PHILIPPE GILLIÉRON
Avocat

Auteur: Wilhelm Avocats | Le : 15 August 2023

Generative AI and copyright: quo vadis?

Generative AIs are all around. The questions surrounding them are numerous. We will limit ourselves here to summarize the issues raised in terms of copyright and the possible answers (some even fairly likely) in the light of current law and the broader questions that these tools raise in this area.

I.   Assessment under current law

a)   Input data training

Large Language Models (LLMS) like ChatGPT and other generative tools require unimaginable data ingestion to achieve the results they do. By way of example, GPT-4 is said to have been trained on trillions of pieces of data made up of both words and images, some of which are protected by copyright.

Can we talk about copyright infringement? This is the opinion of numerous rights holders who, since the end of last year, have launched actions in the USA against [Microsoft, GitHub and OpenAI](#) in connection with the Copilot tool's training, [artists against Stability AI and Midjourney](#) and, before the High Court of London, [Getty Images against Stable Diffusion](#).

First question: does the operator of these platforms have standing to defend?

The question is worth asking, given that, for the image-generating tools in particular, the service provider who supplied them is a third party, namely the German association LAION. This raises two questions:

- Is the LAION association, which, according to its website, has no profit motive, liable to be sued for copyright infringement? At first glance, the answer seems rather negative. This platform, built on the basis of commoncrawl, only provides links to the sites where the images can be found, to the exclusion of any reproduction of these images on its servers. The only remaining question consists of assessing whether this aggregation of links constitutes an infringement of the right to make available:

At European level, art. 17 of Directive 2019/790 does not appear to be applicable to LAION, since the latter cannot be qualified as an online content sharing service provider. Since *Svensson* and *Bestwaters* (C-466/12 and C-348/13), it is admitted that a simple hyperlink to a protected work for a public that already has unlimited access to the work in question does not satisfy the "new public" requirement laid down in Community law, since the work is already accessible to everyone on the original site. Under the ECJ case law, LAION's liability may only come into play if it was aware that the links redirected to works posted online without the consent of their authors, as ruled in the *Sitching Brein* case (C-610/15). In this case, we can assume that, in the majority of cases, the data processed are public and do not come from illicit sources, although such a presumption may be rebutted.

Under Swiss law, the question of whether the insertion of hypertext links to works posted online with the consent of the rights holders infringes their right to make them available remains somewhat undecided, but the majority opinion would lean towards the ECJ case law. In other words, one should distinguish whether the link refers to works that have been posted online with (no copyright infringement) or without (copyright infringement) the rights holder's consent.

Under U.S. law, it is highly likely that, assuming such content is protected, the mere creation of an aggregation of links falls within the notion of *fair use* as defined by § 107 of the *United States Copyright Act*.

All in all, and without going into the matter in detail, it would seem that LAION has good arguments to put forward, should the need arise, to avoid potential liability under copyright law.

- As far as the operator is concerned, there can be little doubt that it reproduces this data to train its algorithm, even if the data may be tweaked and modified to improve training. While the technical operation of the tools will certainly be scrutinized in details in these proceedings, there is serious doubts that platform operators will be able to avoid the accusation of infringement of reproduction rights.

Taking into account the fact that, irrespective of LAION's status in the case of image-generating AIs, the platform operator does indeed reproduce the works on its servers, we can move on to the second question:

Secondly, **are ingested data protected by copyright**?

The answer to this question depends on the type of content and the conditions required for a creation to be considered a "work" protected by copyright in a given state:

- In the case of **text,** the reproduction of a single word or a few words can hardly be perceived as an infringement of a protected "work", but the longer the reproduction, the closer it comes to being a protected "work". As a result, it will be difficult not to recognize a paragraph or chapter as having sufficient individual character to be considered a protected work. This is all the more important than there would be evidence that some models have been trained on entire chapters

[of books](#).

- This protection will be all the more easily granted for excerpts from **musical works** or **audiovisual creations**, for which the protection by copyright bears little doubt.

