

RESEARCH ARTICLE

Copyright in generative deep learning

Giorgio Franceschelli¹  and Mirco Musolesi^{1,2,3,*} 

¹Department of Computer Science and Engineering, Alma Mater Studiorum Università di Bologna, Bologna, Italy

²Department of Computer Science, University College London, London, United Kingdom

³The Alan Turing Institute, London, United Kingdom

*Corresponding author. E-mail: m.musolesi@ucl.ac.uk

Received: 24 June 2021; **Revised:** 04 October 2021; **Accepted:** 13 April 2022

Key words: deep learning; intellectual property; machine-generated artworks

Abstract

Machine-generated artworks are now part of the contemporary art scene: they are attracting significant investments and they are presented in exhibitions together with those created by human artists. These artworks are mainly based on generative deep learning (GDL) techniques, which have seen a formidable development and remarkable refinement in the very recent years. Given the inherent characteristics of these techniques, a series of novel legal problems arise. In this article, we consider a set of key questions in the area of GDL for the arts, including the following: is it possible to use copyrighted works as training set for generative models? How do we legally store their copies in order to perform the training process? Who (if someone) will own the copyright on the generated data? We try to answer these questions considering the law in force in both the United States and the European Union, and potential future alternatives. We then extend our analysis to code generation, which is an emerging area of GDL. Finally, we also formulate a set of practical guidelines for artists and developers working on deep learning generated art, as well as some policy suggestions for policymakers.

Policy Significance Statement

This work provides a comprehensive analysis of the legal issues concerning generative deep learning, including the storage of protected works, for training neural networks, the potential plagiarism risks associated with both the models and the outputs, and the problem of ownership of generated artifacts. Moreover, this work provides researchers with a list of practical and actionable suggestions for ensuring compliance in their daily work. At the same time, it offers an in-depth analysis of several critical points in copyright laws to policymakers. These issues will be central to the debate around these themes in the coming years.

1. Introduction

In recent years, we have witnessed an exponential growth of interest in the field of generative deep learning (GDL; Brown et al., 2020; Devlin et al., 2019; Goodfellow et al., 2014; Isola et al., 2017; Karras et al., 2019; Ramesh et al., 2021; Van Den Oord et al., 2016a; Yu et al., 2017, among others).

GDL is a subfield of deep learning (Goodfellow et al., 2016) with a focus on generation of new data. Following the definition provided by Foster (2019), a generative model describes how a dataset is generated (in terms of a probabilistic model); by sampling from this model, we are able to generate new

data. Nowadays, machine-generated artworks have entered the market (Vernier et al., 2020), they are fully accessible online,¹ and they have the focus of major investments.²

Privacy problems in generating faces or personal data are in the spotlight since few years (a technical survey for privacy in visual monitoring systems can be found in Padilla-López et al. (2015); the use of differential privacy (Dwork and Roth, 2014) has opened a new field of research (Liu et al., 2019; Wu et al., 2019; Xu et al., 2019; etc.), but also privacy in clinical data sharing is considered (see, e.g., Beaulieu-Jones et al., 2019). Ethical debates have, fortunately, found a place in the conversation (for an interesting summary of machine learning researches related to fairness, see Chouldechova and Roth (2020)) because of biases and discrimination they may cause (as happened with *AI Portrait Ars* [O’Leary, 2019], leading to some very remarkable attempts to overcome them, as in Xu et al. (2018) or Yu et al. (2020)).

However, copyright issues—the focus of this article—also arise from the use of deep learning to generate artworks. GDL techniques are trained on a large amount of data, which may be protected by copyright (especially in case of artworks); moreover, they allow for the generation of new samples similar to the training data, which in turn could claim copyright protection (or could represent plagiarism of other works). In this context, it is possible to identify at least three problems: the use of protected works, which have to be stored in memory until the end of the training process (even if not for more time, in order to verify and reproduce the experiment); the use of protected works as training set, processed by deep learning techniques through the extraction of information and the creation of a model upon them; and the ownership of intellectual property (IP) rights (if a rightholder would exist) over the generated works. Although these arguments have already been extensively studied (e.g., Sobel (2017) examines use as training set and Deltorn and Macrez (2018) discuss authorship), this paper aims at analyzing all the problems jointly, creating a general overview useful for both the sides of the argument (developers and policymakers); aims at focusing only on GDL, which (as we will see) has its own peculiarities, and not on artificial intelligence (AI) in general (which contains too many different subfields that cannot be generalized as a whole); and is written by GDL researchers, which may help provide a new and practical perspective to the topic. For doing so, we first provide the reader with a brief review of what GDL means (Section 2). We then present an in-depth discussion of the existing legal frameworks in both the United States and the European Union in relation to storing protected works for using them as training set (Section 3) and, then, owning the copyright of machine-generated works (Section 4). In addition to legal analysis, we present a policy discussion that leads to practical suggestions for artists and developers working in this new field, as well as to considerations about the adequacy of current laws. Finally, even if the main case study will be about artworks, which are by a long margin the main (potentially copyrighted) outputs generated with deep learning, in the last section, we discuss the issues around GDL in relation to source code generation (Section 5).

2. Generative Deep Learning

As with other machine learning algorithms, GDL is based on the availability of a training set, which is a set containing examples of the entities that will be our starting point for the generation of new data. The machine learning model itself has to be *probabilistic* in order to generate new data (if it is deterministic, it will generate each time the same entity). A generative model is intrinsically different from a discriminative one, which represents the most popular (and classic) technique in machine learning. In discriminative modeling, for example, classification or logistic regression, we aim at distinguishing among classes after training on labeled data (it is, in a way, a synonym of supervised learning). Instead, in generative modeling, we use the (unlabeled) dataset to learn to generate new data.

¹ <https://botnik.org/>.

² <https://www.nextrebrandt.com/>.

In general, Foster (2019) defines the generative modeling framework as follows: given a dataset of observations X , and assuming that X has been generated according to an unknown distribution p_{data} , a generative model p_{model} is used to mimic p_{data} in such a way that we can sample from p_{model} to generate observations that appear to have been drawn from p_{data} (and that are suitably different from the observations in X).

GDL is based on the concept of representation learning. The algorithms are used to learn representations of the data that make it easier to extract patterns from them. Differently from classic machine learning techniques, representation learning works directly on unstructured data, not requiring to manually define the features of the input data (necessary to work with structured data). This adds an extra level of autonomy to the algorithm. Deep learning in itself is based on the composition of multiple nonlinear transformations (using multiple stacked layers of processing units) with the goal of yielding more abstract, and ultimately more useful, representations (Bengio et al., 2013). X is typically composed by unstructured data like images or texts.

Several generative models based on deep learning have been proposed in last years (for a more exhaustive list, see Franceschelli and Musolesi (2021b)). However, it is possible to identify a few families of techniques: Variational Auto-Encoders (VAE; first proposed by Kingma and Welling (2014), but also used by Gregor et al. (2015)); autoregressive models (first proposed by Van Den Oord et al. (2016b), but also used by Brown et al. (2020); Devlin et al. (2019); Huang et al. (2019); Parmar et al. (2018); Payne (2019); Radford et al. (2019); Van Den Oord et al. (2016a, 2016c)); sequence prediction models (very well explained by Karpathy (2015), and then used by Jaques et al. (2016); Lau et al. (2018); Potash et al. (2015); Sturm et al. (2016); Yi et al. (2018); Zhang and Lapata (2014); Zugarini et al. (2019)); and Generative Adversarial Nets (GAN; first proposed by Goodfellow et al. (2014) and then used by Vernier et al. (2020); modifications of the first proposal have been done by Brock et al. (2019); Dong et al. (2018); Elgammal et al. (2017); Engel et al. (2019); Isola et al. (2017); Karras et al. (2018, 2019); Yu et al. (2017); Zhang et al. (2019)).

We refer the interested reader to the original papers for a detailed description of the technical aspects of GDL given the scope of this article. The first three techniques have in common the fact that the generator is trained directly on the dataset X of examples (with a very small but not important exception for VAE, since the real generator processes only an encoded version of the original work, and not its real version). On the other hand, the generator of a GAN never sees the real examples: it is trained with the so-called adversarial training. A GAN is composed by two networks: a discriminator and a generator (that is, the generative model). The generator never sees real data, while the discriminator is trained on them, with the goal of learning to distinguish between original data and generated ones; at the same time, the generator tries to generate data that fool the discriminator (so, the goal is to generate data that are recognized as real). The learning process proceeds with the two phases alternatively: when the discriminator becomes better in recognizing original data, the generator has to learn more things about their distribution to fool again the discriminator; and when the generator becomes better in generating seemingly real data, the discriminator has to become better in distinguishing them. At the end, the generator should have learned the model correctly, without training on it (Goodfellow et al., 2014). In any case, data are always necessary as in any other machine learning algorithm. We will discuss the implications of this aspect later in this article.

