

**Comments of the Center for Democracy & Technology
On the World Intellectual Property Organization’s
Conversation on Intellectual Property Protection and Artificial Intelligence**

February 14, 2020

The Center for Democracy & Technology (CDT) thanks the World Intellectual Property Organization (WIPO) for the opportunity to take part in this conversation. CDT is a non-profit advocacy organization working to preserve individual rights and democratic ideals in existing and new applications of technology. The emergence of machine learning as a path toward artificial intelligence has, and will continue to raise difficult legal and moral questions. Many of these questions may challenge our established systems or force us to reassess their purposes and goals in light of emerging technological capabilities. Yet, not every new technology demands new law or policy.

While using the phrase “artificial intelligence” evokes images of super-intelligent robots with human-like capabilities, those concepts have not yet become reality. Instead, much of what we call “AI” now amounts to advanced statistical analysis, and predictions and prognoses based on models derived therefrom. CDT encourages WIPO to maintain a critical approach when assessing the relationship between copyright and machine learning and to base any legal analyses or policy positions on the factual technological underpinnings of how works are used to develop artificial intelligence systems. Further, we encourage WIPO to critically assess the fundamental purposes of copyright systems and the applicability of those purposes in the context of non-human entities such as may emerge in the pursuit of AI.

To that end, CDT suggests utilizing the expertise of the GRAIL Network (Governance Research in Artificial Intelligence Leadership Network), a group of computer scientists, researchers, economists, and legal experts working on a broad set of issues related to the development and use of AI.¹ The network intends to bring policymakers and technical experts together to better align policy development with technical reality and to help technical experts play a more active role in policymaking. We encourage WIPO to consider GRAIL as an additional resource for understanding of the technical, legal, and societal implications of AI.

We address a subset of WIPO’s questions relating to copyright and artificial intelligence below and offer additional inquiries WIPO may wish to include in the next phase of this conversation.

Issue 6: Authorship and Ownership

With respect to authorship of creative works and the ownership of rights therein, CDT opposes any modifications to copyright that could further separate the act of creation from the benefits of copyright protection. We view copyright systems as legal constructions designed around

¹ See www.grailnetwork.org or contact info@grailnetwork.org for more information.

attributes that are, for now, uniquely human. Rather than trying to rebuild these systems under the guise of autonomous machine creativity, we suggest that preserving copyright for the works of entities with human-like goals and incentives better serves the purposes of copyright.

We begin by addressing WIPO's issue prompt and related questions. Although we agree with the introductory statement ("*The policy positions adopted in relation to the attribution of copyright to AI-generated works will go to the heart of the social purpose for which the copyright system exists.*"), we note that both the prompt and questions seem to rely on a suite of unspoken assumptions. CDT suggests that the quality of the ensuing conversation about AI and Copyright would be improved by an explicit acknowledgement of the assumptions implied in the consultation document and a discussion about which values and interests currently recognized by copyright systems should be considered in the context of artificial intelligence.

For example, the phrase "*If AI-generated works were excluded from eligibility for copyright protection, the copyright system would be seen as an instrument for encouraging and favoring the dignity of human creativity over machine creativity,*" indicates that dignity is a concept recognized by copyright and implies that the concept of dignity has meaning for machine authors. If one purpose of copyright is to protect the dignity of authors with respect to their works, then a thorough discussion should begin by explicitly identifying and defining the concept of dignity as recognized by copyright, followed by an acknowledgement of any necessary assumptions, such as the assumptions that machine authors will understand the concept of dignity, hold it in a similar regard, and be capable of acting on their own behalf to protect that interest.

We will discuss additional implied assumptions in response to the questions below, but observe generally that many of WIPO's questions assume clear delineations between works generated autonomously and works with some level of human involvement. We suspect that, even where machines clearly generate works with minimal human involvement, other issues will confound the assignment of any rights in works generated, such as human and machine co-authorship, different machine entities sharing common hardware or software infrastructure, and the potential for ephemeral or transitory machine entities that may only exist momentarily. WIPO also seems to assume sustained definitional clarity between humans and machines as legally recognized entities, but making such distinctions may become increasingly difficult. We address Issue 6 in greater detail below.

(i) Should copyright be attributed to original literary and artistic works that are autonomously generated by AI or should a human creator be required?

