

Generative AI Meets Copyright Scrutiny: Highlights from the Copyright Office's Part III Report

May 20, 2025

On May 9, 2025, the U.S. Copyright Office released a “pre-publication” version of Part III of its highly anticipated Report on Copyright and Artificial Intelligence (AI) (Report). The Report provides a technical overview of how generative AI models are developed, trained, and deployed and how U.S. copyright law, particularly the fair use doctrine, should apply in the context of training generative AI models. The prepublication report states that it was released “in response to congressional inquiries and expressions of interest from stakeholders” and that “[a] final version will be published in the near future, without any substantive changes expected in the analysis or conclusions.”

Copying Can Occur During Training or Use of Generative AI, and the AI Model's “Weights” May Also Infringe

The Report begins by discussing how curating, collecting, downloading, reformatting, transferring, and incorporating copies into AI model training datasets can involve creating multiple copies of protected works. The Report notes that building a training dataset using copyrighted works “clearly implicate[s] the right of reproduction” and further notes that if model outputs are substantially similar to training sources, the output may also implicate a protected right. In these instances, the conduct is presumptively infringing unless the fair use defense applies.

What happens in the middle of the training process is a bit more nuanced, and the extent to which models memorize training examples is disputed (and likely varies across models). However, according to the Report, if the model can generate an identical or nearly identical copy of the underlying work without that expression being provided in the form of a prompt or input, there is a strong argument that the model's “weights” — numerical parameters that determine the importance of dataset features — could implicate the right of reproduction. Model weights that have memorized protectable expression from training data may also infringe the derivative work right.

The Report notes that whether a model's weights implicate the reproduction or derivative work rights turns on whether the model has retained or memorized “substantial protectable expression” from the underlying works. In such an instance, distributing, fine-tuning, or deploying a model could expose developers and downstream users to liability for infringement.

The Fair Use Defense Must Be Evaluated Within the Context of Overall Use

Where copying constitutes prima facie infringement, the next question is whether the fair use defense applies. The fair use analysis considers four nonexclusive factors: (1) the purpose and character of the use, (2) the nature of the copyrighted work, (3) the amount and substantiality copied, and (4) the market effect.

Factor One: The Purpose and Character of the Use Depend on How the AI Model Is Used.

The Report's analysis of the first factor — purpose and character of the use — focuses on identifying the use, transformativeness, commerciality, and lawful access to the work, with transformativeness and commerciality being key elements. On the critical issue of transformative use, the Report relies on the Supreme Court's reasoning in *Warhol v. Goldsmith* that transformative use is a matter of degree. When applied in the context of training AI models, the Report asserts that the analysis depends not just on the training process but also on how the model is used. At one end of the spectrum are research-driven or closed-system uses. For example, scanning books to create a full-text searchable database to support content moderation may be highly transformative. Training models to generate substantially similar copyrighted works, however, may not be. The Report notes that unlike cases where copying was merely a means to remove interoperability barriers, using images or sound recordings to generate substantially similar expressive outputs is unlikely to be transformative unless the work itself is being targeted for comment or parody.

The Report also explains that retrieval-augmented generation (RAG) — which enhances the performance of generative AI models by scraping information from external databases, documents, or the web — requires separate consideration. Unlike pretraining with a large, diverse dataset, RAG retrieves targeted works for the purpose of enhancing the prompt output. RAG is less likely to be transformative where the purpose is to generate outputs that summarize or abridge copyrighted works.

In making this distinction, the Report essentially disagrees with two common arguments that training AI models is inherently transformative. As to the argument that the purpose is not expressive, the Report reasons that because models do more than just statistical pattern recognition (they learn the selection and arrangement of underlying words, images, and sounds), training the models encompasses the “essence” of creative expression. As to the argument that the process is similar to human learning, the Report reasons that the fair use defense does not protect all copying for the purpose of learning and does not distinguish between acts done by a computer or a human.

As to the role of the other critical element of the first factor — commerciality — the Report notes the commerciality turns on whether the use “furthers commercial purposes,” not on the for-profit or nonprofit status of the entity involved in use of the generative AI model.

Factor Two: The Nature of the Copyrighted Work Depends on the Types of Works in the Training Set.