Another very interesting GDL architecture is Google's DeepDream (Mordvintsev et al., 2015). The designers of DeepDream trained a convolutional neural network for image classification, and then they used the network as a generative model working in reverse: the image is not obtained as an output, but through modifications of an image provided in input, where the modifications are done according to the patterns learned by the network. In particular, the input image is modified in such a way that the "excitement" of a selected layer is maximized; this entails the emergence of the patterns that the specific layer has learned to recognize, which might just be simple ornament-like patterns, but also more sophisticated features.

As it is apparent from the overview given above, GDL has very specific characteristics, which have direct implications for copyright. First of all, GDL techniques are trained on large training sets from which information is extracted. These training sets may contain data, which are protected by copyright,

especially if they are artworks. In fact, differently from other learning techniques, they directly learn from the whole work as it is. It means researchers need to store potentially protected works in their expressive form, which is what is protected by copyright, and not by means of features or small chunks if not of their underlined ideas. Second, because of the nature of the generated outputs (especially against the nature of outputs generated by other kinds of AI (Colton, 2008), which simply follow rules explicitly provided by the programmer). They only depend on the technique used and the knowledge acquired, without any kind of human intervention (except for the coding of the program and the choice of input data), making very difficult to determine who shall be considered as the owner of the copyright.

3. Storage of Protected Works and Use for Training

Our analysis starts considering if the storage, reproduction, and therefore the use for training of a protected work by a GDL algorithm violate the copyright, or if it is allowed by US and EU laws.

Sobel (2020) identifies four different categories of uses of data performed by machine learning:

1. uses involving training data not protected by copyright, including works fallen into the public domain (not protected by economic rights anymore);
2. uses involving copyrighted subject matter released under a permissive license or licensed directly from rightholders;
3. market-encroaching uses (whose purpose threatens the market of those data); and
4. nonmarket-encroaching uses (whose purpose is unrelated to copyright's monopoly entitlement).

In the first case, there is no problem in storing and using a work not protected by copyright for this goal. This also applies for works now in the public domain, which happens in the European Union 70 years after the author's death (or the death of the last of the authors), and in the United States 95 years after publication date (if created and published before 1978; otherwise, 70 years after the author's death). The same is also true if the work is protected, but it has been acquired digitally through a license agreement that does not expressly prevent a reproduction with this goal (the second case). Otherwise, for protected works on which we have a lawful access but not in digital form or not for a reproduction (third and fourth cases), the question remains open. To address it, the next two subsections consider it under the US law (3.1) and under the EU law (3.2).³ Finally, the last subsection considers additional issues related on the *outputs* of generative models, and not on their *inputs* (3.3).

3.1. US law

The US Code establishes that the reproduction of a copyrighted work can be allowed if the use can be considered a fair use of the work (17 U.S.C. §107—Limitations of exclusive rights: Fair use). This provision sets the criteria used to determine if the use is fair, that is, the purpose of the use and its economic character, the nature of the work, the amount and substantiality of the portion used, and the impact of the use over its potential market. With these criteria, the law does not state unambiguously what is a fair use and what is not; it provides parameters on which Courts can base their decisions about the fairness of a use. This unpredictability has been criticized, not only because it requires a case-by-case analysis (Netanel, 2008)—and eventually to hire a lawyer (Lessig, 2004)—due to its nature of standard more than of rule (Carroll, 2007), but also because the four factors may fail to drive the analysis and may be instead used to support an independent and antecedent conclusion (Nimmer, 2003). However, the fair use doctrine has also some remarkable strengths. It ensures that two competing public interests are balanced: to incentive

³ Under Berne Convention, works with a country of origin which is a Union country benefit, in all other Union countries, from the same protection as the latter gives to the works of their own nationals. This means that the protection is governed by the laws of the country where protection is claimed, for example, an EU research center should be concerned with EU laws.

the creation of new works, and to improve the public's ability to use or access it (see *Sony Corp. v. Universal City Studios, Inc.*, 464 U.S. 417). In its form, this doctrine helps exclude uses only where exclusivity promotes social welfare (Lunney, 2002). In addition, even if it seems unpredictable, fair use cases tend to be more coherent than expected and could be organized into clusters, which can help in Courts' decisions (see Samuelson, 2009).

Anyway, a deeper analysis of these four criteria in the special case of (generative) deep learning is now necessary. As regards the amount and substantiality criterion, although in GDL we aim at using the entire work (maybe divided into portions, if it is too long like a novel or a song), we have to consider that machine learning refers to the use of ideas, principles, facts, and correlations contained in data given in input; since copyright aims to protect original expressions and not ideas, procedures, or methods, data mining techniques do not use copyright works as works per se but to access the information stored in them, and so the use is not substantial—and neither they constitute a copyright infringement in theory (Kretschmer and Margoni, 2018). Furthermore, Sag (2019) underlines that a *nonexpressive* use such as the one discussed above should be considered fair, since it is just about deriving from them meta-level information and not benefiting from their original expression; however, generative techniques could fall under the definition of *expressive* use, since they could use authors' copyrighted expressions (Sobel, 2017), learning from their creative and expressive choices (Bonadio and McDonagh, 2020). This would invalidate the insubstantiality theory. In addition, we need to consider that the single protected work is used alongside a large number of other protected works: it is rare that the result resembles one of them substantially, presenting its distinctive features. For this reason, the impact of its potential market is typically very small, because it becomes difficult to connect the generated work with the protected ones used during training (this is true in particular in the case of heterogeneous training sets; of course, if the training set is composed by few works from a single author, this consideration is not more valid). The economic character has to be seen considering that this exception is fair only for purposes like research, and so without a real economic character; however, the previously done distinction between market-encroaching uses and nonmarket-encroaching uses acquires significance in dividing between an (almost) sure fair use and a dubious case. Finally, when analyzing the purpose of the use (and its fairness), one needs to consider if it is transformative, the most common reason to assess fair use (Asay et al., 2020): if it adds something new, with a different character, which does not substitute for the original use of the work, the use is more likely to be considered fair.⁴ In particular, the key question to determine a fair use is nowadays if the work is used for a different expressive purpose from that for which the work was created (Netanel, 2011). It is not straightforward to assess this for GDL, but this could eventually be the case of GDL techniques which aim to add a novelty degree in their production (as in Elgammal et al., 2017). Finally, it is important to highlight that, if the use does not fall in the category of fair use, any sort of reproduction of copyrighted works is not permitted,⁵ usually even if it is only on volatile RAM (as judged by many courts, as reported in U.S. Copyright Office, DMCA Section 104 Report 118; however, the Report suggests that in some cases the reproduction of copies on volatile RAM may be considered as a fair use).

As seen, even if without a clear and unambiguous answer, the fair use doctrine certainly offers a support for machine learning researchers (maybe even too much, in comparison with what is allowed to humans [Grimmelmann, 2016a]). For instance, in the field of text and data mining (TDM), as pointed out by Sag (2019) and confirmed by two famous cases (see *Authors Guild, Inc. v. Google, Inc.* and *Authors Guild, Inc. v. HathiTrust*), the transformative nature of the involved process makes it possible to claim fair use. However, it is not straightforward to extend it to GDL, since the use tends to be more expressive than nonexpressive (which makes it more difficult to be considered as fair [Sobel, 2017]). We can only wait until more clarifications (and Courts' decisions) arrive to unravel it.

⁴ <https://www.copyright.gov/fair-use/more-info.html>.

⁵ <https://www.copyright.gov/help/faq/faq-digital.html>.

3.2. EU law

In the European Union, instead, there are two articles from Directive 2001/29/EC that can or cannot permit the reproduction of protected works. In particular, Article 2 states there shall be for authors the exclusive right to authorize or prohibit direct or indirect, temporary or permanent reproduction of their work by any means and in any form, in whole or in part; while Article 5(1) provides an exception to this (exclusive) reproduction right (an exception represents a situation in which a right is not reserved—as it usually is). In fact, it states there shall be an exception from the reproduction right in Article 2 for temporary acts of reproduction which are transient or incidental and an integral and essential part of a technological process, and whose sole purpose is to enable a lawful use of a work which has no independent economic significance (therefore excluding market-encroaching uses). This exception is quite similar to the one provided by the US Code, since it requires no independent economic significance (the third case by Sobel (2020) is probably excluded; reproductions have an independent economic significance if their use generates an additional economic advantage (Triaille et al., 2014), and it seems indeed the case: a neural network, deprived by some of the training data, would probably perform worse) and that the use is lawful (a use is lawful if it is authorized by the rightholder or not restricted by law; we will analyze soon if the use done by GDL is lawful or not), but it adds another constraint: the reproduction must be transient or incidental. Following Riedo (2019), this exception could theoretically be applied to TDM activities (as long as TDM is a lawful use), since during the process data are stored in RAM and then erased when turned off, and so the reproduction is transient. For Schonberger (2018), Article 5(1) arguably covers copying the works for the training process, deleting them at the end of the process. Despite this, it is very likely that the miners, especially if they are researchers, have to retain the data corpus for verification, aggregation with new datasets and further analysis (Chiou, 2019), deleting it only once their work is completed and published; and therefore storing them until the end of the research is very unlikely to be considered under the exception for temporary reproduction.