(ii) In the event copyright can be attributed to AI-generated works, in whom should the copyright vest? Should consideration be given to according a legal personality to an AI application where it creates original works autonomously, so that the copyright would vest in the personality and the personality could be governed and sold in a manner similar to a Corporation?

(iii) Should a separate sui generis system of protection (for example, one offering a

reduced term of protection and other limitations, or one treating AI-generated works as performances) be envisaged for original literary and artistic works autonomously generated by AI?

In CDT's view, a better initial inquiry would be: *For what reasons might we grant copyright in autonomously-generated works?* If the answers to this question pertain to human-centric values or motivations from the perspective of creators, such as financial compensation or moral rights, then the distinction of autonomous generation is unnecessary— human motives and values are the driving force regardless of the level of human involvement. Further, if human owners or developers of autonomous systems are the ultimate beneficiaries of any rights assigned to works generated by those systems, as implied by question (ii), then the grant of copyright in autonomously generated works undermines copyright's central theme of rewarding creators by instead rewarding ownership.

We encourage WIPO to resist the view that all creative works must need legal protection. To that end, we suggest a counterfactual inquiry: *What outcomes may result if no copyright is granted for autonomously generated works?* On this point, CDT disagrees with WIPO's assertion that allowing copyright for machine-generated works would result in the "availability for the consumer of the largest number of creative works." This statement implies a causal relationship based on at least two assumptions: that copyright promotes public availability and that machine creators respond to the financial incentives copyright protection creates. CDT encourages WIPO to elaborate on the assumptions and reasoning behind this assertion.

As to the proposal to recognize autonomous creators as legal personalities, CDT suggests that further inquiry into the implications and consequences of such recognition is necessary. In this regard we propose the following questions:

Would a machine based legal personality be able to advocate for itself or make decisions on its own behalf?

If not, what persons or entities would act on behalf of the personality?

If so, by which goals is the personality motivated?

In either case, would existing legal protections and remedies function as intended?

How would infringement actions work with respect to the AI personality?

Unless creative machines are sufficiently intelligent, autonomous, and empowered to act on their own interests, then the prospect of granting copyright in autonomously generated works is ultimately a question of deciding which humans will reap the financial rewards for either making (for which IP protection already exists) or owning the creative machine. Regardless of which humans benefit, this would create a new rift between authorship and copyright. We urge WIPO to avoid this outcome.

CDT also proposes further inquiries with respect to WIPO's interest in creating a separate legal system to address machine-generated works. We note that the most undesirable aspects of

divorcing copyright's benefits from actual authorship cannot be mitigated by creating a separate system, but offer the following observations and questions to illustrate additional complexities:

How should we assess the term of copyright?

How would this assessment change with the recognition of legal personalities for AI systems?

How should the term of copyright be assessed for works co-authored by a human and a machine?

What role should formalities, such as registration, play?

Would humans be capable of assessing the originality of massive quantities of machine-generated works?

If not, how would our existing concept of originality align with a computer-based system for assessing originality?

In summary, CDT urges WIPO not to create a new copyright system that rewards machine ownership, but to consider instead what qualities those machines should exhibit to justify granting rights in their creations.

Issue 7: Infringement and Exceptions

(i) Should the use of the data subsisting in copyright works without authorization for machine learning constitute an infringement of copyright? If not, should an explicit exception be made under copyright law or other relevant laws for the use of such data to train AI applications?

As a primary matter, CDT suggests that answering this question begins with a careful look at what the “use of the data subsisting” in copyrighted material means in the context of machine learning. When machine learning systems “ingest” data, they are measuring various aspects of the individual elements and mapping the relationships among them. More simply, they extract facts about the data and look for patterns among those facts. It is through this process that systems discern patterns, correlations, and logical relationships in the dataset, and it is from these relationships that rules and predictive models are derived and refined. The use, then, of a copyrighted work is as a source of information representing one or more data points in a larger collection of data points. This kind of use is fundamentally different than a use where some or all of an original work is reproduced as part of the new work.²

Based on this understanding of the use of data in machine learning, CDT suggests that where works used as data elements are protected by copyright, and those elements are reproduced at some stage of the process, such use should not be considered infringing. Even under the

² CDT notes that authors and judges have discussed other computer-related uses of copyrighted works as “non-expressive,” and would agree that the use in question here is also non-expressive. See, e.g. James Grimmelman, *Copyright for Literate Robots*, Iowa Law Review, Vol. 101:657 (2016); *Authors Guild v. Google, Inc.* 804 F.3d 202 (2nd. Cir. 2015). However, using copyrighted works in machine learning differs from those uses because most algorithmic models retain neither the content of nor the data about the works from which they are derived. Rather, those models represent mathematical and logical representations of the rules and patterns observed in the data.

relevant limitations and exceptions, such as fair use and fair dealing, the kinds of uses for works contemplated in current machine learning processes should not result in a finding of infringement.

