**December 22, 2023**

**OpenPolicy comments to the Copyright Office concerning the Ninth Triennial Proceeding on section 1201 exemptions, Class Four**

Item A. Commenter Information.

OpenPolicy appreciates the opportunity to provide commentary concerning the Ninth Triennial Proceeding 1201 exemptions, class 4.

*OpenPolicy*[1] is the world's first policy intelligence and engagement technology platform, aiming to democratize and simplify access to policy engagement for entities of all sizes, by leveraging scale and technology including AI. We strive to make policymaking accessible, affordable, and inclusive for all. We believe that the open and collaborative nature of dialogue and convenings are essential to further policymaking work, which can be significantly contributed by the participation of innovative companies and startups. Indeed, many if not most of the technologies used to evaluate the measurements, testing and audibility of AI, or to secure AI, are currently developed by such innovative companies and startups – these are the communities OpenPolicy represents and engages. OpenPolicy is also in the process of forming a non-profit entity for policy engagement, with a focus on technology policy in areas of AI and security. In this proceeding, OpenPolicy represents the broad interests of leading innovative companies that develop cutting-edge technologies specifically in the area of AI safety, security and governance, and security research. Indeed, our commitment to promoting trusted AI and working collaboratively with the administration is documented in the White House webpage concerning the AI Executive Order, and we will take part in the NIST-established AI safety institute consortium.[2] In addition, OpenPolicy is committed to advancing good-faith security research and processes that advance vulnerability disclosure and handling. In this proceeding, the group is represented by Dr. Amit Elazari, J.S.D. amit@openpolicygroup.com.

Dr. Elazari is also the co-founder of Disclose.io, an initiative dedicated to advocating for legal protections and "safe harbors" for security researchers, which brief has been cited in the CFAA decision in the Supreme Court decision of *Van Buren*. Disclose.io contract language geared to create legal protections for security research has been adopted across 1,000s of bug bounty and vulnerability disclosure programs, including programs managed by federal agencies, and is

[1] www.openpolicygroup.com.
[2] See the White House, What are they saying, https://www.whitehouse.gov/briefing-room/statements-releases/2023/10/31/what-they-are-saying-president-biden-issues-executive-order-on-safe-secure-and-trustworthy-artificial-intelligence/ (Oct. 31, 2023).

referred to in regulatory requirements such as DHS CISA Binding Operating Directive 20-01, and DHS "Secure by Design" best practices.

Dr. Elazari algorithmic and data privacy auditing work, during our research time at U.C. Berkeley has been published in leading academic press, and awarded by policymakers. She was also awarded "data abuse" bounties for Meta and Google for her research.

Item B. Proposed Class Addressed.

OpenPolicy commentary supports the petition for a newly proposed exemption for Class 4: "*Computer Programs–Generative AI Research*".[3] In addition, OpenPolicy would like to express support of the comments submitted by the "Hacking Policy Council" (HPC), represented by Harley Giger, on Dec. 21, 2023, in support of the proposed exemption.

Item C. Overview.

OpenPolicy supports the proposed exemption for generative artificial intelligence (AI) research under Section 1201 of the Digital Millennium Copyright Act (DMCA).[4]

We believe such exemption can be further expanded to good-faith research performed on broader categories of AI systems or deployments, that extend beyond generative AI. We further believe, similar to HPC, that the research permitted should not be confined to findings or concerns related to "bias", but can include broad sets of undesirable social impacts, and other harmful or undesirable unintended outputs in AI systems, from discrimination to "untrustworthy" behavior.[5]

Notably, while a broad set of research activities, or techniques used, and results garnered, can be covered under the existing security exemption,[6] The potential chilling effect resulting from uncertainty if such activities are exempted, and the need to expand the category to broader sets of findings and research methods (some still emerging), suggest consideration of such exemption is warranted.

[3] 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 Dec. 12, 2023).
[4] Copyright Office, Notice of proposed rulemaking, Exemptions to Permit Circumvention of Access Controls on Copyrighted Works, 88 F.R. 72013, 72025 (Oct. 19, 2023).
[5] "AI risk management calls for addressing many other types of risks" NIST, Artificial Intelligence Risk Management Framework, AI Risks and Trustworthiness, pgs. 12, 39, Jan. 2023, https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.100-1.pdf. See also []
[6] As HPC submission clarifies, some forms of research about AI bias and misalignment are already exempt under the good-faith security research exemption, this may include cases where if the researched risks or harm found are related to confidentiality, integrity, or availability of information, or the physical safety of the users of the machines on which the AI system operates, or otherwise can be categorized as a security vulnerability finding.

Notably, the last year has marked unprecedented recognition by policymakers, academics, and industry experts - as well as civil society, of the need to facilitate at-scale third-party auditing of AI. This recognition builds on decades of research in this domain, and is cultivated by the rise of new AI risks (among others, given generative AI), and adoption of AI systems more broadly.

More notably, the U.S. Administration, Executive Order 14110 on AI, included multiple references and taskings related to AI red teaming, the development of new best practices in this domain, and has elevated the need to perform third-party red teaming of AI, as part of the voluntary commitments for AI. Indeed. AI audits are emerging as key requirements in state and federal requirements under proposed and enacted policies, international standards, and industry best practices.[7]

Policy, civil society., academic, and industry research all indicated, across the security, safety and AI domains — the value that is added when such assessment, audit and testing is done by third parties, external to the organization, often absent the "organization consent", as well as potential barriers to obtaining consent for such research.

Ultimately the consideration of an exemption is consistent with industry best practices for AI red teaming and auditing and will advance the Administration priorities articulated across multiple policies, including in Executive Order 14110.

By performing research, and identifying and disclosing issues, flaws and vulnerabilities so they can be addressed, AI research practices advance the security, safety, trustworthiness, transparency and fairness of AI systems, including generative AI.