To sum up, it is difficult to say that TDM activities are allowed under this exception—or at least there is a degree of uncertainty (Geiger et al., 2018b; Rosati, 2018). In addition, Article 5(1) requires the use to have no independent economic significance, therefore excluding for-profit research and private companies. This surely would put the European Union's competitive position as both a research and an industrial area in danger (Geiger et al., 2015), and this is one of the main reasons behind the adoption, in 2019, of a Directive “on copyright and related rights in the Digital Single Market.”

This fundamental Directive tries to answer to all the questions examined above by the means of two Articles.

Directive's Article 3 states there shall be an exception to allow reproductions and extractions of lawfully accessible protected works for performing TDM, if it is made by research organizations and cultural heritage institutions (for the purposes of scientific research and as long as the copies are stored with an appropriate level of security). Notably, the Article states the copies may be retained for the time required for the purposes of scientific research, including for the verification of research results. However, we have to remember that the Article only asks for an exception to the *reproduction* right (i.e., the exclusive right to make direct or indirect, temporary or permanent reproduction of the work by any means and in any form) and the *extraction* right (i.e., the exclusive right to permanently or temporarily transfer all or a substantial part of the contents of a database to another medium by any means or in any form). The Article is not about the *making available* right. On the contrary, it is common in scientific research to make source materials available in order to allow others to verify and repeat experiments. Since this is about the making available right and not the reproduction right (Geiger et al., 2018a), Directive's Article 3 does not allow to publish protected works used during training. In principle, this appears to be correct: the researcher has lawful access to the works, whereas others may not. Making them (or derived versions of them) available means providing others access even if any terms or conditions have not been agreed. However, in practice, this means the verification of research results is not promoted, since it can only be performed by the researchers themselves (Geiger et al., 2018b, among others, highlighted this in vain before Directive's adoption). In this direction, a

good compromise appears to be Article 60d of the German Law on Copyright and Related Rights, which allows the making available of the (normalized and structured) dataset to a “specifically limited circle of persons for their joint scientific research, as well as to individual third persons” for quality assurance (Geiger et al., 2018b). Without a statement like this, in case of a research based on not publicly available data, the only (lawful) way to allow the verification of results will be to provide all the necessary information about the data used⁶ and all the preprocessing steps carried out on them or, even better, the related source code.

In addition, Article 4 states there shall be an exception or limitation to allow reproductions and extractions of lawfully accessible protected works also by other people or institutions, only for the time necessary for the purposes of TDM. Crucially, this exception or limitation is applied only if it has not been expressly reserved by their rightholders in an appropriate manner.

To summarize, the Directive includes the use of a (lawfully accessible) protected work for training among the lawful uses: it allows for research organizations to use it for TDM. In addition, also other entities can do the same, provided that its rightholder has not expressly reserved this right.

Article 3 is undoubtedly able to foster innovation and scientific research. Even if it only adds an exception for the reproduction right and not for the making available right, making it difficult to reproduce and validate research experiments, in our opinion, it is a good compromise between protection and innovation. In addition, Article 4 represents a positive contribution, at least from a theoretical perspective. It is able to encourage innovation also in private environments, avoiding the risk of losing considerable investments (see Hilty and Richter (2017) for a discussion of compelling motivations for opening TDM to private entities). In our opinion, leaving the possibility of reserving this exception to rightholders is essential from an ethical perspective (we will talk about it in Section 5). However, many questions arise when trying to operationalize this Article. As highlighted by Rosati (2019), private developers who want to use protected works to train generative models have to follow these three steps: obtain a lawful access to the data; check if rightholders have not reserved the right to make reproductions for TDM purposes; and retain any copies made only for as long as is necessary for TDM purposes. The first and third steps appear to be reasonable; instead, in our opinion, the second step presents some issues. How could an EU-based developer know if this right has been reserved for a certain protected work? In Recital 18, the Directive suggests reserving those rights through machine-readable means (e.g., metadata and terms or conditions of a website or a service) in case of publicly available content, and contractual agreements or unilateral declarations for other contents. Even if the list of means that is provided might appear as exhaustive, the issue does not appear to be addressed in a satisfactory manner. As we have seen in Section 2, a generative model is typically trained on a very large dataset. In other words, this might translate into the practical solutions of (a) having online databases that allow to filter (by means of metadata) the available works depending on this reservation or (b) directly publishing datasets only composed by *reservation-free* works, making sure at the same time they can be integrated with the reserved ones for research purposes. But are these providers forced to do so? Or will this checking activity fall on developers (in a way, dissuading them by training generative models [Chiou, 2019])? Finally, we have also mentioned Google’s DeepDream, which only requires one image in input. The same can apply to poems or other texts, which can be fed in input to emerging techniques such as VQ-GAN + CLIP (Esser et al., 2021; Radford et al., 2021) to obtain *alien dreams* (see Snell, 2021). Here, only one protected work may be used, not extracted by a database provider, but independently acquired. Will the transaction report if this right has been reserved or not? How would it be possible to discover if a work (with an access acquired before 2019) can be used or not? We believe that there is an urgent need of addressing these questions before the TDM exception could be applied in all of its strength.

⁶ An example is the proposal of Data Cards (Geburu et al. (2018)), whose aim is to address the so-called *documentation debt* (Bandy and Vincent (2021)), which can also have negative consequences from an ethical perspective (Bolukbasi et al. (2016)). This is impacted by the impossibility of making the dataset available.

3.3. *Additional issues*

We should also remember from [Section 2](#) that GDL involves the creation of a probabilistic model describing data of interest, from which we obtain new works through sampling. This peculiarity leads to other critical questions. In fact, the model creation and storage, which contains the extracted probabilistic features, may infringe the copyright of works used for training. Its storage cannot be considered as transient or incidental; therefore, it is allowed only if it does not constitute a (partial) reproduction of protected works, since, in that case, there is no copyright relevant activity (Margoni, 2018).

Deep learning models are usually stored as sets of numerical weights; usually, they do not fall in the category of those that are considered as a partial reproduction of a work. However, if the model is built to mimic in output the input (with some nonsubstantial changes) and it is trained over a protected work, the model would represent that work, at least partially. Moreover, even if the model is not intentionally built to mimic a protected work, it could still end up doing so to an infringing degree: it might reconstruct idiosyncrasies of input data instead of reflecting underlying trends about them (Sobel, 2017). If just trained with the goal of learning how to reproduce works, an overfitted generative model may actually be considered as a direct reproduction of the works.

These issues are not only related to the model itself, but also to the output it can generate. Protected input data are commonly used to train models to generate similar output. Then that output may infringe copyright in the preexisting work or works to which it is similar to (Sobel, 2017). In this context, the importance of adversarial training from GANs and searching for diverging from existing works as done by Elgammal et al. (2017) may tip the balance toward legality. With respect to other classic techniques, the generative part of a GAN never uses protected data; therefore, it is harder to obtain an output representing an expression of protected data. In addition, new techniques with a novelty objective, which tries to increase the distance between outputs and training data, will tend to be more transformative. This may be crucial, since prior appropriation art cases suggest that, if the result is sufficiently transformative, the use may be protected by fair use or may not represent an infringement of the copyright law (Ligon, 2019). However, accidental reproduction of protected works in part might happen, requiring the explicit right-holders' authorization, and not only their not reservation (Sturm et al., 2019). For this reason, in addition to use (new) transformative methods, we suggest to conduct experiments about accidental plagiarism that may be caused by the developed system, as done, for instance, by Hadjeres et al. (2017).

These considerations about the transformative nature of the result seem fundamental to establish a potential copyright infringement in case of using a protected work as the input image to algorithms like Google's DeepDream. If the result of the modifications of the input is an image that substantially resembles it, it is likely to be considered a reproduction in part, and so this might lead to a copyright infringement. However, as reported by Guadamuz (2017), this process typically results in producing new images that do not resemble the original ones. This opens the possibility of considering them sufficiently transformative not to be considered as a reproduction in part. In addition, the fact that they are not the result of creative decisions by the programmers leads to the question that we will try to address in the next section.