If we are to apply the law consistently, perhaps considering the use of copyrighted works from a human perspective would be appropriate. For example, the act of measuring and extracting facts about creative works, if performed by humans, would be considered research because it is the methodical search for new information—facts— *about* the work in question. Likewise, using copyrighted data elements for the purpose of developing, refining, or testing an algorithmic model could equally be construed as “teaching” in the same way that exposing students to the works of previous artists helps develop and refine their artistic sensibilities. For example, in some methods of machine learning, programs are explicitly instructed as to the relationships between data elements and related information. One familiar example of supervised learning is the “mark all the boxes containing x” exercise used by some websites to deter bot users, in which a human identifies and tags portions of images with identifying labels.³ In most jurisdictions, use of works by people for either research or teaching is not considered infringing.

A more structured analysis of these uses under the relevant exceptions and limitations yields similar results. Given the wide variety of fact patterns for using copyrighted data elements to train or test automated systems, a single fair use analysis according to the four-part statutory test may not accurately represent all possible outcomes.⁴ However, we offer the following observations about some aspects many such uses will have in common, with respect to the fair use test.

Factor 1: Purpose and character of use- As noted above, use of copyrighted data elements for developing or training a machine learning system is always for the purpose of either research or teaching. This should weigh strongly in favor of a finding of fair use. CDT is unaware of any cases in which a developer of a machine learning system markets any reproductions or portions of copyrighted data elements used in development. However, to the extent that the ultimate application of the model developed using the copyrighted data elements is commercial in nature, CDT suggests that such commercial use would only weigh against a finding of fair use when the application negatively impacts the market for the work (under Factor 4). We also note the several degrees of separation between the use of a copyrighted work as one data element in a training corpus and the outputs of an automated system in which a model derived from the

³ See, e.g., James O’Malley, *Captcha if you can: How you’ve been training AI for years without realizing it*, TechRadar (Jan. 12, 2018) <https://www.techradar.com/news/captcha-if-you-can-how-youve-been-training-ai-for-years-without-realising-it>

⁴ We perform only a general analysis under the United States doctrine of fair use, but note that the test, including individual elements, and the outcome is similar under our understanding fair dealing. We also suggest that use of works as training data (or a legislative exemption addressing such uses) would fit well within the more general framework for exceptions and limitations set forth in the three-step test of Article 9(2) of the Berne Convention. However, we see no value in creating additional exceptions for machine learning uses in jurisdictions incorporating either fair use or fair dealing into their copyright systems.

training data is used. As this relationship between the work used and the ultimate output of the model becomes less direct, the weight assigned to any commercial uses should decrease.

Factor 2: Nature of copyrighted work- Given the diversity of copyrighted works that may be used as a corpus for machine learning, general analysis under this factor is difficult. However, we note that even for the most creative works of artistic expression, it is the information derived *about* the work and its relationship to other data elements that is relevant to building, training, or testing a model. This analytical approach attenuates the importance of a work’s creativity. From this perspective, the nature of the copyrighted work should not weigh strongly in the analysis.

Factor 3: Amount and substantiality of portion used- As with the second factor, a uniform assessment under Factor 3 is not possible given the diversity of training methods used for machine learning. But even assuming that entire copyrighted works are reproduced in the training process, this factor would not necessarily weigh against a finding of fair use.⁵ Moreover, in most cases no part of any reproductions made in the training or testing of a model become part of the final product. Rather, the models or rules derived from the training data reflect the relationships or commonalities found among the data elements. From this perspective, the works used are a source of information but the content of the works is not “borrowed” in the way that a traditional fair use assessment contemplates. Instead, the use of copyrighted works in machine learning is more like an art student studying the paintings of great masters in a museum and incorporating their impressions of those works into their mental model of what classical paintings look like. Although reproductions may be made in the process, either in the mind of the student or the memory of the computer, the content of the reproduced work does not become part of the final product. Unless some or all of the original work remains in the system, this factor should not weigh against a finding of fair use, regardless of the portion used in the training or testing of a model.