- As for **images,** while some countries require a certain degree of individuality, which may lead to some discussion, others such as Germany and Switzerland protect photographs as such, irrespective of whether or not they have an individual character.

In short, given the volume of data ingested, it will be very difficult for any operator to establish that it does not reproduce copyrighted data in a way or the other to train its algorithm.

This being the case, it should not be forgotten that in terms of the burden of proof, it is up to the plaintiff to establish that the system has been trained on a protected work in which it holds rights; proof is not always easy. From a procedural point of view, it will undoubtedly be interesting to see the possible requests for expert opinions and the production of exhibits, which may require the plaintiff to go fishing for information in order to prove his or her rights after having initiated proceedings...

Third question: can the platform operator invoke an exception?

Since the reproduction of copyrighted content seems difficult to contest, the question arises as to whether the operator can benefit from an exception that would be provided by copyright laws. Here, everything depends on the applicable law, as different countries have opted for different approaches, the vast majority of which are still likely to evolve.

- Japan, for example, has ruled that exploitation of copyrighted content for the purposes of training an algorithm does not infringe copyright.

- In the **European Union,** art. 3 and 4 of Directive 2019/790 authorize text and data mining for scientific research purposes, an exception that is similar to Art. 24d of the Swiss Copyright Act; to benefit from this exception, however, the data must be used for scientific purposes, and thus lead to a scientific publication in the first place. Actors from the private sector may thus find it hard to benefit from such an exception. In contrast to **Swiss law,** article 4 of European Directive 2019/790 however adopts a more favorable approach to the private sector. This provision allows commercial companies to exploit such data for training purposes, provided that the rightful owners have not prohibited such exploitation by appropriate means, in particular by prohibiting scraping. This reservation notwithstanding, European law nevertheless appears to be more favorable than Swiss law.

- Finally, in the **United States,** the whole question is likely to revolve around whether the reproduction of copyrighted data in such a system constitutes *fair use*.

In [Warhol v. Goldsmith](#), a case handed down by the Supreme Court on May 18, 2023, and which I have commented on in this [blog](#), the Supreme Court ruled that a case of *fair use arises* when the original work has been transformed to such an extent that it now serves a different purpose from the one initially contemplated by the original right holder.

It is conceivable that the platform operators' legal counsels will seek to take advantage of this ruling, arguing that the use of data for

training purposes pursues an entirely different purpose from the one initially contemplated by the authors of the original works. Their objective is not to exploit the work for its original goal in mind, but to train an algorithm, without seeking to exploit and "consume" such work in its original purpose. In our opinion, such an argument has a good chance of success. To be continued.

To date, Swiss law therefore appears to be particularly protective of rights holders who, if they manage to meet their burden of proof and demonstrate the use of their works, will have a good chance of being able to invoke an infringement of their copyright, with little chance for operators to be able to invoke any exception whatsoever. The situation will be different in Europe and, even more so, in the United States, culminating in Japan, which now appears to be an Eldorado for training the algorithms of these platforms.

b)   Output

Another question consists of assessing to what extent the output generated by *Large Language Model* islikely to be protected by copyright, and who will own the rights to it.

In my opinion, the answer to this question should not be so complex.

Copyright for the developer of the *Large Language Model?*

It is hard to imagine that the developer of the *Large Language Model* used could claim any copyright whatsoever on the result generated by a user through various prompts. His model may itself be protected by copyright or patent if it satisfies the conditions laid down by these laws, but there is no justification for the developer to be granted any copyright whatsoever on the result generated by a user.

It is also hard to imagine that the developer would be so indelicate as to include in its general terms and conditions an outright transfer of copyright on the results generated by its users, a provision that would caught the attention and ultimately be detrimental to its own model; on the other hand, a license allowing the operator to re-exploit these results to feed its algorithm seems much more plausible. Such typically of [Midjourney](#), whose terms do not go so far as to require its users to guarantee that the results generated do not infringe the rights of third parties.