4. **Copyright of Generated Works**

The remaining question is who, if anyone, would be the owner of the IP rights associated with an artwork produced by a generative model. To answer this question, this section is divided in two parts: a subsection with the current legal analysis ([Section 4.1](#)) and a subsection with possible future addresses and some policy suggestions ([Section 4.2](#)).

4.1. *Legal analysis*

At first, it is important to make some distinctions. If the generative model is used just as a tool (or the human has an active role in the creative process), the human will be considered as the author; this means that if the human is in charge of the intellectual creation (e.g., by setting all the parameters required to

characterize the product, as it is possible with StyleGAN [Karras et al., 2019]) or the product can be considered as a co-creation, then the authorship is assigned to that person. In addition, even if the machine has generated the work independently from the human but the latter has selected and evaluated the outputs, rejecting some works and choosing only the best ones following his/her aesthetic tastes, the human can arguably be considered as the author of the work (Glasser, 2001). As far as works that are fully attributable to a machine are concerned, actually no one would obtain the copyright of machine-generated artworks (Santos and Machado, 2020). In fact, a fundamental requirement for the application of all the actual laws is that of originality. Even if it is not straightforward to find a precise and applicable definition of originality, in the European Union, it has been commonly considered as satisfied when the work is the reflection of author's personality (Deltorn, 2017), whereas in the United States, it could be interpreted as a minimum requiring evidence of a human (intellectual) creativity (Gervais, 2002). It is questionable to say that computer-generated artworks are the result of the personality of someone—or something—leaving the works unprotected.⁷ As a confirmation of this, the Compendium of U.S. Copyright Office Practices establishes that it will not register works produced by a machine or mere mechanical process that operates randomly or automatically without any creative input or intervention from a human author (see Article 313.2), citing, as example, a list of mechanical activities that are the exact opposite with respect to those performed by GDL, and the ones that might be reasonably considered as creative (Palace, 2019). In addition, Spain, Germany, and Australia have formulated a similar criterion, establishing that only works created by humans can be protected by copyright.

On the contrary, the most famous example of a law article for machine-generated artworks is Section 9(3) of the British Copyright, Designs and Patents Act. This section states that in case of a literary, dramatic, musical, or artistic work, which is computer-generated, the author shall be taken to be the person by whom the arrangements necessary for the creation of the work are undertaken (the same criterion is also considered by Ireland, New Zealand, Hong Kong, South Africa, and India). This section has been the subject of an intense debate (see Bond and Blair, 2019). The general consensus is that is difficult, even if not impossible, to find a person who provides necessary arrangements for contemporary machine-generated artworks, i.e., the ones created by GDL (Smith, 2017). The current lack of protection and the nongeneral applicability of British criterion are two of the reasons why the European Parliament and Commission have recently been highlighting the need of a specific law for IP rights in case of machine-generated works.

In particular, in 2020, the European Parliament adopted a Resolution “on IP rights for the development of AI technologies”, which follows a 2017 Resolution with recommendations to the Commission on Civil Law Rules on Robotics, the first step in trying to regulate this chaotic field. In 2017, the European Parliament asked a solution to protect and, at the same time, foster innovation to the Commission, in order to overcome the problem of nonallocating rights explained above, supporting a horizontal and technologically neutral approach to IP. This approach appears to have an influence on the content of the most recent Resolution. Here, the Parliament highlights the importance of creating a regulation to protect IPRs in the field of AI, in order to protect innovation, guarantee the legal certainty, and build the trust needed to encourage investments in these technologies (Points 3 and 6). Then, it suggests not to impart legal personality to AI—so, no rights can be assigned to it (Point 13). On the contrary, it recommends, if copyright is considered as the correct protection for AI-generated works, to assign the ownership of rights to the person who prepares and publishes a work lawfully, provided that the technology designer has not expressly reserved the right to use the work in that way (from the Explanatory Statement of Report A9-0176/2020, on which the Resolution is based; Point 15 specifies that ownership of rights should only be assigned to natural or legal persons that created the work lawfully, but in our opinion, the term “created” in the context of AI-generated works is misleading and unclear, and the explanation reported above is more useful). In addition, it clarifies that this position is legally correct and harmonized with the existing law;

⁷ This is one of the reasons why the World Intellectual Property Organization has recently started a Conversation on Intellectual Property and Artificial Intelligence to discuss the impact of AI on IP. See, for example, https://www.wipo.int/about-ip/en/artificial_intelligence/conversation.html.

the only requirement is to consider the condition of originality as satisfied not only if the process is creative, but even when the result is creative. This derives from the assumption that AI-generated creation and *traditional* creation still have the aim of expanding cultural heritage in common, even if the creation takes place by means of a different act. For this reason, if and when this Resolution is embraced by the European Commission too, the owner of the rights will be the person that has prepared and published the work lawfully.

4.2. Policy suggestions

Even if the current laws do not contemplate machine-generated works for copyright protection, the matter of right attribution has been hugely discussed, not only in terms of law, but also in terms of ethical implications.

The position of not assigning copyright in machine-generated works may appear to be convenient at first; indeed, it does not ask to change anything. It might also help preserve the centrality of human authorship in copyright law (see Mezei, 2020) and stress the importance of what an author should be versus what an author should do (see Craig and Kerr, 2019). Another, more practical reason is that a work should receive copyright protection only if an *author* exists; but to be considered so, the work must include a meaning or a message he/she wishes to convey, and this cannot happen if no one is able to predict the output of the program (Boyden, 2016), as in deep learning models (Ginsburg, 2018). Finally, placing computer-generated works in the public domain can help preserve the centrality of humans in creative fields, since protection would be guaranteed only to work with an intellectual human contribution (Palace, 2019).

At the same time, there are also strong reasons not to leave them unprotected. First of all, though consistent with the traditional concept of an author as a person, denying protection is inconsistent with the historically flexible interpretation and application of copyright laws as technology has developed. AI products should be evaluated following this flexible interpretation too (Butler, 1982). However, the best motivation for the allocation of ownership interests to someone is that the person should be incentivized not for the ideation and creation of the work in itself, but for its public promotion and for making it possible for the computer to create the work (by writing it, training it, or instructing it; see Miller, 1993). If the law considers a machine-generated work as incapable of being owned because of the lack of a human author, there is limited incentive for creating them and making them public. On the contrary, this might lead to potentially malicious behavior, for example, the person that used the algorithm to generate the work might be tempted to lie about the way it was created or change it in order to be considered its author (Samuelson, 1986). Finally, the idea that this would mean to incentivize the proliferation of arts and articles of poor quality, penalizing the role of human artists and journalists (Gervais, 2020), does not convince us at all. If, as stated, the protection of computer-generated works would translate into a larger number of mediocre works, then for humans it will be easier to produce works of quality higher than those that have been generated by machines and clearly stand out. In addition, even if the current copyright laws were thought to regulate scarcity of products created by humans (Hurt and Schuchman, 1966) and not abundance (of machine-generated products), leaving all of them in the public domain could cause more damage to human authors. The possibility of using them for free may persuade clients to do so, even if human artworks are qualitatively better.⁸

As we have seen in the previous subsection, the European Parliament seems to agree with this line of thought, proposing to allocate rights to who has prepared and published the work lawfully. The same conclusion (of the European Parliament) can be reached in different ways. Following Franceschelli (2019), there can be three main individuals involved in the process: the programmer, the person who provides necessary arrangements, and the user. Notice that even if we will consider them separately, in many cases, they are the same person. The programmer is just the person who has

⁸Note also that the art industry feels art lovers would always prefer handmade arts and crafts: <https://www.forbes.com/sites/anniebrown/2021/09/06/is-artificial-intelligence-set-to-take-over-the-art-industry/?sh=78b774c33c50>.

written the code for the machine; the person who provides necessary arrangements can be, for example, the individual who provides instructions for the desired output, or necessary information about the work the machine has to generate; and the user is the person that, legitimately (because, e.g., the individual is the owner of the machine or has acquired it with a license), ultimately runs the machine and asks it to generate an artwork. In our opinion (but also of the European Parliament and of Bohlen (2017); Samuelson (1986); Yanisky-Ravid (2017), among others), the rights should be allocated to the user, that can be considered an alter ego of the “person who prepares and publishes a work lawfully” (even if the person who provides the necessary arrangements can also be seen under the definition of who prepares the work). Although, in a way, it seems counterintuitive, there are different ways to reach this conclusion.