Factor 4: Impact on the market for the work- The wide variety of machine learning applications again makes a uniform fair use assessment difficult. CDT is aware of some generative adversarial networks (GANs) that are capable of producing visual or literary works.⁶ Although it is unclear whether those works have any effect on the market for or value of the works used to train the network, we note that the relationship between the use of copyrighted works to train a GAN and the market for either the original works or the GAN-produced works is indirect at best. As with the other factors, the “use” here is not of the same kind that traditional fair use assessments contemplate, where portions of one work are re-used in another work. In most cases, exposing a machine learning system to works does not result in portions of the training

⁵ *Authors Guild v. Google, Inc.* 804 F.3d at 221 (2nd. Cir. 2015) (Noting that although Google copied entire works, it did not make them available to the public in a way that undermined the authors’ interests in distribution.)

⁶ See, e.g. *These Works of Art Were Produced By Artificial Intelligence*, Duke Today, (Mar. 18, 2019) <https://today.duke.edu/2019/03/these-works-art-were-created-artificial-intelligence>.

dataset becoming embedded in either the models or any output of the system.⁷ Instead, it is information about the training dataset that shape the models and the resulting outputs.

Existing exceptions and limitations like fair use offer more flexibility than special-purpose exceptions.

Despite the unconventional use case machine learning presents, CDT believes fair use or fair dealing can adequately address the legality of such uses. Further, fair use offers advantages over other possible legal mechanisms for allowing the use of copyrighted works in the context of machine learning. For example, in places where databases and datasets have additional use restrictions, such as the *sui generis* database protection in the European Union, policymakers have been pressed to create exceptions to accommodate machine learning.⁸ Yet despite prolonged negotiations, the resulting exception for text and data mining (TDM) is so rigid and restrictive as to prevent many beneficial uses of datasets.⁹

In particular, the TDM exceptions suffer from two major flaws. First, the exception created in Article 3 limits the uses to “non-commercial” uses by a narrow set of “research organisations and cultural heritage institutions.”¹⁰ These limitations prevent many legitimate uses of copyrighted works in datasets that pose no risk to authors’ ability to exploit their own works. Second, Article 4, which is more permissive in terms of uses and users, allows authors to deny use of their works.¹¹ This ability creates logistical and practical barriers to uses of large datasets in which some, but not all, authors have opted out of the exception because the additional efforts required to identify, cull, or negotiate individual licences for works make using the dataset unfeasible. In comparison, fair use allows more legitimate, non-infringing uses of copyrighted works while still protecting the interests of the rightsholders and does not suffer from the rigid definitional limitations or the practical barriers presented by statutory exceptions such as Art. 3 and 4.¹²

(ii) If the use of the data subsisting in copyright works without authorization for machine

⁷ CDT acknowledges that some outputs of GAN systems may be similar to existing copyrighted works, but we take no position as to the legal implications of that possibility as a general matter.

⁸ Directive (EU) 2019/790 of the European Parliament and of the Council on copyright and related rights in the Digital Single Market and amending Directives 96/9/EC and 2001/29/EC, Art. 3, 4 (Apr. 17, 2019). (“DSM Directive”)

⁹ See, CDT, CDT’s Concerns on the European Commission’s Proposal for a Directive on Copyright in the Digital Single Market, (2017), https://cdt.org/wp-content/uploads/2017/02/CDT_Concerns_EC_Proposal_Directive_on_Copyright_DSM.pdf. See also, Bernt Hugenholtz, The New Copyright Directive: Text and Data Mining (Articles 3 and 4), Kluwer Copyright Blog (July 24, 2019) <http://copyrightblog.kluweriplaw.com/2019/07/24/the-new-copyright-directive-text-and-data-mining-articles-3-and-4/>

¹⁰ DSM Directive, Art. 3(1).

¹¹ DSM Directive, Art. 4(3).

¹² See, Vadym Kublik, *EU/US Copyright Law and Implications on ML Training Data*, Valohai, <https://blog.valohai.com/copyright-laws-and-machine-learning>, last visited Feb. 14, 2020.

learning is considered to constitute an infringement of copyright, what would be the impact on the development of AI and on the free flow of data to improve innovation in AI?

