

Kluwer Copyright Blog

Generative AI and copyright: Convergence of opt-outs?

Paul Keller (Institute for Information Law (IViR)) · Thursday, November 23rd, 2023

Machine readable opt-outs from TDM



As we head into the last month of the current EU legislative term, there are [increasing signs that EU lawmakers are unable to agree on the AI Act](#), which was supposed to be one of the crowning digital policy achievements of Ursula von der Leyen’s Commission. [Recent media reports suggest that the Parliament and Member States remain at loggerheads](#) over how and (if) the law should regulate so-called foundation models. While this discussion focuses mainly on the tension between innovation and and safety concerns related to such systems, it is also relevant from a copyright perspective, as efforts to introduce transparency requirements related to the use of copyrighted works for training generative AI models are part of the broader set of requirements aimed at foundation models, whose fate is now in jeopardy.

The latest [council presidency compromise proposal](#) (circulated before negotiations broke down) included two different requirements that would apply to providers of generative AI models and that would need to be met before such models could be made available in the EU. Providers of such models would need to “prepare and make publicly available a sufficiently detailed summary of the content used to train the model or system and information on the provider’s internal policy for

managing copyright-related aspects” and they would need to demonstrate “that adequate measures have been taken to ensure the training of the model or system is carried out in compliance with Union law on copyright and related rights, in particular with regards to Article 4(3) of Directive (EU) 2019/790.”

While the first obligation is an evolved version of the language contained in the European Parliament’s report on the AI Act, the second obligation is new. It would introduce the first (and only) explicit reference to the existing EU copyright framework by requiring providers of generative models to demonstrate compliance with Article 4(3) of the CDSM Directive, which allows creators and other rightholders to explicitly reserve the use of their works for text and data mining, including the reproductions necessary for the use of works to train generative AI models (hereafter referred to as “opting out”).

If adopted in this form, such a provision would significantly strengthen the position of creators and rights holders to prevent or license the use of their works for the purpose of training generative AI models as foreseen in the CDSM Directive. This would also reinforce the importance of machine-readable opt-outs for the EU approach to regulating the use of copyrighted works for training (generative) AI models.

Opting out – but how?

This development brings additional focus to the question of how these machine-readable opt-outs should work in practice. As we have shown in a [recent policy brief on defining best practices for opting out of machine learning \(ML\) training](#) (co-authored with my colleague Zuzanna Warso), it is largely unclear how such opt-outs will work in practice, as there are currently no generally recognized standards or protocols for the machine-readable expression of the rights reservation. There are a number of emerging approaches to this issue, ranging from [protocols developed by publishers](#), [services developed by artist-led startups](#), and specifications proposed by AI companies [1,2,3], but there is currently a lot of uncertainty about which of these will be supported by AI model developers. As a result there is significant uncertainty for creators and rights holders about the practical benefits of investing in working with any of these tools and standards.

In [our policy brief](#) we therefore argue that there is an urgent need for the European Commission to intervene in this space and provide guidance on machine-readable rights reservations. We suggest that as a first step, the Commission should publicly identify data sources, protocols and standards that allow authors and rightholders to express a machine-readable rights reservation in accordance with Article 4(3) CDSM, that are freely available and whose functionality is publicly documented. Such a list of standards would provide clarity to rightholders and more certainty to ML developers seeking to understand how to comply with their obligations under Article 4(3) of the CDSM Directive.

Convergence on opt-outs?

To date, the major players in the field of generative AI have been largely silent on how they intend to comply with the obligations under the EU copyright framework. Most of the public discussion about the legal status of using copyrighted works to train generative AI systems has focused on an

increasing number of lawsuits challenging current practice under the U.S. copyright system and whether the use of copyrighted works to train generative AI systems constitutes fair use.

In this context, it is interesting to look at the responses submitted by leading AI companies to the [Notice of Inquiry \(NOI\) on Artificial Intelligence and Copyright issued by the U.S. Copyright Office on August 30](#). Among the nearly 10,000 responses, there are responses from all the big names in generative AI. Not surprisingly, they all argue that the use of copyrighted works to train their systems should be considered fair use. But behind this first line of defense, a number of the major players (Open AI, Microsoft, Google, Stability AI, and Hugging Face) concede that there is a need to respect opt-outs, at least on a voluntary basis.

In this context, [Google](#), [Microsoft](#), and [Open AI](#) all point to the introduction of their own proprietary standards that allow rights holders (“web publishers”) to opt out of having their works used to train specific AI models. These standards are based on extensions to the robots.txt protocol that allow web publishers to exclude works published online from being included in the training dataset of a small number of generative AI models owned by these companies. In addition, Open AI notes that it maintains a web form where creators and other rights holders can request the exclusion of visual works from the training dataset that powers the DALL-E model.