Yanisky-Ravid (2017) and Bohlen (2017) suggest an analogy with the ownership of economic rights in case of software produced by an employee, the so-called work-made-for-hire doctrine: as the employer is entitled to exercise all economic rights in an employee’s computer program (if the creation is part of the scope of his/her employment or it is commissioned by the employer), the user is the person who actually causes the creation of the work. Therefore, it is possible to say that the user has *employed* the computer for his/her creative endeavors. In this way, rights can be allocated preventing works from falling into the public domain regardless of the extent of human creative contribution. In addition, Hristov (2017) and Gurkaynak et al. (2017) suggest to use the analogy with work-made-for-hire doctrine to allocate rights, but with a potentially different conclusion. In fact, they suggest that the equivalent of the employer has to be the programmer or the owner of the generative program, because they are those who really need to be incentivized, rejecting the chance of assigning rights to the end user. However, their definition of end users is quite different from that considered here: here, we consider as users the persons who have lawful access to the generative program, and can lawfully use it to generate works. We assume that they are the owners of the program (who, following Wu (1997), should be the copyright owners), possibly because (a) they are also its developers; (b) they have acquired it from the developers; or (c) following the same doctrine, the developers are their employees. In any case, they are the lawful owners of its economic rights. Alternatively, we assume that they have acquired it via license. In this way, there is no need of an extra economic incentive for the developers (they have already been paid or have chosen to release the product freely) or for the owners (they have already been paid for the license or have chosen to release the product freely). On the contrary, the users have paid for using the generative program, or they can use it freely because of a particular license. In general, we believe that, even if the generative program has been published and it is accessible by millions of end users, Terms of Service written by the publisher should regulate these kinds of problems. We suggest future publishers combine their generators with Terms of Service in order to clearly identify the owner of the generated products (and the associated terms).

Vice versa, Samuelson (1986) (but also see Denicola, 2016) states that even if users may not be the ultimate market-makers for machine-generated works, they are in the same position as traditional authors since they take the initial steps that will bring a work into the marketplace (and into its exterior form). Since society has an interest in making these works available to the public, the most effective solution is to give incentives to users to make them available and accessible to others.

Finally, we can reach the same conclusion also working by elimination. The programmer is responsible of the machine creative abilities and, for other kinds of AI (e.g., rule-based systems), it might seem to be enough to establish the originality requirement—and therefore the ownership—in the programmer (Farr, 1989). However, as regards GDL, he/she just creates the potentiality for the creation of the output, but not its actuality (Samuelson, 1986). It would be like trying to assign copyright to the teaching artist of a painter, instead of to the painter himself/herself, or, using the analogy proposed by Ralston (2004), to claiming a knife manufacturer is more responsible for a murder than the person who wielded the knife. The person who provides necessary arrangements can be difficult to identify, and sometimes the generative model may not have such a person associated to it, due to the complexity of deciding which are necessary arrangements and which are just useful arrangements (Franceschelli, 2019). A simple example to understand why it is so difficult to use this definition to assign ownership is now presented.

Consider the already mentioned Botnik and its creative keyboard.⁹ When we open it, it starts with *John Keats* as the source, and it starts suggesting words according to John Keats' texts on which a neural model was previously trained. Then, we can start selecting each time the word in the first position between suggestions, composing a new (and hopefully creative) text. Notice that we could choose the word among different options, and this selection would mean that we would be recognized as the authors (using Botnik just as a tool), be we did not. Now, which shall be considered the necessary arrangements? The only two things we have done as users were to open the website and compulsively click to the first suggested word. Is it enough to consider our actions as necessary arrangements? Of course not. So, necessary arrangements shall be the ones performed by who has loaded John Keats' poems and trained the network; or maybe the ones performed by who has decided that the preset source was John Keats one. But no poems would arise from the creative keyboard without our incredibly simple operations, and it does not seem reasonable to leave, a priori, the ownership of these kinds of machine-generated works to someone who was not involved in the materialization of the work—that is, what the law shall protect. No solution seems reasonable, in these cases, in order to assign the rights to who has provided necessary arrangements; of course, there can be other cases in which it might seem the right choice, but it would be better to have a rule with the most general applicability, even if Grimmelmann (2016b) suggests a case-by-case analysis to deal with the heterogeneity of computer-generated works. For this reason, we discard the person who provides necessary arrangements. On the contrary, allocating rights to the user seems not to have any particular flaws; of course, he/she may not have provided any creative contribution, but, as explained above, this does not seem a valid argumentation. For all of these reasons, we think the user should be considered as the owner of (economic) rights.

Finally, an additional consideration about copyright allocation must be done. One of the most explored creative fields by AI researchers is (video)game design. With respect to GDL, it typically concerns the use of generated images (Tilson and Gelowitz, 2019), characters (Jin et al., 2017), or soundtracks¹⁰ inside games; in these cases, all the conclusions drawn until now are still valid, and no additional considerations are required. However, a growing application is the so-called procedural content generation, where the game scenarios are dynamically generated during game (Liu et al., 2021). Although this task is technically very similar to image generation (with an additional complexity provided by the need of dynamic adaptation and of complexity growth), an additional consideration about copyright allocation is needed. In fact, in procedural content generation, it is not immediate to identify the user (as we have intended him/her). Of course, there is the player, who is a *game* user; but the copyright allocation as stated above is about *generative model* user. In this case, the algorithm which generates the game content is not directly used by a person; it is directly used by the game code, and therefore indirectly by the game programmer, who has employed the GDL techniques not to generate content statically, but dynamically. By considering the programmer of the game code as the user of the generative model, the conclusion drawn during this section remains generally applicable in our opinion (and also according to previous literature, e.g., Bridy, 2012).

5. Code Generation

Finally, a quite different GDL use case is about code generation. In particular, GitHub (and OpenAI) Copilot¹¹ has caused a great debate about copyright implications (Guadamuz, 2021). Copilot is an AI system able to autocomplete lines of code, but also to generate entire blocks and functions just providing comments or signatures. It is based on an autoregressive model (very close to GPT-3 [Brown et al., 2020]) trained on English texts and source codes from publicly available sources as GitHub's public repositories, but not exclusively.

⁹ <https://botnik.org/apps/writer/>.

¹⁰ <https://cordis.europa.eu/article/id/421438-ai-composers-create-music-for-video-games>.

¹¹ <https://copilot.github.com/>.

From a legal perspective, GitHub Copilot introduces some additional complexities that are of particular interest for our analysis of copyright in generative learning. With respect to the acquisition and the storage of works used during training, we note that the fact a work is publicly available does not mean it is in the public domain, or it is released under a permissive license—see Points 1 and 2 in Section 3. For instance, the contents of GitHub’s public repositories, which are not associated a license, are intended to be under copyright’s law—see Points 3 and 4 in Section 3. The GitHub’s Terms of Service states that GitHub can process content shared in public repositories as needed to provide the Service, which includes all the applications, software, and products provided by GitHub—and therefore also includes Copilot. However, content from external sources is also used. In order to lawfully exploit these sources, their use must fall under the definition of “fair use” (or, in the European Union, must be considered a “lawful use”), as highlighted above. GitHub itself claims that training machine learning models on publicly available data is considered fair use across the machine learning community; however, it is not straightforward to confirm so. The fair use is determined based on the four criteria examined in Section 3.1. In this case, the use is not for research purpose, and it seems more expressive than non-expressive. In addition, its economic character is not negligible (Copilot is free to use, but companies may use it). The public availability of the works helps satisfy the second criterion for fair use. Then, the work is entirely used during training, but the substantiality of the use is questionable, as discussed before. Finally, the effect upon the potential market depends on the model itself. If it cannot substantially reproduce an existing source code or, if it can, it is able to identify it and refer the user to the original source, also this fourth condition is satisfied and the use could presumably be seen as a fair use.

We also note that publicly available contents have been released under licenses like GNU GPL, with the goal of protecting freedom.¹² These licenses are chosen in order to avoid any commercial use of the free software, asking to release the derivative work with the same license, in order to foster freedom (and innovation). This is in clear contradiction with the application of the fair use doctrine in case of training a model which could be used for economic purposes. In this way, the question becomes more about ethics than about law: should fair use doctrine be applicable to deep learning independently from the economic character of the application (as it is: fair use is an on/off switch (Ginsburg, 2014) and once established the use is fair, nothing prevents the user to do so) or should developers have the opportunity of reserving the right to this kind of utilization? The European Union appears to be aligned to the second option. As discussed in Section 3.2, Directive’s Article 4 considers this use as lawful, unless the rightholder has not expressly reserved this exception. It is not completely clear how to practically reserve it; in our opinion, an interesting possibility will be to augment current licenses and to deal with this right as with the others: as a license specifies if a commercial use is allowed or not, it can also specify if a training use (not only for research purposes) is allowed or not. Of course, this sort of solution could work in the United States only if they will decide that this use can be reserved, and that the fair use doctrine is not always applicable.