Considering the use of copyrighted works for machine learning to be an infringement would add both friction and liability into the process of training and testing algorithmic systems. This additional friction and liability would effectively increase the potential and actual costs of developing and testing new systems, which would tend to slow the development and improvement of AI systems. Additionally, increasing the transactional costs of using datasets containing copyrighted works will encourage developers to turn to lower cost, lower friction alternatives. Unfortunately, these data sets are smaller, older, and may not accurately reflect the world as it is. Using them to train and test algorithms incorporates their limitations and bias into the rules and models from which AI systems are built. We discuss the issue of bias in more depth in response to Issue 10.

Considering the use of copyrighted works to train or test algorithms to be an infringement of copyright would also tend to preserve the concentration of large training corpora in the possession of a few firms. Many of the largest datasets are possessed by a small set of companies. These firms acquired this data, much of which is subject to copyright, directly from their subscribers who grant licenses for the use of their works as a condition of using the company's products and services. While those licenses may allow the company (and affiliates) to make use of the gathered data, they may not provide for broader use. Therefore, even if the company that collected the data wishes to make the data more broadly available for machine learning uses, anyone not covered by the licensing agreements could face liability for using protected works included in the dataset.¹³

Although the liability for infringement and the friction of licensing would tend to inhibit the free flow of data and negatively impact the development and innovation of machine learning applications, it is unclear what positive impacts this legal perspective might provide. Therefore, and in keeping with the analysis provided in response to Issue 7(i), CDT suggests that, at least in cases where none of the training data set is reproduced in or by the final product, using copyrighted works to train algorithmic models should not be considered an infringement of copyright.

(iii) If the use of the data subsisting in copyright works without authorization for machine learning is considered to constitute an infringement of copyright, should an exception be made for at least certain acts for limited purposes, such as the use in non-commercial user-generated works or the use for research?

(iv) If the use of the data subsisting in copyright works without authorization for machine

¹³ CDT acknowledges that, in some cases, it may be desirable to have a legal mechanism to prevent some uses of data to guard against violations of privacy or to protect other interests unrecognized by IP law. We limit our comments here strictly to the context of data uses in machine learning with respect to those interests recognized by United States IP law.

learning is considered to constitute an infringement of copyright, how would existing exceptions for text and data mining interact with such infringement?

(v) Would any policy intervention be necessary to facilitate licensing if the unauthorized use of data subsisting in copyright works for machine learning were to be considered an infringement of copyright?

(vi) How would the unauthorized use of data subsisting in copyright works for machine learning be detected and enforced, in particular when a large number of copyright works are created by AI?

As noted in response to Issue 7(i), purpose-built exceptions such as might be developed to account for uses of copyrighted works for machine learning often provide unsatisfactory results. Further, WIPO has identified three additional considerations that add significant practical complications: reconciling existing exceptions, structuring mechanisms for licensing, and detecting and enforcing against infringement. CDT offers no comment on the issue of how exceptions might interact with a legal perspective that use of copyrighted works for machine learning constitutes infringement. As to the need for and practicality of policy interventions to facilitate licensing, we note that creating a new licensing mechanism to enable authorized uses of copyrighted works for machine learning would require solving at least three outstanding issues related to copyright and digital media. We present those here in the form of additional questions WIPO might consider:

First, by what means can an entity (social media platform, rights management group, collection agency, etc) correctly attribute works to their original authors? Absent some type of formalities that create a record of both originality and authorship, this task seems either impossible or destined for massive inaccuracies and abuses. Even the copyright attribution systems employed by sophisticated providers of user-generated content services demonstrate the inadequacies of these systems and their potential for abuse.¹⁴ In addition to the difficulties presented by attributing works to authors claiming copyright in works, there is also the problem of “orphan works” for which no author can be identified. Although some jurisdictions have implemented extended collective licensing programs to address this issue, it is unclear what values or interests such a system would protect in the context of machine learning, beyond the impliedly de minimis commercial value of the works for these uses. Establishing such a system for the use of orphan works in machine learning would almost certainly create a windfall for collecting agencies, but would not result in meaningful compensation or control for uncredited authors.

