WHITE PAPER

The Remuneration of Music Creators for the Use of Their Works by Generative AI

Professor Daniel J. Gervais

Prepared for the International Council of Music Creators (CIAM)
and Fair Trade Music International (FTMI), April 2024

From the Author

The Purpose of This Paper

This paper explains the challenge of ensuring ongoing compensation for music creators and their industry partners once most of the existing music has been used to train large language models. The paper does two main things: It explains the technology and how current law applies to it, and then proposes a possible new right. The proposed new right is outlined at a fairly high level of generality in order to focus the discussion on its desirability rather than the exact mechanics of its implementation. Unfortunately, the text has to delve into a number of complex legal doctrines that may not be easy to read for readers without legal training. However, the legal analysis section ends with a summary box intended for those readers who may not need to know all the details of these doctrines.

Daniel Gervais, April 2024

From the Music Creator

A New Right of Remuneration for Generative Artificial Intelligence (GAI)

From the invention of the player piano to the dawn of the internet, advances in technology have resulted in the expansion of rights for music creators. Thanks to these necessary “new” rights, music creators today are remunerated for a host of 20th century advances in music production and exploitation including vinyl records, tapes and CDs; radio and TV broadcasts; streaming; downloads and other uses of musical works.

Now, in the face of the enormous challenge posed by generative artificial intelligence, we believe an additional right of remuneration vested in individual human creators is needed. This new right would be instrumental in providing a sustainable future for our creative community, and to preserving our diverse cultures and identities around the world.

The key points are as follows:

- CIAM and its global partner alliances support all existing rights and their prospective application to GAI.
- We propose an additional right of remuneration for the ongoing use of human created works by GAI platforms. This new right would be vested in the original human creator, who can then assign this right to a CMO, administrator or relevant party.
- GAI platforms must keep track of and be transparent regarding the use of specific human created works. Attribution of this kind helps provide the basis for the fair distribution of revenues to the appropriate creators and rights holders.

The following white paper, written by Professor Daniel Gervais of Vanderbilt Law School, provides a detailed legal analysis of the benefits of a new right of remuneration for GAI, its promise in supporting the sustainability of human authorship of creative works, and related legal issues.

Sincerely,

Eddie Schwartz
President, The International Council of Music Creators (CIAM)

Executive Summary

Generative AI (GenAI) applications challenge humans on the very terrain that has distinguished us from other species for millennia: our ability to create literary and artistic works to communicate new ideas to one another, whether as works of music, art, literature, or journalism. We urgently need to find a way to avoid irreparable damage to this crucial facet of human existence - a *sine qua non* for human progress - an ability that tends to be honed over time by creators who have the time to do so and learn from experience, which often means that they can live off the fruits of their labor. The stated aim of the paper is to find a way for creators to retain agency as their life's work is taken without their consent to create "content" that can compete with them in the marketplace.

The best way for creators to generate a decent stream of ongoing revenue for the use of their copyrighted works by GenAI applications is to be paid when the datasets used to train GenAI containing their works are used to create new "content". This should take the form of a license. For this to happen, there must be a right that can be licensed.

From a legal point of view, the discussion revolves around which rights apply to the training (text and data mining) and to the production of literary and artistic works. In almost all cases, the development of a Large Language Model (LLM) implies the creation of at least one copy of the data that the machine uses for its training. This has several advantages, including increased speed of access and the ability to examine and make changes to the dataset. From a copyright perspective, this implies one or more reproductions. In the case of copyrighted works, this means that the right of reproduction has been infringed unless a license has been obtained or a statutory exception applies.

What is often misunderstood is that this reproduction of the copyrighted work continues to exist in modified form (i.e., a second reproduction occurs) in the dataset created during the training process. This second dataset is the one used by the LLM to produce its outputs. It consists of the creation of "tokens" based on the material used for training.

The outputs of an LLM may infringe both the right of reproduction and the right to prepare derivative works, also known as the right of adaptation (and its close cousin, the right of translation). An adaptation includes, for example, a musical arrangement or a film based on a novel. The exact scope of the concept of derivative works in this area is controversial.

Against this background, existing copyright law provides a partial solution for authors and other right holders for four main reasons.

First, there are different national exceptions and limitations to copyright rights in relation to text and data mining (TDM) i.e., the “input” or training stage, which delineate what companies producing LLMs can and cannot do without a license. In the United States, where many of the best-known LLMs have been created, there is (and will continue to be for years to come in the absence of a licensing regime) doubt about the scope of fair use in this context.

Second, although the copying that occurs during the training of GenAI systems typically occurs only a few times for each GenAI dataset or LLM model, some major models (such as OpenAI’s) are moving toward the creation of an infrastructure layer, that is, a dataset that can be used by other companies and individual users.

This dataset contains, as mentioned above, a complete or partial copy of the material used for training, which implies possible liability for users who make a copy. Nevertheless, the number of copies of copyrighted material used to create the dataset will be limited.

Third, the reproduction right and/or the derivative work right is more easily applied to certain GenAI outputs that are a copy or adaptation of a substantial portion of one or more identifiable pre-existing works in the dataset. If this is true, only a relatively small percentage of GenAI outputs are likely to infringe the reproduction right, the derivative work right, or both.

Fourth, as a matter of copyright law, there is no protection per se for a “style” or “sound” (e.g., a person’s distinctive voice), although statutes and various legal doctrines may provide protection against this form of appropriation.

Despite these legal complexities, there is a deep sense among many authors and performers that the creation of datasets containing their tokenized works without consent or compensation is an unfair situation, a misappropriation, for which they expect the law to provide a remedy. Unfortunately, while the law of misappropriation exists, it is not internationally harmonized and is unlikely to be any time soon. There is a related view that anything created using a data corpus containing tokenized copyrighted material is a “derivative” of the dataset, and in a layman’s sense this is the case, since no output would be generated if it weren’t “derived” from the dataset by the GenAI application. Unfortunately, the legal terms “adaptation” and “derivative work” are likely to be interpreted more narrowly by the courts. Rights holders seeking to correct what they perceive to be an injustice will undoubtedly pursue avenues based on existing laws, including copyright, publicity rights, and misappropriation claims. These may lead to settlements for the use of existing material, including compensation for “past sins”.

