provision (17 U.S.C. § 1204), which apply to willful violations of the anti-circumvention provisions in Section 1201 committed for purposes of commercial advantage or private financial gain. The Department of Justice also has an interest in how artificial intelligence and related systems are used across a variety of its work, from its impact on privacy to its potential use in ways that could adversely affect civil rights.

As part of the Department's Criminal Division, CCIPS advises other federal prosecutors on the application of the DMCA and has brought criminal DMCA charges in a variety of cases involving circumvention of technological protection measures and trafficking in devices designed or marketed for use in circumvention. CCIPS is also responsible for prosecuting federal crimes involving unauthorized access to computers, damage to information systems, and other offenses under the Computer Fraud and Abuse Act (CFAA), among other statutes.

[1] Library of Congress, Copyright Office, Notice of Proposed Rulemaking, Exemptions to Permit Circumventions of Access Controls on Copyrighted Works, 88 Fed. Reg. 72,013, 72,027 (Oct. 19, 2023) ("NPRM"). We note that this year, the Department of Justice's Antitrust Division has also filed separate comments (jointly with the Federal Trade Commission) related to other proposals in the DMCA rulemaking process. While we concur with our Department of Justice colleagues' comments, we refer to our comments below and our previous comments as from "CCIPS" to avoid confusion with the Antitrust Division's separate letter.

In comments submitted in previous rulemaking rounds, CCIPS has underscored the importance of security research in identifying vulnerabilities in computer systems and has expressed its support for proposals that effectively expanded the DMCA exemption for good faith computer security research by removing and clarifying certain limitations. In the seventh triennial rulemaking process in 2018, CCIPS expressed support for a limited expansion of the exemption for computer security research effective at the time.[2] In the eighth triennial process in 2021, CCIPS submitted comments supporting renewal of the research exemption, as well as certain modifications of the requirement, such as replacing the requirement that a circumvention not violate "any applicable law" in order to qualify for the exemption with a clarification that the research exemption is not a safe harbor from, or defense to, liability under other applicable laws.[3] CCIPS reiterates its recognition of the importance of good faith computer security research, and again recommends renewal of the research exemption.

This year, one petition, submitted by Jonathan Weiss of Chinnu, Inc., proposes a new exemption for "Security Research Pertaining to Generative AI Bias."[4] Mr. Weiss notes the rapid advancement of generative AI models and concerns that such models may perpetuate or exacerbate systemic biases related to race, gender, and ethnicity, and similar factors. He proposes a new exemption for circumvention of technological measures that control access to copyrighted generative AI models, solely for the purpose of researching biases, as well as for the "sharing of research findings, techniques, and methodologies that expose and address biases in these AI models." Mr. Weiss also suggests limitations on the exemption to prevent misuse, including requirements of "no malicious intent," prioritization of data privacy, and collaboration with stakeholders.

Public comments on Mr. Weiss's proposal for an exemption for generative AI bias research from the Hacking Policy Council (HPC),[5] OpenPolicy,[6] HackerOne,[7] and an

[2] Letter from John T. Lynch to Regan Smith, June 28, 2018 (available at https://www.copyright.gov/1201/2018/USCO-letters/USDOJ_Letter_to_USCO.pdf).

[3] Letter from John T. Lynch to Regan Smith, February 9, 2021 (available at https://www.copyright.gov/1201/2021/comments/reply/Class%2013_Reply_Department%20of%20Justice.pdf).

[4] Jonathan Weiss, Petition for New Exemption Under 17 USC 1201, Copyright Office, 9th Triennial Rulemaking, https://www.copyright.gov/1201/2024/petitions/proposed/New-Pet-Jonathan-Weiss.pdf (last accessed Mar. 26, 2024).

[5] Harley Geiger, Hacking Policy Council comments – Ninth Triennial Proceeding, Class 4, Dec. 21, 2023 (available at https://www.copyright.gov/1201/2024/comments/Class%204%20-%20Initial%20Comments%20-%20Hacking%20Policy%20Council.pdf).

[6] Dr. Amit Elazari, OpenPolicy comments to the Copyright Office concerning the Ninth Triennial Proceeding on section 1201 exemptions, Class Four, Dec. 27., 2023 (available at: https://www.copyright.gov/1201/2024/comments/Class%204%20-%20Initial%20Comments%20-%20OpenPolicy.pdf).

