March 18, 2024

Short Reply Comment to US Copyright Office
Ninth Triennial Section 1201 Proceeding (2024)
Class 4: Computer Programs–Generative AI Research.

Dear Sir or Madam:

HackerOne Inc. (HackerOne) submits the following reply comments in response to the US Copyright Office's Ninth Triennial Proceeding on Section 1201 of the Digital Millennium CopyRight Act (DMCA). As noted in our initial comments, HackerOne supports the proposed exemption for Class 4: Computer Programs—Generative AI Research. We support the exemption language proposed by the Hacking Policy Council, and we encourage clarification on any overlap with the existing Section 1201 exemption for good faith security research. HackerOne believes the contributions of independent good faith research are a fundamental driver of holistic improvements to the technology ecosystem.

These reply comments focus on the value of an exemption for good faith AI research under DMCA Section 1201 alongside, as a complement to but not an equivalent of, research access programs and "bias bounty" programs. Research access programs and bias bounty programs are employed by AI service providers to incentivize researchers to identify and report algorithmic flaws in an AI system. Such programs are similar in concept to "bug bounty" programs, which focus on incentivizing researchers to identify and report security vulnerabilities. Research access programs and bias bounty programs are key methods for AI system owners and operators to correct trustworthiness issues in the AI system before those issues are exploited by bad actors or cause problems for users.

However, research access programs and bias bounty programs, much like their security-focused counterparts such as penetration tests and bug bounty programs, are most effective as key components of a holistic AI safety and trustworthiness program. By their nature, these programs are often limited in availability and scope. While the AI owners or operators that participate in such programs tend to be forward-looking and diligent regarding security and trustworthiness, AI systems are rapidly growing in adoption and many organizations do not participate in such programs. In general, the AI owner or operator pays a fee to participate. As with bug bounties, the programs generally permit circumvention of technological protection measures only in specified circumstances that the AI system owner or operator approves, such as on specified assets, using specified methodologies, and subject to terms regarding reporting and disclosure. (See Copyright Office, Eight Triennial Proceeding, Recommendation of the Register, October 2021, page 244.) In addition, research access programs tend to be made available only to certain individuals and entities, not the full community of good faith researchers.

Even if the copyright owner of an AI system participates in a research access program or a bias bounty, such programs tend to not cover every instance of the AI system. The rapid adoption of AI systems means that instances of such systems may be deployed on many platforms that are separate from the copyright owner of the AI system. For example, an e-commerce platform may license a generative AI system from the copyright owner, making the AI system accessible to users with accounts on the e-commerce platform. Good faith AI research may take place on the AI system "instance" (e.g., the copy of the software running on the e-commerce platform) rather than directly through the copyright owner. Research on an AI system instance can be valuable because the instance may exhibit unique

properties as a result of its API or the context in which the instance operates. In such circumstances, an independent researcher may not have explicit permission from both the copyright holder and the instance host, and it may not always be apparent to the researcher which entity owns the copyright to the AI system.

Research access and bias bounty programs are highly valuable tools for organizations to address bias and trustworthiness issues in AI. However, the availability of these focused tools should be a complement to, not a substitute for, an exemption under DMCA Section 1201 for independent AI research or red teaming performed in good faith. Just as the existing DMCA Section 1201 exemption for good faith security research enables a “see something, say something” approach to security vulnerabilities that makes the software and systems we all use more secure, the proposed exemption would have a similar safety- and trustworthiness-enhancing impact on the generative AI ecosystem. HackerOne supports the community of ethical researchers that independently help identify algorithmic flaws so that they can be mitigated for the benefit of the public.

Respectfully Submitted,

Ilona Cohen
Chief Legal and Policy Officer
HackerOne