This paper examines the applicable norms of international copyright law, and considers an additional option, namely the creation of a *right of remuneration* for creators to compensate for the use of LLMs created using their copyrighted works to produce commercially available “content” that can compete with the material on which

the machine was trained. The proposed right should vest in them though it would remain assignable or licensable. For example, when a music streaming service filled a stream with AI-produced music, it would pay for the use of the copyrighted works in the dataset used by its generative AI model. This would be another adaptation of the copyright framework to a major technological change, as copyright has consistently done for more than two centuries. Indeed, it would be strange if copyright did not adapt to what is perhaps the most consequential technological change in history.

To be clear, this proposed solution does not preclude a licensing regime for the reproduction(s) that occur during the TDM process, which is already the subject of litigation in several jurisdictions. What it does is add a clearly defined, ongoing layer of compensation for the benefit of music creators and rights holders for GenAI systems that produce material in competition with the creators of the copyrighted material on which they were trained.

I. The Legal and Technological Framework of GenAI

A. Technological and Factual Overview

Artificial Intelligence (AI)[1] can perform a multitude of functions, including operating customer service call centers, medical imaging, and driving autonomous vehicles. Generative AI (GenAI) is a subset of AI used to refer to AI systems whose primary function is to generate "content" that mimics literary and artistic works produced by humans and protected by copyright.

GenAI is often used to refer to a Generative Pretrained Transformer (GPT) language models such as ChatGPT or Meta's LLaMA Generative AI also includes diffusion models used in image and video generation, for example StabilityAI. This paper focuses mostly on LLMs, although brief references to diffusion models are also included.

Another term often used as a synonym for AI (although this is not entirely accurate) is "machine learning" (ML). ML is a type of AI that involves a process by which a computer "learns" from a data set. This process, which applies to both transformer and diffusion models, can be supervised by humans, but sometimes it is not, as in the case of deep learning. In such cases, the machine learns on its own. Research on major models suggests that the larger the model, the better the results.

This means that in order to build powerful models, it is advantageous to use massive computing power and significant resources to create a very large dataset for the machine to "learn" from. The need for such investment suggests that large players are likely to dominate the field for the foreseeable future.

The data used to create the dataset is typically copied locally (i.e., where the model is being created), both to speed up the learning process and to allow access to the original dataset (e.g., to remove or add data). In creating a Transformer architecture model, the computer breaks down data, often consisting of literary or artistic works, into smaller representations of the words or music, called *tokens*. For example, a language model can break down the words in a text into tokens, which can be letters, syllables, words, or phrases, depending on the algorithm. As with other forms of machine learning, once the

1 The OECD defines an AI system as "a machine-based system that, for explicit or implicit objectives, infers, from the input it receives, how to generate outputs such as predictions, content, recommendations, or decisions that can influence physical or virtual environments. Different AI systems vary in their levels of autonomy and adaptiveness after deployment". OECD Recommendation of the Council on Artificial Intelligence (2019).

computer has “learned” enough, humans can, but do not have to, review and refine the results.[2]

GPT models, such as the one used by OpenAI, implement a tokenization technique called Byte-Pair Encoding (BPE). BPE creates tokens by combining commonly occurring pairs of characters until a desired vocabulary size is reached. The larger the data set, the better the systems tend to perform. Tokens are different from one model to another because the “tokenizers” used to create them are different. Tokens are essentially numbers. The machine converts the sequences of words from the original text into “vectors” called “word embeddings”, which are ordered sets of numbers, like rows or columns in a spreadsheet. The embedding of a token identifies its relationship to other words in the sentence. These embeddings indirectly *preserve the original representation* (for example, natural language) on which they were trained. In other words, these embeddings are representations of large chunks of text, or even entire works, on which the machine was trained. Embeddings are often stored in the dataset in one form or another. The GPT (or “Transformer”) technology can reproduce entire chunks of works.

Thus, contrary to a common misconception, the process of ingesting text to train an LLM does not involve the destruction of the copied material. Rather, it involves breaking down the copyrighted material into smaller units while preserving the relationships of words or other elements within those units. It is through these word embeddings that the AI system captures and stores the relationships of sequences of words, sounds, pixels, etc. This representation is key to the semantic properties of LLMs. For example, chunking allows the machine to understand that the relationship between, say, “Washington” and “United States” is the same as the relationship between “Rome” and “Italy,” even though they are lexicographically very different words[3].

When a copy of a work is tokenized, very often some information (“metadata” or “Rights Management Information (RMI)”[4]) contained in the original digital file (for example, the name of the author and publisher and the place and year of publication) is deleted (or “omitted”) from the local copy made for training purposes.

Simply put, an LLM is a giant prediction machine. It uses the tokenized/chunked data set to predict the “next best word” (or pixel, or chord) in response to a prompt or fine-tuning instruction. An analogy may be helpful. Imagine taking a book and using scissors to cut it into pieces, each piece containing one or a few words. These are similar to the “tokens” created during the training of an LLM. However, it would be incorrect to think of the tokenization process as throwing all the paper tokens into a big box and shaking it. The tokenization process maintains relationships between the tokens, as if the pieces of paper had little threads that kept them “related”. While this analogy is obviously imperfect, it is useful in explaining the obvious, namely that

2 In contrast, diffusion models break up images in pixels and that, to simplify, the machine then learns to reassemble.

3 The author is grateful to Dr Babis Marmanis (CCC) for the background information. All errors are mine, however.

4 The official international definition is “information which identifies the work, the author of the work, the owner of any right in the work, or information about the terms and conditions of use of the work, and any numbers or codes that represent such information, when any of these items of information is attached to a copy of a work or appears in connection with the communication of a work to the public”. WIPO Copyright Treaty (1996), art 12(2).

even after works have been tokenized, an LLM may regurgitate long excerpts from a single work in its dataset, i.e., a long string of tokens that matches an existing "string" in the original work. The fact that LLM vendors may try to program algorithms to avoid such output should not obscure the fact that the strings of words (or pixels, etc.) still exist, albeit virtually, in the tokenized dataset. This has legal implications discussed below. It opens the door to the argument that not only the original dataset in which copyrighted works were copied, but also the tokenized dataset are both reproductions, and that a reproduction of the tokenized dataset would itself constitute a new reproduction of the copyrighted works used to create it. This may have direct implications for licensing and who may need a license.