[7] Ilona Cohen, HackerOne Short Comment to US Copyright Office, Ninth Triennial Section 1201 Proceeding (2024), Dec. 26, 2023 (available at: https://www.regulations.gov/comment/COLC-2023-0004-0058).

interdisciplinary group of academic researchers[8] have expressed support for Mr. Weiss' specific proposal while also suggesting that an exemption for AI research should be expanded in various ways. For example, HPC recommends that an exemption for research on generative AI systems should not be confined to research on "bias," but should extend to research on discrimination and other harmful or undesirable outputs in AI systems."[9] OpenPolicy supports a similar research exemption but suggests that it should cover not only "generative AI" but other types of artificial intelligence as well, citing the statutory definition of "artificial intelligence" found in 15 USC § 9401(3).[10] Similarly, the comments from the group of academic researchers recommended an exemption that covers more than generative AI, because "most AI systems are not generative and would benefit from further trustworthiness research by independent researchers."[11] CCIPS agrees that limiting any applicable exemption to security research to "generative AI" risks confusion, especially given that "artificial intelligence" is used to describe a number of related fields of research and engineering. Accordingly, we support the exemption applying to forms of AI research beyond "generative AI."

CCIPS recognizes that the widespread adoption and deployment of AI in various applications has the capacity to influence decision-making and affect behavior on a large scale, in ways ranging from beneficial to harmful or unlawful. Independent research on the functioning and security of AI systems, often called AI "red-teaming"—including research into the generation of outputs that perpetuate or exacerbate bias and discrimination,[12] the generation of outputs that result in or encourage unlawful conduct or harm, and the vulnerability of AI systems to manipulation and misuse—will likely be essential to ensuring the integrity and safety of AI systems.[13] In much the same way that computer security research has helped protect the integrity

[8] Comments from Researchers Affiliated with MIT, Princeton Center for Information Technology Policy, and Stanford Center for Research on Foundation Models, Ninth Triennial Proceeding, Class 4, Mar. 19, 2024 (available at: https://www.copyright.gov/1201/2024/comments/reply/Class%204%20-%20Reply%20-%20Kevin%20Klyman%20et%20al.%20(Joint%20Academic%20Researchers).pdf).

[9] Hacking Policy Council comments at 2, 4. CCIPS recognizes that the potential harms that may result from the use of AI will be heavily dependent on the context of the use and may involve intervening factors beyond the technical "output" of an AI system, but for purposes of this letter, will borrow the HPC's use of the term "outputs."

[10] OpenPolicy comments at 3–4.

[11] Academic Researchers at 2.

[12] See *United States v. Meta Platforms, Inc., f/k/a Facebook, Inc.*, 1:22-cv-5187 (S.D.N.Y.), where the Department of Justice settled its complaint alleging that Meta's housing advertising system discriminates against Facebook users based on their race, color, religion, sex, disability, familial status, and national origin, in violation of the Fair Housing Act.

[13] "Artificial Intelligence red-teaming is most often performed by dedicated 'red teams' that adopt adversarial methods to identify flaws and vulnerabilities, such as harmful or discriminatory outputs from an AI system, unforeseen or undesirable system behaviors, limitations, or potential risks associated with the misuse of the system." Executive Order No. 14110, Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence, 88 Fed. Reg. 75191 (2023) at Sec.

of computer systems and networks on which the public rely, good-faith research into bias and other potentially harmful outputs in AI models can serve a similar, critical, role. For example, good faith research can help reveal unintended or undisclosed collection or exposure of sensitive personal data, or identify systems whose operations or outputs are unsafe, inaccurate, or ineffective for the uses for which they are intended or marketed by developers, or employed by end users. Such research can be especially significant when AI platforms are used for particularly important purposes, where unintended, inaccurate, or unpredictable AI output can result in serious harm to individuals.

CCIPS recognizes that research on the vulnerabilities and other functions of AI system can be done for malicious purposes, or with disregard to risks to public safety, privacy, or security, and so we emphasize that *good faith* in the conduct of such research is meaningful, and good faith AI research should incorporate, among other things, measures to minimize such risks. As we have noted in the context of bad-faith computer security research, unauthorized access to AI systems or exploitation of vulnerabilities in such systems for malicious purposes may be subject to criminal prosecution under a range of statutes. Like good-faith computer security research that might circumvent technological protection measures used to protect copyrighted works but does not result and is not intended to result in infringement, CCIPS believes that good faith research on potentially harmful outputs of AI and similar algorithmic systems should be similarly exempted from the DMCA's circumvention provisions.

While the existing exemption for computer security research covers many types of research focused on the security and integrity of AI models, we recognize that it may not be sufficiently broad in its current form to exempt research that falls outside of "security" concerns. Therefore, we recommend that the Copyright Office consider clarifying the existing exemption to ensure its application to good-faith security research regarding AI systems and other, similar, algorithmic models, but also consider how best to clarify or amend the existing exemptions to cover good-faith research into bias and other harmful and unlawful outputs of such systems.

CCIPS appreciates the Copyright Office's efforts to address these issues as part of the ninth triennial rulemaking proceeding, and for the opportunity to submit these comments.

Sincerely,

John T. Lynch, Jr.

John T. Lynch, Jr.
Chief

3(d), https://www.whitehouse.gov/briefing-room/presidential-actions/2023/10/30/executive-orderon-the-safe-secure-and-trustworthy-development-and-use-of-artificial-intelligence/