Another important consideration is related to potential copyright infringement. As observed in Section 3.3, training on protected works could cause a reproduction in parts of the model or, more probably, in the generated output. It has been shown that Copilot quotes existing content very rarely, and it mostly quotes very common code, for which it is almost impossible to detect a substantial reproduction (Ziegler, 2021). However, since in principle it may happen, it is praiseworthy that GitHub and OpenAI are building a tracker to detect the rare instances of code that is repeated from the training set. We hope it will help users identify the source and author of the code, so that they can directly use it and check the corresponding license, which might require that certain conditions are satisfied.

Finally, the issue related to the ownership of generated code does not introduce any additional complexity. Currently, Copilot works only as a tool, asking the user to test, review, and complete the code; therefore, the user is lawfully considered as the author—and as the owner. If, in the future, it will not require active supervision anymore and the program will be considered as the author, then it will just become another example of how the work-made-for-hire doctrine perfectly fits GDL. The user, who also

¹² www.gnu.org/philosophy/open-source-misses-the-point.html.en.

provides necessary arrangements (in the form of comments and signatures), is employing Copilot to perform a task within the scope of its employment; therefore, even with a higher degree of autonomy, the user should be considered as the owner of related rights.

6. Conclusion

In this article, we have explored the most important problems concerning copyright in relation to GDL, trying to understand how the current laws can or cannot permit some common practices. In particular, in [Section 3](#), we have analyzed if and when, under US and EU law, it is possible to store protected works with the goal of training a generative model; the conclusion is that fair use doctrine and 2019 EU's Directive seem to allow this, but with some reservations about the general applicability (for the United States) and a few practical obstacles (for the European Union). Then, in [Section 4](#), after having explained that at the moment a GDL output is not protected by copyright, we have also explored future directions in terms of design of legislative frameworks and how these works shall be protected in near future.

In conclusion, from a practical point of view, as far as researchers are concerned, we suggest (a) to pay attention at the information that is stored (and for how long) during the whole process (and, in case of private use, to carefully check terms or other conditions which could prevent the TDM exception); (b) to try to diverge from the dataset used during training, in order to avoid the risk of a reproduction in part, but also in order to strengthen a transformative use claim; (c) to clarify their position through Terms of Service if other users would be able to use their generative model; and (d) of course, to keep update about the evolution of the legislative frameworks at national and international levels.

Acknowledgment. A preprint version of the article is available in arXiv (Franceschelli and Musolesi, 2021a).

Funding Statement. This work received no specific grant from any funding agency, commercial, or not-for-profit sectors.

Competing Interests. The authors declare no competing interests exist.

Author Contributions. G.F. and M.M. designed the study, wrote the first draft, and approved the final version of the manuscript.

Data Availability Statement. No data, code, or other materials have been used.

Ethical Standards. The research meets all ethical guidelines, including adherence to the legal requirements of the study country.

References

- Asay CD, Sloan A and Sobczak D (2020) Is transformative use eating the world? *Boston College Law Review* 61(3), 905–970.
- Bandy J and Vincent N (2021). Addressing “documentation debt” in machine learning research: A retrospective datasheet for BookCorpus. *NeurIPS Track on Datasets and Benchmarks*. Online: Curran Associates, Inc., Vol. 1.
- Beaulieu-Jones B, Wu Z, Williams C, Lee R, Bhavnani S, Byrd J and Greene C (2019) Privacy-preserving generative deep neural networks support clinical data sharing. *Circulation: Cardiovascular Quality and Outcomes* 12, e005122.
- Bengio Y, Courville A and Vincent P (2013) Representation learning: A review and new perspectives. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 35, 1798–1828.
- Bohlen HM (2017) *EU Copyright Protection of Works Created by Artificial Intelligence Systems*. Master's thesis. Bergen, Norway: University of Bergen.
- Bolukbasi T, Chang K-W, Zou J, Saligrama V and Kalai A (2016) Man is to computer programmer as woman is to homemaker? Debiasing word embeddings. In *NeurIPS*. Red Hook, NY: Curran Associates, Inc., pp. 4356–4364.
- Bonadio E and McDonagh L (2020) Artificial intelligence as producer and consumer of copyright works: Evaluating the consequences of algorithmic creativity. *Intellectual Property Quarterly* 2, 112–137.
- Bond T and Blair S (2019) Artificial intelligence & copyright: Section 9(3) or authorship without an author. *Journal of Intellectual Property Law & Practice* 14(6), 423–423.
- Boyden BE (2016) Emergent works. *The Columbia Journal of Law & the Arts* 3(39), 377–394.
- Bridy A (2012) Coding creativity: Copyright and the artificially intelligent author. *Stanford Technology Law Review* 5, 1–28.
- Brock A, Donahue J and Simonyan K (2019). Large scale GAN training for high fidelity natural image synthesis. In *ICLR*. New Orleans, LA: OpenReview.

- Brown T, Mann B, Ryder N, Subbiah M, Kaplan JD, Dhariwal P, Neelakantan A, Shyam P, Sastry G, Askell A, Agarwal S, Herbert-Voss A, Krueger G, Henighan T, Child R, Ramesh A, Ziegler D, Wu J, Winter C, Hesse C, Chen M, Sigler E, Litwin M, Gray S, Chess B, Clark J, Berner C, McCandlish S, Radford A, Sutskever I and Amodei D** (2020) Language models are few-shot learners. In *NeurIPS*. Online: Curran Associates, Inc. Vol. 33, pp. 1877–1901.
- Butler TL** (1982) Can a computer be an author—Copyright aspects of artificial intelligence. *Hastings Communications and Entertainment Law Journal* 4(4), 707.
- Carroll MW** (2007) Fixing fair use. *North Carolina Law Review* 85, 68.
- Chiou T** (2019) Copyright lessons on machine learning: What impact on algorithmic art. *Journal of Intellectual Property, Information Technology and Electronic Commerce Law* 10, 398–412.
- Chouldechova A and Roth A** (2020) A snapshot of the frontiers of fairness in machine learning. *Communications of the ACM* 63(5), 82–89.
- Colton S** (2008) Creativity versus the perception of creativity in computational systems. In *AAAI 2008 Spring Symposium*. Stanford, CA: AAAI Press.
- Craig CJ and Kerr IR** (2019) The death of the AI author. In *Osgoode Legal Studies Research Paper*.
- Deltoro J-M** (2017) Deep creations: Intellectual property and the automata. *Frontiers in Digital Humanities* 4(3).
- Deltoro J-M and Macrez F** (2018) *Authorship in the Age of Machine Learning and Artificial Intelligence*. Centre for International Intellectual Property Studies (CEIPI) Research Paper No. 2018-10.
- Denicola RC** (2016) Ex machina: Copyright protection for computer-generated works. *Rutgers Law Review* 69, 251–287.
- Devlin J, Chang M-W, Lee K and Toutanova K** (2019) BERT: Pre-training of deep bidirectional transformers for language understanding. In *NAACL 2019*. Minneapolis, MN: Association for Computational Linguistics, pp. 4171–4186.
- Dong H-W, Hsiao W-Y, Yang L-C and Yang Y-H** (2018) MuseGAN: Multi-track sequential generative adversarial networks for symbolic music generation and accompaniment. In *AAAI*. New Orleans, LA: AAAI Press, Vol. 32(1), pp. 34–41.
- Dwork C and Roth A** (2014) The algorithmic foundations of differential privacy. *Foundations and Trends in Theoretical Computer Science* 9(3–4), 211–407.
- Elgammal A, Liu B, Elhoseiny M and Mazzone M** (2017) CAN: Creative adversarial networks, generating “art” by learning about styles and deviating from style norms. In *ICCC*. Atlanta, GA: Association for Computational Creativity.
- Engel J, Agrawal K K, Chen S, Gulrajani I, Donahue C and Roberts A** (2019). GANSynth: Adversarial neural audio synthesis. In *ICLR: OpenReview*. New Orleans, LA.
- Esser P, Rombach R and Ommer B** (2021). Taming Transformers for High-Resolution Image Synthesis. *CVPR*. Online: IEEE, pp. 12868–12878.
- Farr EH** (1989) Copyrightability of computer-created works. *Rutgers Computer and Technology Law Journal* 15(1), 63–80.
- Foster D** (2019) *Generative Deep Learning*. Newton, MA: O’Reilly.
- Franceschelli G** (2019) *I, Artist. Opere d’arte e Intelligenza Artificiale. Il Curioso Caso del Diritto d’autore*. Scapezano, Italy: Ventura Edizioni.
- Franceschelli G and Musolesi M** (2021a) Copyright in generative deep learning. *Preprint*, arXiv:2105.09266 [cs.CY].
- Franceschelli G and Musolesi M** (2021b) Creativity and machine learning: A survey. arXiv:2104.02726 [cs.LG].
- Geburu T, Morgenstern J, Vecchione B, Vaughan JW, Wallach H, Daumé III H and Crawford K** (2021) Datasheets for datasets. *Communications of the ACM* 64(12), 86–92.
- Geiger C, Bulayenko O, Hassler T, Izumenko E, Schonherr F and Seuba X** (2015) The resolution of the European Parliament of 9 July 2015: Paving the way (finally) for a copyright reform in the European Union? *European Intellectual Property Review* 37(11), 683.
- Geiger C, Frosio GF and Bulayenko O** (2018a) Text and data mining in the proposed copyright reform: Making the EU ready for an age of big data? *International Review of Intellectual Property and Competition Law* 49, 814–844.
- Geiger C, Frosio GF and Bulayenko O** (2018b) The exception for text and data mining (TDM) in the proposed directive on copyright in the digital single market—Legal aspects. [Requested by the JURI Committee] (European Parliament, February 2018).
- Gervais DJ** (2002). Feist goes global: A comparative analysis of the notion of originality in copyright law. *Journal of the Copyright Society of the U.S.A.* 49, 949–981.
- Gervais DJ** (2020) The machine as author. *Iowa Law Review* 105, 1053–2106.
- Ginsburg JC** (2014) Fair use for free, or permitted-but-paid? *Berkeley Technology Law Journal* 29, 1383–1446.
- Ginsburg JC** (2018) People not machines: Authorship and what it means in the Berne convention. *International Review of Intellectual Property and Competition Law* 49, 131–135.
- Glasser D** (2001) Copyrights in computer-generated works: Whom, if anyone, do we reward? *Duke Law & Technology Review* 24.
- Goodfellow I, Bengio Y and Courville A** (2016) *Deep Learning*. Cambridge, MA: MIT Press.
- Goodfellow I, Pouget-Abadie J, Mirza M, Xu B, Warde-Farley D, Ozair S, Courville A and Bengio Y** (2014) Generative adversarial nets. In *NeurIPS*. Red Hook, NY: Curran Associates, Inc., pp. 2672–2680.
- Gregor K, Danihelka I, Graves A, Rezende D and Wierstra D** (2015) DRAW: A recurrent neural network for image generation. In *ICML*. Lille, France: Online: Curran Associates, Inc. Vol. 37, pp. 1462–1471.
- Grimmelmann J** (2016a) Copyright for literate robots. *Iowa Law Review* 101, 657–681.
- Grimmelmann J** (2016b) There’s no such thing as a computer-authored work—And It’s a good thing, too. *Columbia Journal of Law & the Arts* 39, 403–416.