The purpose of GenAI is to produce *outputs*. As just noted, the outputs are produced by using the dataset to predict the next best word, musical chord, pixel, etc. to produce the type of output that is requested. The request for such output usually comes from prompts, often written by humans. Professional or semi-professional users of GenAI can be called prompt engineers – a booming business.

Some GenAI providers (Google, Microsoft and OpenAI) have provided "indemnifications" to users of their systems for copyright infringement. The indemnifications are reportedly backed by automated "filters" that will ensure that no infringing output is generated. Two questions come to mind. First, humans and courts often have difficulty determining what is and is not infringing. The idea that this can be entirely automated is thus unconvincing. Second, the legal text supporting the indemnifications for infringing outputs often contains significant exclusions.

For example, OpenAI's Terms of Service exclude indemnification for outputs which the "Customer or Customer's End Users knew or should have known the Output was infringing or likely to infringe, (ii) Customer or Customer's End Users disabled, ignored, or did not use any relevant citation, filtering or safety features or restrictions provided by OpenAI, (iii) Output was modified, transformed, or used in combination with products or services not provided by or on behalf of OpenAI, (iv) Customer or its End Users did not have the right to use the Input or fine-tuning files to generate the allegedly infringing Output, (v) the claim alleges violation of trademark or related rights based on Customer's or its End Users' use of Output in trade or commerce, and (vi) the allegedly infringing Output is from content from a Third Party Offering."[5]

Whether the "should have known" clause imposes a duty on users to check whether a particular output may be infringing is unclear, but the standard is certainly open to various interpretations. Moreover, excluding any material that the user *modified* is noteworthy, as many users are likely to tweak the machine's output. Even a format change might be sufficient to exclude the application of the protection. Indemnifications offered by Google and Microsoft also contain important limitations. Google's indemnity clause excludes customer uses "after receiving notice of an infringement claim", for example.[6]

5 https://openai.com/policies/service-terms, s 3(b), updated November 6, 2023

6 https://workspace.google.com/intl/en/terms/service-terms/, updated 20 November 2023

There is ample support, therefore, for the claim in a recent Forbes article that "if you read the fine print, the protections offered are narrower than what's suggested by the PR."[7]

Let us now turn to the international legal framework.

B. The International Legal Framework

This section provides an overview of the international applicable legal framework.

1. Relevant legal instruments, notions, and institutions

- Almost every country recognizes certain copyright rights, including the right of reproduction and the rights of adaptation and translation, the latter two sometimes referred to together (e.g., in United States law) as the "right to prepare derivative works".

- The Berne Convention is the main international treaty in the field of copyright (authors' rights). It contains an obligation for its 181 member States to provide the above-mentioned rights in their national law[8]. It is administered by the World Intellectual Property Organization (WIPO).

- For members of the World Trade Organization (WTO), the obligations contained in the Berne Convention are enforceable through the WTO dispute-settlement system, as most of the substantive provisions of the Berne Convention were incorporated into the 1994 Agreement on Trade-Related Aspects of Intellectual Property Rights (TRIPS Agreement)[9]. The TRIPS Agreement also includes a right of reproduction in phonograms (sound recordings) and excludes ideas from the scope of copyright protection[10].

- The Berne Convention and the TRIPS Agreement limit the ability of states and other parties to these instruments (for example, the EU itself is a party to the WCT and WPPT and is a member of the WTO) to restrict copyright under an important legal doctrine known as the three-step test. This test is discussed separately below in subsection C of this section.

- The Berne Convention contains a mandatory exception, namely the quotation right. It also excludes protection for "news of the day or [...] miscellaneous facts having the character of mere items of press information."[11]

- Copyright applies to literary, artistic, and musical works (also called "compositions"). Rights in musical works are originally held by the composer/lyricist/songwriter but are often shared with a publisher. Important forms of exploitation of musical works are administered by Collective Management Organizations (CMOs), including the right of public performance (live), the right of communication to the public (e.g., by broadcast or online transmission), and the

7 Brad Stone, "AI Legal Protections May Not Save You from Getting Sued", Forbes, 13 November 2023.

8 Berne Convention for the Protection of Literary and Artistic Works (1971), arts 8, 9(1),and 12.

9 TRIPS Agreement, art. 9.1 and 9.2.

10 TRIPS Agreement, art. 14.1 *in fine*.

11 Ibid. arts 10 and 2(8), respectively.

right of “mechanical reproduction”, a subset of the reproduction right that often takes the form of licensing copies for the purpose of distribution on devices such as vinyl records, CDs, and so on. CMOs administering rights in musical works are generally members of the International Confederation of Societies of Authors and Composers (CISAC).

- The reproduction right in literary works is often administered by CMOs known as Reproduction Rights Organizations (RROs). They are generally members of the International Federation of Reproduction Rights Organizations (IFRRO).

- The 1996 WIPO Copyright Treaty (WCT-115 parties as of this writing) protects rights beyond those contained in the Berne Convention. It provides an exclusive right to “make available” protected works online (described in the treaty as part of the right of communication to the public), as well as a distribution right applicable to copies of works[12].

- The WCT also provides rights against the removal of RMI and against the circumvention of technological protection measures (TPMs)[13].

- In addition to the rights just described, which focus on the economic exploitation of musical works, most countries recognize a moral right for authors of musical works, which protects them against misattribution or mutilation of their works.

- Musical works are very often exploited in the form of a recorded performance of the work. In most countries, these sound recordings are also protected by a related right. This related right is not copyright proper, but it is largely equivalent in that it also confers an exclusive right of reproduction and (for performers) a moral right. These rights are protected under the 1961 Rome Convention (97 member states as of this writing) and the 1996 WIPO Performances and Phonograms Treaty (WPPT--112 member states as of this writing). The United States is not party to the Rome Convention, but it is party to the WPPT. US law recognizes sound recordings as copyright works—that is, not (formally) as objects of neighboring rights--although without a full set of exclusive rights[14].