Of the respondents, only Google explicitly positions its opt-out protocol as an implementation of Article 4(3) of the CDSM Directive:

Google-Extended is an example of an approach that complies with the European Union Digital Single Market Copyright Directive, and specifically Article 4’s reference to machine-readable opt-out tools.

While it is interesting to see an explicit acknowledgement of the need to comply with Article 4(3) in this response, it is almost certainly a misrepresentation that google-extended complies with the requirements set forth in that article. Article 4(3) CDSM allows rightholders to reserve the right to “reproductions and extractions” of their works “for the purposes of text and data mining”. For “content made publicly available online”, this must be done in a machine-readable form. It does not – as the response implies – give creators the privilege to opt out of having their works used as training data for specific models operated by individual companies, in a manner determined by the company training those models. [As creators have pointed out](#), such a model-specific opt-out mechanism is worthless to them because it would require them to repeatedly provide opt-outs for each entity that trains models, which would consume disproportionate resources.

In the passage above, Google is used as an example, but the issue is equally present in the approaches taken by Open Ai and Microsoft, which also suffer from offering opt-outs that are both model-specific and must be expressed in a form specified by the companies. And these are far from being the only entities engaged in training activities that fall within the scope of Article 4(3) of the CDSM Directive.

This situation highlights the urgent need for a standardized way for rightholders to opt their works out of such training activities. Such a standard must not be model provider specific and must apply to all uses of the work covered by Article 4(3). The current fragmentation of the field into model provider specific pseudo-standards shows that the development of such a standard cannot be left to them, but needs to happen in a setting with broader stakeholder representation. Such a process would ideally be initiated or supported by the European Commission.

Standards setter or standards taker?

Another development, illustrated by AI companies' responses to the Copyright Office's NOI, is that there is some convergence toward accepting that opt-outs play an important role in the governance of AI training datasets for generative AI systems. As highlighted above, most of the major players in the field acknowledge this in their responses, which point to fair use as the relevant framework, but also indicate that in practice they respect opt-outs in some form. At least two of them – Anthropic and Open AI – also explicitly point to the need to consider “harmony and interoperability of copyright approaches among major economies” (from [the Anthropic submission](#)).

It is precisely at this point where the balanced legislative approach adopted by the EU in the 2019 CDSM directive could become a global template. The approach (consisting of Articles 3 and 4 CDSM Directive) takes into account the interests of the scientific research community (who benefit from the Article 3 exception), creators and rightholders who actively manage their works (who have the right to opt out from all other types of uses). It also takes into account the interests of AI developers and users of AI tools (many of them creators themselves) who retain access to the wealth of content that is shared online but not actively managed.

In this context, it is also worth noting that some fair use scholars have recently begun to suggest that the legal access and opt-out requirements established by the CDSM Directive would need to be taken into account when determining the fair use status of the use of copyrighted works for ML training^[1]. This again signals some convergence of approaches across different copyright traditions.

The possibility of the EU approach becoming a global template makes it all the more important to complete the EU framework by identifying standards for opting out. Without a generally accepted standard (or set of standards), the system of balances contained in the TDM exceptions will not be able to survive contact with the reality created by the sudden emergence of generative AI as a major new technology paradigm.

Note – This post has been updated to remove a footnote that incorrectly claimed that the Google-extended mechanism only applies to two specific models. It does in fact apply to a much wider range of models. The author regrets the error.

^[1] See for example Professor Mathew Sag's thoughts on Fairness and Fair Use in Generative AI delivered during the 12th Annual Peter A. Jaszi Distinguished Lecture on Intellectual Property in September of this year.

To make sure you do not miss out on regular updates from the Kluwer Copyright Blog, please [subscribe here](#).

Kluwer IP Law


The **2022 Future Ready Lawyer survey** showed that 79% of lawyers think that the importance of legal technology will increase for next year. With Kluwer IP Law you can navigate the increasingly global practice of IP law with specialized, local and cross-border information and tools from every preferred location. Are you, as an IP professional, ready for the future?

Learn how **Kluwer IP Law** can support you.

79% of the lawyers think that the importance of legal technology will increase for next year.

Drive change with Kluwer IP Law.

The master resource for Intellectual Property rights and registration.



2022 SURVEY REPORT
The Wolters Kluwer Future Ready Lawyer
Leading change

This entry was posted on Thursday, November 23rd, 2023 at 10:27 am and is filed under [Artificial Intelligence \(AI\)](#), [European Union](#), [Legislative process](#), [Text and Data Mining \(TDM\)](#). You can follow any responses to this entry through the [Comments \(RSS\)](#) feed. You can leave a response, or [trackback](#) from your own site.