


Kluwer Copyright Blog

Copyright and Generative AI: What Can We Learn from Model Terms and Conditions?

Gabriele Cifrodelli (University of Glasgow) and Lilian Edwards (Newcastle Law School) · Wednesday, July 24th, 2024

Although large, general purpose AI (GPAI) or “foundation” models and their generative products have been around for several years, it was ChatGPT’s launch in November 2022 which captured the public and media’s imagination as well as large amounts of venture capital funding. Since then, large models generating not just text and image but also video, games, music and code, have become a global obsession, touted as set to revolutionise innovation and democratise creativity, against a background of media frenzy. Google, Meta and now even Apple



have integrated foundation model technology into their lead AI-generated image by DALL-E 3 (through Microsoft Copilot) products, albeit not without controversy. based on Gabriele Cifrodelli’s prompt: ‘Terms and Conditions on a cracked computer screen’

The relationship between copyright and generative AI (genAI) has turned out to be one of the most controversial issues the law has to resolve in this area. Two key issues have generated much argument, relating respectively to the inputs to and outputs from large models. On the first, substantial litigation has already been launched concerning whether the data used to train these models requires payment or opt-in from creatives whose work has been ingested, often without consent. While creative industries claim their work has been not only stolen but specifically used to replace them, AI providers continue, remarkably, to insist that the millions of images ‘fed’ to the AI can be used without permission as part of the “social contract” of the Internet. The results of these disputes are likely to take years to work through and may have very different outcomes in different jurisdictions given the very wide scope of fair use in the US compared to (inter alia) the EU. Turning to outputs, courts and regulators have already been asked repeatedly (and usually

answered no) as to whether genAI models, especially Text-To-Image (T2I) models, can be recognised as the creators of literary or artistic works worthy of some sort of copyright protection.

These two points have generated substantial policy and academic discussion. But less attention has been paid to how AI providers regulate *themselves* by their terms and conditions – what is known as [private ordering in the contractual context](#). AI large model providers regulate their users via a variety of instruments which range from the arguably more legally binding terms and conditions (T&C or terms of service (ToS)), privacy policies or notices and licenses of copyright material, through to the fuzzier and more PR-friendly but less enforceable “acceptable use” policies, stakeholder “principles” and codes of conduct. While study of social media and online platform private ordering is a very well-established way to find out how providers deal with copyright, data protection and consumer protection, studies of generative AI T&C have been slower to get going. Study of ToS is crucial because in most cases, pending the resolution of litigation or novel legislation, they will effectively be what governs the rights of users and creators. Yet especially in the business-to-consumer or “B2C” context, these ToS have often been reviled as largely unread, not understood, and creating an abusive relationship of imbalance of power in monopolistic or oligopolistic markets. Indeed, Palka has named T&C of online platforms “[terms of injustice](#)” and argued they should no longer be tolerated. With this background, we chose to run a small pilot as soon as possible to see what terms were being imposed by generative AI providers, and whether the results were indeed deleterious for users and creators.

Our pilot empirical work in January-March 2023 mapped ToS across a representative sample of 13 generative AI providers, drawn from across the globe and including small providers as well as the large globally well-known firms such as Google and OpenAI. We looked at Text-to-Text models (T2T – e.g. [ChatGPT](#)); Text-to-Image models (T2I – e.g. [Stable Diffusion](#) and [MidJourney](#)); and Text-to-Audio or Video models (T2AV e.g. [Synthesia](#) and [Colossyan](#)). We analysed clauses affecting user interests regarding privacy or data protection, illegal and harmful content, dispute resolution, jurisdiction and enforcement, and copyright, the last of which provided perhaps our most interesting results and which is the focus of this blogpost.

Drawing on emerging controversies and lawsuits, **we broke our analysis of copyright clauses into the following questions:**

1. Who owns the copyright over the outputs and (if any indication is found) over the inputs of the model? Is it a proper copyright ownership or an assigned license?
2. If output works infringe copyright, who is responsible (e.g. user, service)?
3. Did model providers undertake content moderation (e.g. prompt filtering) to try to reduce the risk of copyright infringement in outputs?

Question 1 gave inconsequential results **re inputs**. There was almost no reference to ownership of training data that had come from parties other than the contractual partners. [ChatGPT](#), for example, defined inputs restrictively to mean prompt material and recognised the user’s ownership. We had hoped perhaps naively for some indication of the rights of creators in relation to copyright works used to train the models *ex ante* but of course since these lay outside the model – user relationship we found almost nothing. Interestingly, at the time of our study the issue of whether users of a primary service could by default be required to provide their data to help train and retrain the large models being developed by the service provider had not become as acute as it has more recently, e.g. in relation to [Adobe](#), [Meta](#) and [Slack](#). We hope to return to this theme in future work.