- The WPPT also contains a “three-step test” for limitations and exceptions to the rights it contains[15].

- A reproduction of a work does not have to be identical to a pre-existing work to infringe. Substantial similarity is sufficient, although in many countries any element of a work

12 WIPO Copyright Treaty (1996), arts 8 and 6, respectively.

13 Ibid, arts 11 and 12.

14 Sound recordings do not have public performance rights, but right holders have rights in respect of digital transmissions. 17 U.S.C. sec 106(6).

15 WPPT, art 16(2).

that is considered standard or routine may be copied without infringement. Reproduction of a sound recording, on the other hand, may require evidence that the actual sounds of a protected recording were used.

- International treaties, including the Berne Convention and the TRIPS Agreement, contain a "national treatment" obligation. This means that countries bound by these international instruments cannot discriminate against foreign rightholders (but they can treat them better than their own nationals). However, this obligation only applies to the rights protected by those instruments and only to nationals of other parties to the instrument. For example, the 1961 Rome Convention (mentioned above) contains an obligation to remunerate performers, producers of sound recordings, or both, for the broadcasting or communication to the public of protected sound recordings (which the Convention refers to as "phonograms"). Because the United States is not a party to this Convention, U.S. performers and producers cannot expect the same rights as nationals of countries that do provide this right.

- Licensing is a legitimate way to use copyrighted material, and some major companies have licensed the use of copyrighted works as data to train LLMs.

- Finally, it should be noted that the above list is not exhaustive. There are other potentially relevant instruments, such as the 1971 Phonograms (Geneva) Convention, but the relevant framework consists of the instruments and concepts described above.

2. Application of the International Framework to Generative AI[16]

Despite differences in national laws - in particular judicial and/or legislative restrictions on copyright - the creation and use of LLMs may go beyond what those restrictions allow. Moreover, as explained above, even after the "tokenized" dataset is created, in addition to the liability under international rules for removing rights management information, a partial copy of the original works remains in the tokenized dataset.

Responsibility for copyright infringement may lie with the people who train, sell, or use the machine, rather than with the machine itself, since the machine is not a legal entity. This will depend on the facts (who did what) and the rules of secondary liability in each jurisdiction, which are not harmonized under international law.

Let us take a deeper look.

(a) General considerations

Potential instances of liability for copies made during the machine learning process include copying the original "data" (i.e., copyrighted works) and creating and copying the tokenized dataset.

16 This White Paper does not address the issue of "machine authorship", that is, the idea that an LLM can autonomously provide the originality required to obtain copyright protection. I have explained elsewhere why I strongly disagree with this proposition. See The Machine as Author, 105 Iowa Law Review 2053-2106 (2020); The Human Cause, in Research Handbook on Intellectual Property and Artificial Intelligence (R. Abbott, ed), (Edward Edgar, 2022) pp 21-38; and for a shorter take on the subject, Humans as Prompt Engineers, Kluwer Copyright Blog, 14 June 2023. This does not mean that a human cannot an AI machine as a tool to create, as the above-mentioned Iowa Law Review article explains.

It is important to understand at the outset of the analysis that just because something is publicly available online does not mean that it is free to use. There may be cases, such as Creative Commons licenses, where the terms of use allow certain uses or waive economic copyright rights altogether, but this is a case-by-case determination. Therefore, training an LLM on publicly available material does not lead to the conclusion that the training was not infringing. Several jurisdictions, including the European Union, Japan, Singapore, and Switzerland, have enacted specific laws regarding the copyright aspects of machine learning. These laws provide guidance on what can and cannot be done legally without a license. In the United States, there is some uncertainty about the scope of the "fair use" exception. This uncertainty forces both copyright owners and users to hedge against possible negative litigation outcomes. The impact of LLMs on the market for copyrighted material and the ability of AI to create commercially competitive results is undoubtedly an important consideration in these discussions.

Regardless of how much creators receive in royalties or damages for the "ingestion" and potential ongoing copying in tokenized LLM datasets, this may be a one-time compensation payment for any massive Large Language Model (LLM) already trained, because there is only one human timeline to "ingest". This means smaller payments for ingesting future work.

Put another way, while compensation for "past sins" may be very significant-especially in jurisdictions such as the United States that provide for statutory damages-once the past (going back about 100 years) has been ingested and paid for, how can authors and their partners secure ongoing payments from large AI companies? If a few large models are offered as an "infrastructural layer" to smaller users, this may not require a new ingestion of the past, but only much smaller *ongoing payments* for the ingestion of future works. This would mean that human creators and their partners in the music industry (publishers, producers) might find themselves out of the main financial loops generated by LLMs. If LLMs are used to replace human creators, the result could be catastrophic for those creators and for all of us who depend on their work.

(b) Input liability

When applying international legal norms to the *input* (training) phase, two rules of the Berne Convention (incorporated into the TRIPS Agreement and this enforceable at the WTO) are relevant. First, is the taking of one or more "chunks" of a copyrighted work comparable to a quotation, which article 10 of the Convention clearly allows ("It shall be permissible to make quotations…")? The answer is no, because the chunking that occurs during the training of an LLM is actually the chunking of an entire work. Article 10 requires that citation be "compatible with fair practice". It seems difficult to argue that quoting an entire book, for example, is compatible with fair practice. Indeed, one could argue that cutting a book or other work into a series of snippets *representing the entire work* is hardly a series of "quotations".

Indeed, when article 10 was added to the Berne Convention (in 1948), an example given was "short quotations from newspaper articles and periodicals" and the General Report of the Revision Conference noted

that “only short fragments can be involved”[17]. Furthermore, a quotation must indicate the source and name of the author to be legal under the Convention[18]. Thus, the quotation exception does not apply[19]. The only door that remains open for an exception in national law is the second Berne rule that is most relevant in this context, namely the three-step test (discussed below). The quotation right may, however, be relevant, in the analysis of output infringement in the next section.