The Report states that the second factor — the nature of the copyrighted work — depends on the “model and work at issue.” Observers have commented that the second factor rarely plays a substantial role in fair use balancing. The Report notes that most AI models are trained on a variety of types of works and concludes that if the works are more expressive or previously unpublished, this factor would weigh against fair use.

Factor Three: The Amount and Substantiality Copied Should Consider Guardrails Against Infringement and What Content Is Made Public.

The third fair use factor examines how much of a copyrighted work was used and whether that amount was reasonable in light of the purpose. Here, the Report notes that AI model training usually entails full or nearly full copying of entire works and makes use of their expressive content for training, which weighs against fair use. However, the Report suggests that developers can mitigate the presumption against fair use by showing that the copying was functionally necessary to a transformative purpose and that effective guardrails were used to prevent the output of protected expression. It is worth noting that the Report also finds that the presence of technical guardrails is relevant to the first factor as a means of limiting a model's ability to reproduce copyrighted material and the risk of market substitution.

Factor Four: Market Effect Depends on Outputs That May Impact the Market Through Lost Sales, Dilution, and Licensing Fees.

As to the fourth fair use factor — effect on the market — the Report evaluates different ways in which the use of copyrighted works in training AI models can affect the market value of protected works and addresses broader claims that the public benefits of unlicensed training might shift the fair use balance. Here, the Report identifies three categories of potential harm: lost licensing opportunities, lost sales, and market dilution.

In particular, while the first and second categories are typically considered in the fair use analysis, the Report notes that some commentators also advocated for consideration of the potential harm caused by market dilution (i.e., where even those outputs that are not substantially similar to a specific copyrighted work could nevertheless compete in the market for that type of work). The Report authors appear to have been persuaded by this novel theory, noting that “stylistic imitation made possible by [the original work’s] use in training may impact the creator’s market,” and warn that “the speed and scale at which AI systems generate content pose a serious risk of diluting markets for works of the same kind as in their training data.” However, even the Report acknowledges that this position is “uncharted territory” and no court has yet embraced it as a reason to deny fair use.

The Report Advocates for Voluntary Licensing, Collective Bargaining, and Policy Reform.

The Report also discusses various licensing solutions for AI model training, including voluntary, collective, and compulsory licenses, and a statutory “opt-out.” In giving its recommendation, the Report stress the importance of recognizing that training involves a wide variation of works, which will affect the feasibility of the licensing regime. Voluntary licensing may be feasible where there are large volumes of copyrightable material or a limited number of copyright owners. Collective licensing may be feasible and could reduce transactional costs if appropriate safeguards against anticompetitive behavior are implemented and could be used to “preserve some ability to block unwanted uses or negotiate terms” should Congress ever consider an exception or limitation for AI training. However, the Report further recognizes that compulsory licenses could hamper flexible and creative market-based solutions and are arguably inconsistent with the basic requirement of consent to use copyrighted works.

Taken as a whole, the prepublication report on AI model training takes a measured approach but appears to favor copyright owners — most notably in its endorsement of the novel market dilution theory of harm.

CONTACTS

If you have any questions regarding this Sidley Update, please contact the Sidley lawyer with whom you usually work, or

Lauren M. De Lilly , Partner	+1 213 896 6085, ldelilly@sidley.com
Nima H. Mohebbi , Partner	+1 310 595 9648, nima.mohebbi@sidley.com
Rollin A. Ransom , Partner	+1 213 896 6047, rransom@sidley.com
Randi W. Singer , Partner	+1 212 839 8422, randi.singer@sidley.com
Kristina Martinez , Managing Associate	+1 213 896 6123, kmartinez@sidley.com
Christina H. Strohmman , Associate	+1 213 896 6685, christina.strohmman@sidley.com

Attorney Advertising—Sidley Austin LLP is a global law firm. Our addresses and contact information can be found at www.sidley.com/en/locations/offices.

Sidley provides this information as a service to clients and other friends for educational purposes only. It should not be construed or relied on as legal advice or to create a lawyer-client relationship. Readers should not act upon this information without seeking advice from professional advisers. Sidley and Sidley Austin refer to Sidley Austin LLP and affiliated partnerships as explained at www.sidley.com/disclaimer.

© Sidley Austin LLP