Concerning *outputs* however, the results were more interesting. In almost every model studied, ownership of outputs was assigned to the user, but in many cases, an extensive license was also granted back to the model provider for coexisting use of the outputs. The terminology was often very similar to that familiar from the ToS of online user-generated content (UGC) platforms like Google and Meta. T2I model [Lensa](#), e.g., granted the user ‘a perpetual, revocable, nonexclusive, royalty-free, worldwide, fully-paid, transferable, sub-licensable license to use, reproduce, modify, adapt, translate, create derivative works’. By contrast, T2I [Nightcafe](#) simply prescribed that once the content was created and delivered to the user, the latter owned all the IP Rights. [Stable Diffusion](#) adopted a commonly known open-source license, the [CreativeML Open RAIL-M license](#), that allowed its users not just rights over their generated output artworks but also to deliver and work with the Stable Diffusion model itself.

In T2T services, OpenAI’s [ChatGPT](#) assigned to the user all the ‘right, title and interest in and to Output’. [Bard](#), [Simplified](#) and [CLOVA Studio](#) also assigned ownership to users. By contrast, the company Baidu – proprietor of [Ernie Bot](#) – identified itself as the owner of all IP rights of the API service platform and its related elements, such as ‘content, data, technology, software, code, user interface’. Unusually, [DeepL](#), an AI translation service, did ‘not assume any copyrights to the translations made by Customer using the Products’.

Why were providers so willing to give away rights over the valuable outputs of their services, especially when for consumers at this stage of genAI development, the services were largely free?

Question 2 gave us some clues. In almost every model or service studied, the *risk* of copyright infringement in the output work was left, with some decisiveness, with the user. For instance, [Midjourney’s T&C](#) used entertainingly colourful language:

‘[i]f you knowingly infringe someone else’s intellectual property, and that costs us money, we’re going to come find you and collect that money from you’.

So what we found was a Faustian bargain whereby users were granted ownership of the outputs of their prompts but only so long as they also took on all the risk of copyright infringement suits from upstream creators whose work had been absorbed into training sets. Yet infringement risks will come near exclusively from the contents of the training datasets, often gathered without notice or permission from creative content providers, and whose contents are often a proprietary secret where users have no idea of any arrangements for consent or compensation. This seems the essence of an unfair term.

We argue in our full report that AI providers are thus positioning themselves, via their ToS and to their sole benefit, as “neutral intermediaries”, similarly to search and social media platforms. They trade ownership over outputs in exchange for assignment of risk to users, making their profits not from outputs but from subscription and API fees, and quite likely in future, just like online platforms, advertising. Yet genAI providers are not platforms; they do not host user generated content, but simply provide as a service *AI generated content*. We call this a ‘platformisation paradigm’, a deceptive practice whereby AI providers claim the benefits of neutral host status but without the governance increasingly imposed on these actors (e.g. in Europe through the [Copyright in the Digital Single Market Directive](#) and the [Digital Services Act](#)). As of February 2024, EU online platforms (not just very large ones or “VLOPs”!) have to make their ToS and content moderation actions public and also [take into account the rights and interests of users when interpreting and enforcing their ToS](#). None of these new rules ameliorating the “terms of injustice”

Palka refers to, [apply to genAI providers](#) (at least unless the services are incorporated into services subject to the DSA such as GPT incorporated into Microsoft's Bing, a Very Large Online Search Engine (VLOSE)).

The platform paradigm is reinforced in optics by the way almost every model provider except the smallest undertook content moderation, with notice and take down arrangements the norm (**Question 3** above). Again, although users would bear the risk of liability associated with outputs, model providers invariably exercised their own discretion in assessing what output or behaviour violate the ToS, and what the sanction might be (site ban, for example) (see for instance, [Nightcafe](#)).

In conclusion, while academics, legislators and judges are arguably seeking to balance the interests of creators whose work is used to build genAI models, the providers who build them and the rights of users of these services, ToS analysis provides a familiar sight of one-sided contracts of adhesion, written in legalese to minimise risk and maximise control to service providers masquerading as platforms to evade regulation. We argue this situation needs addressing, at least by analysis from consumer protection law but quite possibly by reflection on how the DSA can be extended to govern generative AI and foundation models. Another solution may be to take up these points in the code of conduct for GPAI providers which the Commission now has nine months to draft – but [since that process already seems to have been co-opted by the AI companies themselves](#), we do not hold out much hope in that direction.

This blog post is based on the findings of a pilot empirical work conducted between January and March 2023 funded by the EPSRC Trusted Autonomous Systems Hub. You can find the full report [here](#).

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 Wednesday, July 24th, 2024 at 9:00 am and is filed under [Artificial Intelligence \(AI\)](#), [Infringement](#), [Ownership](#)

You can follow any responses to this entry through the [Comments \(RSS\)](#) feed. You can leave a response, or [trackback](#) from your own site.