Before looking at liability for outputs, let us look briefly at some arguments used to justify the free use of copyrighted material for training commercial LLM applications. First, there is the argument that by tokenizing copyrighted works, the machine is using only the ideas, not the expression, contained in those works. Ideas are not protected by copyright. But as the technological overview makes clear, the opposite is true. It is the *expression* (words) that the machine copies. Another specious argument is that an LLM “learns” just like a human, and that humans do not infringe copyright when they learn.

This is not a useful comparison for at least three reasons. First, a human reading a book does not copy; the machine does, before and during its training. Second, copyright law imposes limits, and at some point (beyond what national law allows) requires a license when a human copies a book to learn from it. Third, when humans create, they do not simply regurgitate what they read in the work of other authors, which they can do (without permission) to the extent that it is a fair quotation. Human authors use their life experience and myriad other factors; the machine does not.

Finally, it should be noted that both the WCT and the WPPT require the existence of legal remedies against unauthorized removal or alteration of electronic RMI. This information is often removed during the machine learning process. This may constitute a separate cause of action against the maker of the generative AI application, although there is no “knowledge” element to establish liability. For example, the WCT limits the obligation to provide a remedy to acts done “knowing, or with respect to civil remedies having reasonable grounds to know, that it will induce, enable, facilitate or conceal an infringement of any right covered by this Treaty or the Berne Convention”[20]. This issue has yet to be fully litigated but a similar requirement under chapter 12 of Title 17 of the US Code is before several US courts[21].

(c) Output liability

There is some confusion in analyzing copyright liability for LLM outputs. This may be because several rights may be involved. The most obvious is the right of reproduction. Here, the basic analysis is straightforward: has the machine produced an output that is identical or substantially similar to a pre-existing protected work? If so, it infringes,

17 Berne Convention Centenary (WIPO 1986), at 180. The WIPO *Guide to the Berne Convention* (WIPO, 1978, at 58), defines quotations as “reproducing extracts from a work either to illustrate a theme or defend some proposition or to describe or criticize the work quoted from”.

18 Berne Convention, art 10(3).

19 Another potentially relevant limit is that the exception only applies to works “lawfully made available to the public”. Ibid, art 10(1).

20 WCT, art 12(1).

21 Chapter 12 is not part of the U.S. Copyright Act proper. It was added by the *Digital Millennium Copyright Act* of 1998.

unless the defendant can show that an exception applies or that it did not have access to the work.

The second right is the right to make "derivative works," such as adaptations or translations. For example, imagine an LLM producing a translation of the latest Booker or Goncourt prize-winning novel into, say, Chinese or Spanish. This is clearly a violation of the translation right (a subset of the derivative work right). It is also likely to be a violation of the right of reproduction and possibly the right of distribution-not to mention the moral rights of the author, especially if the derivative uses unattributed material.

There is little doubt that a court would both issue an injunction to prevent distribution of the unauthorized translation and, if appropriate, order payment of damages[22]. In all such cases (both reproduction and adaptation/derivation), rights in identified (or at least identifiable) works are being infringed.

There is an argument that any output of an LLM application that has learned from copyrighted material is necessarily "derivative" of the training material and therefore an infringement of the derivative work right. In this analysis, which I can only sketch out here, one must be careful to distinguish between the law in the United States, where the Copyright Act contains a somewhat unique definition of the term "derivative work", on the one hand, and international norms regarding translation and adaptation, on the other.

Although the idea that all outputs of the LLM are "derivative" is fully defensible under a colloquial definition of the term "derivative", it seems much less likely that courts will find under existing law that this colloquial definition is similar to the narrower legal notion of "derivative work". Indeed, courts that have already expressed an opinion have disagreed with this view[23].

3. The Three-Step Test

As noted in sub-section (a) above, parties to the Berne Convention, the WCT and WPPT and members of the WTO must follow the three-step test when adopting exceptions to copyright rights in their national law.

The three-step test originated in the last substantive revision of the Berne Convention in 1967[24]. Its purpose was to provide both a limit and guidance to legislators regarding exceptions and limitations but only to the right of reproduction. The importance and scope of the test were increased significantly in 1994 when the TRIPS Agreement adopted versions of the three-step test for exceptions to *all* copyright rights (article 13)[25]. The test now also appears in multilateral intellectual property instruments including the WCT and WPPT and in many

22 There is a peculiar debate in the United States as to whether a machine, which cannot be an author, can ever infringe the derivative work right, because some courts and scholars have taken the view that the derivative work right can only be infringed if the derivative work is itself an original work, and only human creators can provide this originality. I will not belabor the point but let me state that I disagree as I explain here: Gervais, Daniel (2022), "AI Derivatives: The Application of the Derivative Work Right to Literary and Artistic Productions of AI Machines", 52:4 Seton Hall Law Review 1111.

23 In November 2023, in a case brought by Sarah Silverman, a judge rejected the idea that any output created from a corpus of her material was derivative as a matter of (US) copyright law.

24 Berne Convention, art. 9(2). An Appendix for developing countries was added at Paris in 1971.

25 In addition to designs (article 26.2) and patent rights (article 30), and partly to trademark rights as well (article 17).

bilateral and regional trade agreements[26]. The test is enforceable in the area of copyright through the World Trade Organization's binding dispute settlement mechanism. Two cases involving the test have already been decided by the WTO, and the WTO found that the exceptions challenged in those disputes were inconsistent with the test. Thus, despite ongoing disagreements concerning the operation of the WTO, governments have tended to avoid adopting a new exception or limitation to intellectual property protection without considering the potential role of the three-step test.

It is not necessary to explain the test in detail here. Suffice it to say that an *exception or limitation* to authors' rights included in national law can be challenged at the WTO. The key part of the legal analysis revolves around the impact of the exception or limitation on *prospective* markets is critical. Interpreting the three-step test to apply only to established markets would discourage investment in new technologies and new markets. Conversely, to consider harm to any prospective market, no matter how remote, would render virtually all restraints and exemptions inconsistent with the second step. The test does not go that far. First, the interference must affect an income stream (whether actual or prospective) that is sufficiently close to normal exploitation. Second, the prospective market must be reasonably foreseeable.