Copyright for the users of *Large Language Model?*

It is primarily the user who should be able to claim copyright on any results generated.

In my opinion, nothing should preclude the possibility that, where appropriate, the output is protected by copyright if it displays a sufficient level of individuality.

Having done the exercise myself, the average user will quickly realize that it can actually be difficult to produce an image that corresponds to the one you were hoping for, and that the sequence of prompts to reach such a result requires a certain amount of dexterity.

Having said that, it's important to note that, in my opinion, it's not the sequence of prompts that should be protected, any more than the individual brushstrokes of a Picasso or Van Gogh to construct their painting are. Prompts, like brushstrokes, are at first sight more akin to a form of method and thus, in itself unprotectable. Only the result counts.

This leaves open the issue to know whether the output, assuming it is copyrighted, constitutes a derivative work of any work on which the algorithm was trained. This notion, familiar to the various legal systems reviewed here, would certainly deserve a more detailed analysis. Suffice it to say here that this should only be the case if the individual features justifying copyright protection of the ingested

work (*input)* can be found in the output. The fact that the user had no intention to do so, was acting in good faith and was unable to ascertain what data was actually ingested and used to generate the result delivered to him or her is irrelevant. Here again, only the result counts, based on an objective assessment.

II.   Legislative policy considerations

While the foregoing remarks relate to a legal assessment based upon the existing law, certain legislative policy considerations are nonetheless worth noting.

A basic question consists of assessing in the first place whether the training of these algorithms on potentially protected works should be admitted or not. While Japan has made its choice, this is not the case in most countries, which have not modified their copyright laws in this respect.

It is a tricky question:

To admit that such training amounts to a copyright infringement, and therefore leads to an obligation to remunerate the owners whose works are ingested, means recognizing the obligation to compensate creators and, more broadly, to support human creation, which some believe is in danger due to the advent of these platforms.

In the absence of such recognition, human beings may no longer have the incentive to invest their time for results that an algorithm can in many cases generate more quickly and at lower cost, bearing in mind that not all artists who make a living from their work are Picassos or Miyazakis who have nothing to worry about.

Protecting artists, however, means potentially giving *Big Tech* the upper hand, since the latter could, depending on how remuneration is determined, be the only ones in a position to compensate owners. The result would be an even greater concentration of power in the hands of these players, to the detriment of the smaller ones.

On the other hand, considering that creators whose works are ingested have no reason to be compensated for an exploitation that is far removed from the primary purpose of their works – a choice that Japan has made – is potentially encouraging competition and innovation by a much larger number of companies, but it is also potentially endangering human creation for a large majority of artists.

Needless to say, the choice is a difficult one. To date, initiatives have tended to come from the private sector, as witnessed, for example, by the [current negotiations between Google and Universal Music](#) for Google to take over excerpts of musical tracks or artists' voices on compositions they have never sung (which, incidentally, is more a matter of personality rights than copyright) to train its algorithm. The music and audiovisual industries, and in [particular the players in these sectors who are likely to be replaced by AIs](#), have put themselves at risk by having overlooked the importance of *peer-to-peer*. They are now trying to take the bull by the horns so as not to be outdone. It is doubtful, however, that such isolated approaches are the best ones, however welcome they may be as preliminary steps to address the situation.

Should we continue to explore collective management and the adoption of tariffs calculated on an egalitarian basis, taking into account, for example, the size of potential platform developers and their sales figures? This is undoubtedly one of several possible avenues.

In my opinion, however, such an effort needs to be coordinated on an international level, as a fragmented approach by country can only be detrimental to the greatest number of developers, users and authors. It's a safe bet that progress in this area will be rapid if OpenAI, among others, doesn't go bankrupt, a [risk now underlined by some](#) in view of the prohibitive development costs of these *Large Language Models* and their difficulties in finding a profitable financial model to date – a situation which is not the first to arise, however, as Haskel and Westlake pointed out in their book *Capitalism without Capital.* To be continued...

Ndlr : this paper was written without using an AI tools.