- Guadamuz A** (2017). Do androids dream of electric copyright? Comparative analysis of originality in artificial intelligence generated works. *Intellectual Property Quarterly* 2, 169–186.
- Guadamuz A** (2021) Is GitHub's Copilot potentially infringing copyright? Available at <https://www.technollama.co.uk/is-githubs-copilot-potentially-infringing-copyright>. Retrieved May 9, 2022.
- Gurkaynak G, Yilmaz I, Doygun T and Ince E** (2017) Questions of intellectual property in the artificial intelligence realm. *The Robotics Law Journal* 3(2), 9–11.
- Hadjeres G, Pachet F and Nielsen F** (2017) DeepBach: A steerable model for Bach chorales generation. In *ICML*. Sydney, NSW, Australia. JMLR: Online: Curran Associates, Inc. Vol. 70, pp. 1362–1371.
- Hilty R and Richter H** (2017) *Position Statement of the Max Planck Institute for Innovation and Competition on the Proposed Modernisation of European Copyright Rules Part B Exceptions and Limitations (Art. 3—Text and Data Mining)*. Max Planck Institute for Innovation & Competition Research Paper No. 17-02.
- Hristov K** (2017) Artificial intelligence and the copyright dilemma. *IDEA: The IP Law Review* 57(3), 24.
- Huang C-ZA, Vaswani A, Uszkoreit J, Simon I, Hawthorne C, Shazeer N, Dai AM, Hoffman MD, Dinculescu M and Eck D** (2019) Music transformer. In *ICLR*. New Orleans, LA: Online: Curran Associates, Inc.
- Hurt RM and Schuchman RM** (1966) The economic rationale of copyright. *The American Economic Review* 56(1/2), 421–432.
- Isola P, Zhu J-Y, Zhou T and Efros AA** (2017) Image-to-image translation with conditional adversarial networks. In *CVPR*. Honolulu, HI. IEEE, pp. 5967–5976.
- Jaques N, Gu S, Turner RE and Eck D** (2016) Generating music by fine-tuning recurrent neural networks with reinforcement learning. In *NeurIPS*. Barcelona, Spain: Curran Associates, Inc.
- Jin Y, Zhang J, Li M, Tian Y and Zhu H** (2017) Towards the high-quality anime characters generation with generative adversarial networks. In *NeurIPS Workshop*. Long Beach, CA: Curran Associates, Inc.
- Karpathy A** (2015) The unreasonable effectiveness of recurrent neural networks. Available at <http://karpathy.github.io/2015/05/21/rnn-effectiveness>. Retrieved May 9, 2022.
- Karras T, Aila T, Laine S and Lehtinen J** (2018) Progressive growing of GANs for improved quality, stability, and variation. In *ICLR* Vancouver, Canada.
- Karras T, Laine S and Aila T** (2019) A style-based generator architecture for generative adversarial networks. In *CVPR*. Long Beach, CA: IEEE/CVF, pp. 4396–4405.
- Kingma DP and Welling M** (2014) Auto-encoding variational bayes. In *ICLR*. Banff, Canada.
- Kretschmer M and Margoni T** (2018) Data mining: Why the EU's proposed copyright measures get it wrong. Available at <https://www.iusinitinere.it/of-data-rights-and-boundaries-text-and-data-mining-under-eu-copyright-law-23891>. Retrieved May 9, 2022.
- Lau JH, Cohn T, Baldwin T, Brooke J and Hammond A** (2018) Deep-speare: A joint neural model of poetic language, meter and rhyme. In *ACL* Melbourne, Australia: Association for Computational Linguistics, pp. 1948–1958.
- Lessig L** (2004) *Free Culture: How Big Media Uses Technology and the Law to Lock Down Culture and Control Creativity*. New York: Penguin Press.
- Ligon SE** (2019) AI can create art, but can it own copyright in It, or infringe? Available at lexisnexis.com/practicalguidance/the-journal/b/pa/posts/ai-can-create-art-but-can-it-own-copyright-in-it-or-infringe. Retrieved June 18, 2021.
- Liu J, Snodgrass S, Khalifa A, Risi S, Yannakakis GN and Togelius J** (2021) Deep learning for procedural content generation. *Neural Computing and Applications* 33, 19–37.
- Liu Y, Peng J, Yu JJ and Wu Y** (2019) PPGAN: Privacy-preserving generative adversarial network. In *ICPADS*. Tianjin, China: IEEE.
- Lunney GSJ** (2002) Fair use and market failure: Sony revisited. *Boston University Law Review* 82, 975–1030.
- Margoni T** (2018) Artificial intelligence, machine learning and EU copyright law: Who owns AI? CREATE Working Paper.
- Mezei P** (2020) From Leonardo to the next Rembrandt—The need for AI-pessimism in the age of algorithms. *UFITA* 2, 390–429.
- Miller AR** (1993) Copyright protection for computer programs, databases, and computer-generated works: Is anything new since CONTU? *Harvard Law Review* 106(5), 977–1073.
- Mordvintsev A, Olah C and Tyka M** (2015) Inceptionism: Going deeper into neural networks. Available at ai.googleblog.com/2015/06/inceptionism-going-deeper-into-neural.html. Retrieved June 18, 2021.
- Netanel NW** (2008) *Copyright's Paradox*. Oxford, NY: Oxford University Press.
- Netanel NW** (2011) Making sense of fair use. *Lewis & Clark Law Review* 15, 715–771.
- Nimmer D** (2003) “Fairest of them all” and other fairy Tales of fair use. *Law and Contemporary Problems* 66, 263–288.
- O’Leary D** (2019) How algorithms can seem racist. Available at <http://mindmatters.ai/2019/11/how-algorithms-can-seem-racist/>. Retrieved June 18, 2021.
- Padilla-López J, Chaaoui A and Flórez-Revuelta F** (2015) Visual privacy protection methods: A survey. *Expert Systems with Applications* 42, 4177–4195.
- Palace VM** (2019) What if artificial intelligence wrote this? Artificial intelligence and copyright law. *Florida Law Review* 71, 217–242.
- Parmar N, Vaswani A, Uszkoreit J, Kaiser L, Shazeer N, Ku A and Tran D** (2018) Image transformer. In *ICML*. Stockholm, Sweden: Curran Associates, Inc., Vol. 80, pp. 4055–4064.
- Payne C** (2019) MuseNet. Available at <http://openai.com/blog/musenet>. Retrieved June 18, 2021.