The test also focuses on the extent to which the "legitimate interests" of authors are affected. Authors may have a legitimate interest, beyond a particular work, in being able to continue to work as authors and thus earn a reasonable living from their creative endeavors[27]. Since an exception in national law can be examined to determine its compatibility with the test, a systemic analysis of its impact not only on a particular work (as opposed to a national court case concerning the infringement of one or more identified works), but also on the broader interests of creators, seems warranted.

The test is relevant when restricting an existing right that must be made available under international law (such as the right of reproduction or the right of public performance), but not when creating a new right. Its impact is more likely to be felt when analyzing restrictions on copyright rights to enable training (text and data mining), although it would be relevant if rights applicable to AI outputs that infringe copyright (making a copy, adaptation, public performance or communication to the public of a protected work) were also restricted.

26 For example, art 11.18.3 of the Regional Comprehensive Economic Partnership (RCEP).

27 The author is grateful to Professor Jane Ginsburg for this insight.

Key Takeaways of Part I

1. Training a Large Language Model (LLM) also known as text and data mining usually involves the copying of the training data
2. If the training data consists of copyrighted material, this is a reproduction and an infringement of the copyright owner's rights unless an exception in national law applies, such as fair use (US)
3. The copyrighted material is "chunked" during the training into tokens, numerical representations of the material, but also embeddings or "vectors" that preserve all or part of the original works. This also constitutes a reproduction
4. If a LLM produces an output that is substantially similar to or derivative of one or more works used during its training, this also constitutes an infringement of the copyright owner's rights, unless an exception applies
5. The training of LLMs using copyrighted materials may involve the removal of Rights Management Information, which would constitute another violation of the copyright owner's rights
6. The indemnifications and defense obligations offered by major AI companies to users of their systems for possible copyright infringement contain significant limitations
7. Any exception to copyright rights in national (or regional) legislation to allow TDM must be compatible with the "three-step test"

II. Finding a Fair Way Forward for Creators

A. Authors' Rights Have Always Adapted to Technology

The history of copyright is one of constant adaptation to technological change. When authors' rights were first established, the underlying premise was to provide authors and their industry partners (publishers) with the means to live from the fruit of their labor by creating a viable marketplace for copies and later, for live public performances of music and theatre.

When the player piano was invented, the right of reproduction was recast accordingly. When radio was invented, the performance right was extended to broadcasting and later adapted to cable retransmission. When cinema was invented, this new category of "work" was recognized, as were computer programs a few decades later. Most of these changes were enshrined in the successive revisions of the Berne Convention, while others were reflected in the TRIPS Agreement[28]. When the World Wide Web tsunami hit authors and other copyright holders, the world came together very quickly in 1996 to adopt the WCT and the WPPT, reflecting the "making available" option that the Internet made possible by creating an exclusive right. These instruments also added a layer of protection for rights management information.

Placing all these adaptations on a timeline and then comparing them, it is clear that new exclusive rights (e.g., broadcast) and remuneration rights (certain cable retransmissions) were designed to ensure that authors would have a say, or at least a right to be compensated, for new commercially significant uses of their works. Indeed, most commercial uses of copyright material are subject to authors' rights, except in cases where a license is unlikely to be granted but there is a societal interest in allowing the use, such as parody.

Now the most profound technological change in history is upon us[29]. A technology that can produce commercially competitive content that is likely to displace some human-created works. It can do this because it has absorbed the works of human authors. The stakes could not be higher.

28 For a more complete history, see Daniel Gervais, *Restructuring Copyright: A Path Towards International Copyright Reform, Revised and updated edition* (Edward Elgar, 2019).

29 Strangely, voices opposed to any adaptation the current framework sometimes say in the same breath that current law is fine, and that AI is too big a change for copyright to be able to adapt.

Creators should be front and center as the normative foundation for the future development of the legal framework surrounding Generative AI. The survival of their economic agency is at stake, and that should matter to all of us.

Music, like other forms of art and literature, can help us understand our world. Delegating to machines the task of helping us understand and interpret our world has profound consequences. It is through this human interpretation and agency that humans can become true actors in the world, and ultimately change it. Delegating this very task to machines is therefore pregnant with implications for the future because it changes the arc of the future. In other words, if most of what we read, watch, or listen to comes from machines, much will be lost, and the loss may be irreversible, as established creators lose the ability to live off their work and fewer newer creators enter the fray. At the same time, AI promises much innovation and the production of new public goods such as new medicines.

Thus, the idea of "stopping" LLM technology seems both unrealistic and undesirable. But the development of LLM *must* ensure the continued existence of human creators using the same system that has worked for over 200 years, namely copyright[30].

30 This term is used here very broadly to refer to authors' rights proper, but also related rights and certain sui generis.

B. A Right to Remuneration

As explained above, music creators and their industry partners may receive compensation for training LLMs using copyrighted material as a result of court rulings or settlements in pending lawsuits. These can be very consequential payments. They could include some form of licensing for the use of future copyrighted content to train commercial LLMs that produce music and other forms of literary and artistic content. This was ostensibly the goal of the European Union in adopting Article 4 of the *Copyright in the Digital Single Market* Directive, namely, to encourage a licensing regime for commercial text and data mining for other copyright holders who wish to opt out[31].

However, establishing *ongoing payments* sufficient to compensate the creators of the material used to train the LLM that may then use that material to endanger the livelihood of those very creators, is critical. This may not be possible based solely on payments for "past sins" and smaller payments for ongoing use for training purposes. The possible infringement of copyright in outputs may be limited in cases where no existing right (reproduction, adaptation, translation) in specific works has been infringed. A license to retain a copy post training and in the tokenized dataset, with annual renewals, might capture some of the activity, but it might not capture what matters most in this context: the actual production of commercially usable "content" that competes with

31 Directive (EU) 2019/790 of the European Parliament and of the Council of 17 April 2019 on copyright and related rights in the Digital Single Market. PE/51/2019/REV/1, OJ L 130, 17.5.2019, p. 92–125, Recital 18 ("Rightholders should remain able to license the uses of their works or other subject matter falling outside the scope of the mandatory exception provided for in this Directive.")

creators in the marketplace. A smaller dataset used thousands of times to produce music should lead to more remuneration for creators than a large dataset used only sporadically to produce content destined for commercial use. Some recent research suggests that, specifically for music, a small-scale model may work well, so that only a relatively small number of tokenized musical works could be used to produce a large quantity of potentially competing material[32].