While many categories of such research may fall under existing exemptions, the ambiguity rather current prohibition of section 17 U.S.C. 1201(a)(1)(A) on circumvention of technological measures that control access to computer programs can restrict independent AI research, id such research identifies broader unintended consequences or leverage methods that extend from security research, can create a stifling effect on such research, chilling a critical societal activity at a time it is needed most. It is well documented that legal concerns arise and can restrict such research if permission of the computer program copyright holder is required to conduct such research.

Finally, we believe such exemption should be expanded to research applicable to all AI systems, not just generative AI. Notably, policies such as Executive Order 14110 recognize, such testing and

[7] See also, Hacking Policy Council, AI red teaming - Legal clarity and protections needed, Dec. 12, 2023, https://assets-global.website-files.com/62713397a014368302d4ddf5/6579fcd1b821fdc1e507a6d0_Hacking-Policy-Council-statement-on-AI-red-teaming-protections-20231212.pdf.

red-teaming activities are recommended for AI systems, not just emerging generative AI or foundation models. We therefore propose the good faith research exemption applies to AI systems broadly defined, which further includes generative AI.

The below proposal, building on HPC proposal, encourages the Register of Copyrights to formulate an exemption in support of research where the researched "unintended consequence", undesirable outcome, bias or "misalignment" (to borrow HPC term), may not directly affect security or safety (e.g. a prompt injection or data position attack on generative AI system to produce undesirable societal impact, where the AI engages in discrimination, or creates synthetic abusive material).

We propose the following exemption which builds on HPC proposal, leveraging the definitions used in the Executive Order, with slight adaptation to expand the scope to all AI systems and broader research categories:

> (i) Computer Programs, where the circumvention is undertaken on a lawfully acquired device or machine on which an AI system operates, or is undertaken on a computer, computer system, or computer network on which an AI system operates with the authorization of the owner or operator of such computer, computer system, or computer network, for the purpose of good-faith AI research.[8]
>
> *[Note: given concerns regarding the ambiguity associated with the use of the term "solely" in the security research exemption, we suggest removing it or clarifying it, given it is expected much of the social research on AI may be performed in commercial auditing or assessment settings (e.g. as ordered by a third party, including a regulatory agency, or in a "bug bounty")].*
>
> (ii) For purposes of paragraph (i), the term "artificial intelligence" or "AI" has the meaning set forth in 15 U.S.C. 9401(3): a machine-based system that can, for a given set of human-defined objectives, make predictions, recommendations, or decisions influencing real or virtual environments.[9]
>
> (iii) For purposes of paragraph (i), the term "AI system" means any data system, software, hardware, application, tool, or utility that operates in whole or in part using AI.[10]

[8] *Id.* at 201.40(b)(16)(i).
[9] See White House, Executive Order on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence, Section 3(b), Oct. 30, 2023, www.whitehouse.gov/briefing-room/presidential-actions/2023/10/30/executive-order-on-the-safe-secure-and-trustworthy-development-and-use-of-artificial-intelligence.
[10] Id. at Section 3(e).

(iv) For purposes of paragraph (i), "good-faith AI research" means accessing a computer program solely for purposes of good-faith testing, investigation, audit or assessment of biased, discriminatory, or otherwise unintended consequence or harmful outputs in, or of, an AI system, where such activity is carried out in an environment designed to avoid any harm to individuals or the public, and where the information derived from the activity is used primarily to promote the trustworthiness of the AI system, or the safety of those who use such systems, or their trust in such systems, and is not used or maintained in a manner that facilitates copyright infringement.[11]

(v) Good-faith AI research that qualifies for the exemption of this section may nevertheless incur liability under other applicable laws, including without limitation the Computer Fraud and Abuse Act of 1986, as amended and codified in title 18, United States Code, and eligibility for that exemption is not a safe harbor from, or defense to, liability under other applicable laws.[12]

Item D. Technological Protection Measures and Methods of Circumvention.

As articulated in algorithmic auditing research work, various AI auditing methods can entail circumvention of technological protection measures on code, or software.

In the matter of Sandvig v. Barr, it was documented how the copyright owner an AI system may under its terms of service, require a user account or gateway, where by the terms of use of prohibit bypassing any protective measures, prohibit reverse engineering, or limited to ability to create users or scrape data, which limits certain AI auditing methods.

Other terms may prohibit research altogether. These barriers are documented in the context of the Sandvig v. Barr proceedings, cited by the Supreme Court decision in the matter of Van Buren, in the context of the Computer Fraud and Abuse Act, and referred to in legal and broader academic surveys and research. HPC submission well documented some of the other methods that may entail access to the AI system for Audit, as we are delighted to facilitate a meeting with such communities to provide an additional overview.

[11] See 37 CFR 201.40(b)(16)(ii).
[12] Id. at 201.40(b)(16)(iii).

Item E. Asserted Adverse Effects on Noninfringing Uses.

On this item we refer the Copyright Office to HPC comments and express support in their articulation of adverse effects.

Notably, the administration, industry and best practices already identified the need for such AI research, as well as prominence to society – ensuring a proper exemption for AI research is a critical, consistent step in ensuring a robust ecosystem of safety and trustworthiness AI community needed to ensure AI advancements align with societal interests, and would drive coherence with U.S. broader policies. It is essential that policies that seek to encourage and require such as AI audits and research– do so while protecting the research community that performs such research, and cultivating its diversity. .

Finally, we believe the consideration of such exemption is essential for the U.S.G continued leadership in promoting the secure, safe, and trusted deployment of Artificial Intelligence.

Regards,

Dr. Amit Elazari, OpenPolicy





openpolicygroup.com info@openpolicygroup.com