- Potash P, Romanov A and Rumshisky A** (2015) GhostWriter: Using an LSTM for automatic rap lyric generation. In *EMNLP*. Stroudsburg, PA: Association for Computational Linguistics, pp. 1919–1924.
- Radford A, Kim JW, Hallacy C, Ramesh A, Goh G, Agarwal S, Sastry G, Askell A, Mishkin P, Clark J, Krueger G and Sutskever I** (2021) Learning transferable visual models from natural language supervision. In *ICML*. Online: Curran Associates, Inc.
- Radford A, Wu J, Child R, Luan D, Amodei D and Sutskever I** (2019) Language models are unsupervised multitask learners. *OpenAI Blog* 1(8), 9.
- Ralston WT** (2004) Copyright in computer-composed music: HAL meets Handel. *Journal of the Copyright Society of the U.S.A.* 52 (3), 281–308.
- Ramesh A, Pavlov M, Goh G, Gray S, Voss C, Radford A, Chen M and Sutskever I** (2021) Zero-shot text-to-image generation. In *ICML*. Online: Curran Associates, Inc., Vol. 139, pp. 8821–8831.
- Riedo M** (2019) Of data rights and boundaries: Text and data mining under EU copyright law. Available at <https://www.iusinitinere.it/of-data-rights-and-boundaries-text-and-data-mining-under-eu-copyright-law-23891>. Retrieved June 18, 2021.
- Rosati E** (2018) The Exception for Text and Data Mining (TDM) in the Proposed Directive on Copyright in the Digital Single Market—Technical Aspects. [Requested by the JURI Committee] (European Parliament, February 2018).
- Rosati E** (2019) Copyright as an obstacle or an enabler? A European perspective on text and data mining and its role in the development of AI creativity. *Asia Pacific Law Review* 27(2), 198–217.
- Sag M** (2019) The new legal landscape for text mining and machine learning. *Journal of the Copyright Society of the USA* 66, 291–363.
- Samuelson P** (1986). Allocating ownership rights in computer-generated works. In *Symposium Cosponsored by University of Pittsburgh law Review and the Software En on the Future of Software Protection*. Pittsburgh, PA: University of Pittsburgh Press, pp. 1185–1228.
- Samuelson P** (2009) Unbundling fair uses. *Fordham Law Review* 77(5), 2537–2621.
- Santos C and Machado A** (2020) Intellectual property on works of art made by artificial intelligence. *International Journal of Advanced Engineering Research and Science* 7(12), 49–59.
- Schonberger D** (2018). Deep copyright: Up - and downstream questions related to artificial intelligence (AI) and machine learning (ML). In *Droit d'auteur 4.0/Copyright 4.0*. Geneva/Zurich: Schulthess Editions Romandes, pp. 145–173.
- Smith L** (2017) AI and IP: Copyright in AI-generated works (UK law). Available at talkingtech.cliffordchance.com/en/ip/copyright/ai-and-ip-copyright-in-ai-generated-works-uk-law.html. Retrieved June 21, 2021.
- Snell C** (2021) Alien dreams: An emerging art scene. Available at <http://ml.berkeley.edu/blog/posts/clip-art/>. Retrieved September 8, 2021.
- Sobel B** (2017) Artificial intelligence's fair use crisis. *Columbia Journal of Law and the Arts* 41, 45–97.
- Sobel B** (2020) A taxonomy of training data: Disentangling the mismatched rights, remedies, and rationales for restricting machine learning. In *Artificial Intelligence and Intellectual Property*. Oxford: Oxford University Press.
- Sturm BL, Iglesias M, Ben-Tal O, Miron M and Gomez E** (2019) Artificial intelligence and music: Open questions of copyright law and engineering praxis. *The Artist and Journal of Home Culture* 8, 115.
- Sturm BL, Santos JF, Ben-Tal O and Korshunova I** (2016) Music transcription modelling and composition using deep learning. *Preprint*, arXiv:1604.08723 [cs.SD].
- Tilson A and Gelowitz CM** (2019) Towards generating image assets through deep learning for game development. In *CCECE*. Edmonton, Canada: IEEE, pp. 1–4.
- Triaille J-P, d'Argenteuil JdM and de Francquen A** (2014) Study on the legal framework of text and data mining (TDM). Available at <https://op.europa.eu/en/publication-detail/-/publication/074ddf78-01e9-4a1d-9895-65290705e2a5/language-en/>. Retrieved September 8, 2021.
- Van Den Oord A, Dieleman S, Zen H, Simonyan K, Vinyals O, Graves A, Kalchbrenner N, Senior A and Kavukcuoglu K** (2016a) WaveNet: A generative model for raw audio. In *ISCA. Speech Synthesis Workshop*. Seoul, Korea: ISCA, p. 125.
- Van Den Oord A, Kalchbrenner N and Kavukcuoglu K** (2016b) Pixel recurrent neural networks. In *ICML*. New York, NY: JMLR., pp. 1747–1756.
- Van Den Oord A, Kalchbrenner N, Vinyals O, Espeholt L, Graves A and Kavukcuoglu, K** (2016c) Conditional image generation with PixelCNN decoders. In *NeurIPS*. Red Hook, NY: Curran Associates Inc., pp. 4797–4805.
- Vernier G, Caselles-Dupré H and Fautrel P** (2020) Electric dreams of Ukiyo: A series of Japanese artworks created by an artificial intelligence. *Patterns* 1(2), 100026.
- Wu AJ** (1997) From video games to artificial intelligence: Assigning copyright ownership to works generated by increasingly sophisticated computer programs. *AIPLA Quarterly Journal* 25(1), 131–180.
- Wu B, Zhao S, Chen C, Xu H, Wang L, Zhang X, Sun G and Zhou J** (2019) Generalization in generative adversarial networks: A novel perspective from privacy protection. In *NeurIPS*. Vancouver, Canada: Curran Associates, Inc., Vol. 32, pp. 307–317.
- Xu C, Ren J, Zhang D, Zhang Y, Qin Z and Ren K** (2019) GANobfuscator: Mitigating information leakage under GAN via differential privacy. *IEEE Transactions on Information Forensics and Security* 14(9), 2358–2371.
- Xu D, Yuan S, Zhang L and Wu X** (2018) FairGAN: Fairness-aware generative adversarial networks. *Big Data*. Seattle, WA: IEEE, 570–575.
- Yanisky-Ravid S** (2017) Generating Rembrandt: Artificial intelligence, copyright, and accountability in the 3A Era—The human-like authors are already here—A new model. *Michigan State Law Review*. <http://doi.org/10.2139/ssrn.2957722>

- Yi X, Sun M, Li R and Li W** (2018) Automatic poetry generation with mutual reinforcement learning. In *EMNLP*. Stroudsburg, PA: Association for Computational Linguistics, pp. 3143–3153.
- Yu L, Zhang W, Wang J and Yu Y** (2017) SeqGAN: Sequence generative adversarial nets with policy gradient. In *AAAI*. San Francisco, CA: AAAI Press, pp. 2852–2858.
- Yu N, Li K, Zhou P, Malik J, Davis L and Fritz M** (2020) Inclusive GAN: Improving data and minority coverage in generative models. In *ECCV*, Vol. 12367. Cham: Springer.
- Zhang H, Goodfellow I, Metaxas D and Odena A** (2019) Self-attention generative adversarial networks. In *PMLR*, Vol. 97, pp. 7354–7363.
- Zhang X and Lapata M** (2014) Chinese poetry generation with recurrent neural networks. In *EMNLP*. Stroudsburg, PA: Association for Computational Linguistics, pp. 670–680.
- Ziegler A** (2021) GitHub Copilot: Parrot or crow? Available at <https://docs.github.com/en/github/copilot/research-recitation>. Retrieved July 3, 2021.
- Zugarini A, Melacci S and Maggini M** (2019) Neural poetry: Learning to generate poems using syllables. In *ICANN*. Munich, Germany: Springer, pp. 313–325.