Moreover, although in jurisdictions such as the United States with high statutory damages for infringement of registered works, even a low probability of infringement may well justify a licensing arrangement, this may not be so obvious in other jurisdictions. Nevertheless, the sense of unfairness to all creators whose works are used to train the machine and then to compete with those same creators, and the presence of a massive societal risk, remain.

As the UK House of Lords noted in a February 2024 report:

"We do not believe it is fair for tech firms to use rightsholder data for commercial purposes without permission or compensation, and to gain vast financial rewards in the process. [...] The point of copyright is to reward creators for their efforts, prevent others from using works without permission, and incentivise innovation[33].

To address the inequity of using decades of work by human creators to create AI systems that may *compete with those same creators* for decades to come, without stifling innovation in AI, a licensing regime must be established that generates sufficient income for creators to continue to practice their craft. The legal question is what right the subject of such a license would be.

One possible way forward is to create a new right to remuneration *for the use of copyrighted material to create an LLM, which is then used to create competing content*. This would have several advantages:

- The right to remuneration would mean that the training of LLMs could continue essentially unhindered;
- Commercial LLM providers would pay for what may be the most valuable input (in addition to the people, computers, programming, electricity, etc. they already have to pay for);
- The right would not apply to organizations using LLM technology for research. For example, a university using a LLM to process a corpus of medical literature to identify potentially useful new molecules would not have to pay a fee; and
- Creators and their industry partners would be appropriately compensated for the use of their life work when the technology trained by their human creativity and labor is used to compete with them."

32 Liwei Lin, Gus Xia, Junyan Jiang, Yixiao Zhang, "Content-based Controls for Music Large Language Modeling", arXiv:2310.17162, 26 October 2023.

33 House of Lords, Large language models and generative AI, 2 February 2024, at paras 245-246.

A proposal along similar lines has been made by Martin Senftleben, namely a "single equitable remunerationl to be paid by providers of a generative AI system if the literary and artistic output generated by the system has the potential to serve as a substitute for a work[34]. Indeed, proposals to compensate creators of material used for training should focus on the substitution potential. The idea is not to prevent the training of language models, but rather to compensate creators when their own work is used to compete with those creators.

C. Application of the International Legal Framework to a New Right to Remuneration

There are several important legal considerations when considering the establishment of the new right. They are sketched out below.

1. Subject matter & right(s)

The right would apply to musical works (compositions) and could be limited to works protected by copyright. As a new right, it would belong to creators; indeed, the normative basis for the new right and the very reason for its existence depends to a significant extent on the ability of creators to retain their agency when their work is used by LLMs to produce content that may end up competing with them in the marketplace.

Authors would be able to transfer this right, whether to a publisher, a CMO or some other entity. As the above-mentioned House of Lords report states: "The [IP] code must ensure creators are fully empowered to exercise their rights, whether on an opt-in or opt-out basis"[35].

The new right would *not* apply to copying for training (TDM or "input phase") as this matter is already covered by existing law and will become clearer as courts determine the liability of GenAI companies in this regard, including the interpretation of fair use in the United States and articles 3 and 4 of the CDSM Directive in the EU. The new right would apply to *the use of the tokenized dataset to create material that can be used to compete with the material used for training, defined as material of the same nature that can be generated for or made available to the public*. For greater efficiency, a single licensing agreement could apply to the training (TDM) *and* output. This new right recognizes that, as a matter of fairness, LLM outputs are "derived" from the training dataset, but in most cases are unlikely to be considered adaptations or "derivative works" – unless these terms are reinterpreted or even redefined. Thus, AI companies would pay the creators, to use their works to create material that can be used to compete with those creators. Non-commercial research uses would be exempted.

The new right is compensation for the use of works, not a license permitting outputs that are substantially similar to, or direct adaptations or translations of, the works used for training. Allowing unrestricted

34 Martin Senftleben, *Generative AI and Author Remuneration, International Review of Intellectual Property and Competition Law 54* (2023), pp. 1535-1560 (revised February 2024).

35 Ibid, at para 252.

modifications or new material generation by AI that is substantially similar to protected copyright works could infringe the moral rights of a creator if the content produced by AI is not attributed (if copied) or is misattributed (if modified) to the creator of the work so copied or modified.

2. National treatment/non-discrimination

If the right is established as a separate right from copyright, i.e. as a "*sui generis*" right, it would not be subject to national treatment obligations. Instead, countries could opt for material reciprocity, i.e. they would only pay rightholders from countries that have a similar right. This could create an incentive for those countries to do so. If the right is established under copyright law, it would be subject to national treatment obligations (non-discrimination against foreign right holders). An analysis of the exact scope and reach of national treatment obligations under the Berne Convention and the TRIPS Agreement is beyond the scope of this White Paper.

3. Compulsory license

If the proposed right was established as a *sui generis* right to remuneration, that would be valid under international law. If established as a copyright right (i.e., for the use of copyrighted works), a compulsory license could also be defended on several grounds. Many countries have had compulsory (or statutory) licenses in place for years without challenge. State practice is relevant to the interpretation of international obligations. Moreover, compulsory licenses existed when the Berne Convention was last revised and were not made illegal. On the contrary, the Convention explicitly allows them in several cases. Finally, and perhaps more importantly, this new right would be defensible as an *extension of authors' rights* rather than a limitation of an existing right.

4. Comparison with a levy system

One could consider instead a levy system based on the amount of material copied at the input stage, as was done many years ago for private copying in a number of countries. A levy system seems a second best for at least two reasons. First, and most importantly, the levy system would have to find a way of responding to the *actual use* of the LLM to produce commercial output, as opposed to the amount of data used for training, since the former seems more directly correlated to the market impact on creators. Second, a levy would necessarily require the use of a proxy for distribution, while the proposed new right opens up the possibility of getting usage data, as explained below. Put differently, a major justification for a levy system may be that it is impossible to obtain output data linked to any material in the training data. However, this is not entirely true.

5. Distribution

An argument often used against *any* right to remuneration or compulsory license is that distribution cannot be done effectively and fairly. This argument is important because the credibility of a new right would depend in part on its ability to reach its intended beneficiaries.

CMOs could distribute funds based on metrics to be determined. Ideally, usage data would be made available for this purpose. Many LLMs *can* be programmed to identify the source material. It is true that this may reduce the efficiency of the systems and increase their cost, and will be fought by the AI industry. Rather than fear-mongering, a conversation should be had about the real possibilities of the technology, which is simultaneously presented as extraordinarily powerful and yet incapable of producing this kind of data, and about the real costs and efficiency losses that might be involved. Concrete options include using a second AI, which may be called a compliance or detection agent, that would "report" on the output produced by the LLM.

There are many other examples and cases in society where producers of a particular technology have been required to take certain actions. For example, decades ago, we required car manufacturers to put catalytic converters on cars for public health reasons. But catalytic converters make cars less efficient and increase costs and the car industry fought them on those precise grounds. When a technology threatens the livelihoods of millions of human creators, and potentially prevents the emergence of a new generation of creators, it doesn't seem unreasonable to suggest that a similar approach is warranted.

However, the creation of a system to compensate creators could also be based on appropriate proxies. There have been several attempts over the past decades to distribute funds generated on behalf of creators, such as private copying levies. There have unquestionably been deficiencies in some cases in the distribution of those funds, but lessons were learned. Credible proxies can be created. Successful proxies exist for private copying and public lending for example, factoring in commercial success and availability.

A better solution still would be to base distribution on more data more directly connected to use. Here, part of the solution depends on transparency. If the data on copyrighted material used to create an LLM is made available (even on a confidential basis), this would provide crucial data points.

As noted above, the new right should vest initially in creators. As also explained above, ensuring that the monies reach individual creator is a key driver of the proposal. Any distribution system used should reflect this and ensure proper attribution. In the European Union, the obligation in the AI Act to provide information about copyrighted material used for training is a step in the right direction. The possibility of imposing an obligation to provide more granular data with appropriate safeguards should be considered in every jurisdiction.

Usage data could be generated (but not necessarily made available to a user[36]). This is a technological question that would need to be examined at the appropriate point in time in the future, as the exact way in which technology will be developed cannot be determined with certainty. However, it is possible to program at least some LLMs to identify the sources used to produce a certain output. If a licensing

36 The determination of specific works that the LLM drew from to create an output in response to one or more prompts/finetuning may be relevant for the protection of moral rights (and possibly an application of the attribution requirement to mention a source contained in article 10 of the Berne Convention when strings of words etc. are reused). This is a matter that a license for input (training) could regulate.

framework offered protection to users (producers and users of LLMs), then that licensing agreement could provide an obligation to provide granular data. The exact nature and scope of the granularity required would depend on technological evolution and a better understanding of available options depending on the model and technology used. Providing adequate data to inform distributions would be part of the license negotiation process. The data could be aggregated to protect the trade secrets of the AI provider, if applicable. It could identify how often each tokenized work was "pulled", but not by whom. Absent such data, the best option is probably to provide a list of all works used for training (which must be available) and devise an appropriate distribution model based on outputs, other uses (eg streaming) etc.[37]

Finally, arguments that creators will not get "enough" to justify such a right, or that it is simply "too complicated" to administer, do not hold water. First, not knowing how much will be generated by the new right or how much a creator considers "enough" means that the first argument is merely an unsupported assertion. In collective management contexts, there is always a "long tail" of creators who receive very little. But by far the most widely distributed and accessed productions are those of creators who receive significant amounts when all forms of exploitation of their works are combined and remunerated. As AI training and substitution continue to develop, they will become a major form of commercial exploitation of the work of creators. Second, if data can be generated by the LLM (or a second AI working in tandem with it) as explained above, then usage data can be easily processed, for example by a CMO. Even if proxies are used, CMOs have demonstrated their ability to manage creators' rights over several decades. Pointing to mistakes made is not a rebuttal. What industry has not made mistakes? The CMO system demonstrates that monies can be collected, data processed, and funds distributed worldwide.

37 As is the case with other distribution schemes, it may make sense for CMOs to keep a small portion of the funds (10%) to help creators deal with the transition to a world dominated by AI. This could include education and training on how to use LLMs, which can be used to create autonomously but can also be used as collaborative tools by human creators.

Daniel Gervais is a professor of law and director of the Intellectual Property Program at Vanderbilt Law School in Nashville, TN. He has spent 10 years researching and addressing policy issues on behalf of the World Trade Organization, the World Intellectual Property Organization, the International Confederation of Societies of Authors and Composers, and the Copyright Clearance Center. He is also the author of *The TRIPS Agreement: Drafting History and Analysis*, a leading guide to the treaty that governs international intellectual property rights.

CIAM
MUSIC CREATORS OF THE WORLD

The International Council of Music Creators (CIAM) was created in 1966 to advocate for the cultural and professional aspirations of music creators from all repertoires and regions. Focusing on economic and legal interests, CIAM serves as a forum for the exchange of ideas, information and practical advice. CIAM also has alliances in Latin America, Africa, Asia-Pacific, Europe and North America which lobby for the interests of music creators in their regions.

www.ciamcreators.org

ALCAMúsica AMA APMA ecsa MCNA

INTERNATIONAL
FAIR TRADE MUSIC
INTERNATIONAL

Fair Trade Music International is an independent non-profit founded in 2014 and backed by more than 500,000 music creators from around the world. Led by an international board, it combines global networking, education and awareness initiatives to promote the emergence of an ethical, sustainable and transparent music ecosystem that is equitable for everyone in the music value chain.

www.fairtrademusicinternational.org