Second, what impact would a licensing mechanism have on the existing models for data aggregation? The largest initial aggregators of data in the form of copyrighted works (social media platforms) already include licensing agreements in their terms of use. Establishing a system with the intent to provide monetary compensation to rightsholders would likely not

¹⁴ See Electronic Frontier Foundation, *The Mistake So Bad, That Even Youtube Says Its Copyright Bot “Really Blew It”*, <https://www EFF.org/takedowns/mistake-so-bad-even-youtube-says-its-copyright-bot-really-blew-it>, (last visited Feb. 14, 2020).

change this practice, but could negatively impact the ability of newer, smaller data aggregators to compete. In CDT's view, it is unlikely that the potential value of any licensing royalties would be enough to counter the network effects enjoyed by the dominant data aggregators.

Third, assuming the attribution problem can be solved and that a viable alternative licensing mechanism can be devised, to what extent could a rightsholder negotiate the terms of a license? Individual negotiations (on a per work or per author basis) would create prohibitively high transaction costs for all parties. But a compulsory blanket licensing approach would largely eliminate rightsholders' abilities to control the uses of their works. Either outcome would undermine at least one of the dual purposes of the mechanism- facilitating authorized use of works, but in contexts approved by rightsholders.

Finally, we address the issue of detecting and enforcing against infringing uses of works in the context of machine learning. Although most uses in this context should not be considered infringing, we offer the following observations.

First, detecting such uses is and will be prohibitively expensive, in terms of time and resources. This is especially true given that the majority of such uses do not significantly impact an author's ability to capitalize on the commercial value of their works in traditional markets and that, even if a market for the use of works in machine learning existed, the value of individual works is vanishingly small in this context. Even where authors may wish to control use of their works on the basis of one or more moral rights, they are unlikely to succeed in most cases. Developers do not always document which data they used, nor does every machine learning project result in a publicly accessible application. Even if authors become aware of an application for which there is sufficient documentation to indicate that one or more of the author's works may have been used in developing the application, it may be impossible to determine whether a given work in a dataset was actually reproduced or otherwise used in the process. In CDT's experience, the most commonly used datasets involving copyrighted works are composed of works for which authors have granted some type of general license, rendering successful infringement claims even more unlikely.¹⁵

Second, we note that the relationship between the first and second parts of the question is unclear. How do the outputs of an automated system relate to the individual elements of the dataset used to develop that system? The question may contemplate systems designed to produce new works in the style of the works included in the training data. If so, then perhaps the question is probing human artists' ability to review all such generated works for any that are substantially similar to their own. But there are many types of machine learning applications, many of which produce outputs with little relation to the elements in their training data.

Issue 8: Deepfakes

¹⁵ See, *Common License Types for Datasets*, <https://help.data.world/hc/en-us/articles/115006114287-Common-license-types-for-datasets>, last visited Feb. 14, 2020.

Should the copyright system take cognizance of deepfakes and, specifically, (i) Since deepfakes are created on the basis of data that may be the subject of copyright, to whom should the copyright in a deepfake belong? Should there be a system of equitable remuneration for persons whose likenesses and “performances” are used in a deepfake?

CDT fears that this issue ventures too far afield from the interests protected by copyright. We encourage WIPO to more fully articulate the relationship between algorithmically manipulated video and the interests of performers from a copyright perspective. We also reject the notion that a person’s likeness is or should be protected by copyright, even though recordings of performances or other depictions of a person’s likeness may be. Instead, tort and privacy law are better suited to protect a person’s interest in their own likeness.

Although deepfake technology raises many difficult questions, it is not clear from the questions posed in the consultation document which of the interests protected by copyright are implicated. To the extent that deepfake videos include reproductions of copyrighted material, it would seem that existing copyright law is sufficient to address any infringements such reproduction may constitute. However, the question as presented may be read to imply that a person’s likeness, appearance, or mannerisms are subject to copyright protection. Although unstated, the question may also refer to one or more “moral rights.” In either case, encourage WIPO to be more explicit as to which rights protected by copyright may be implicated. To the extent that WIPO contemplates a potential expansion of copyright protections as a way to “take cognizance of” deepfakes, we encourage more explicit articulation and explanation of why such expansion would be necessary and what forms any proposed expansion may take.

Finally, we take issue with both the premise and the phrasing of part (i)- “since deepfakes are created on the basis of data that may be subject to copyright, to whom should the copyright in a deepfake belong?”- which seems to imply that facts (data about works) are subject to copyright and that the rights to new works created from or inspired by older works might belong to someone other than the creator of the new works. The applications that produce “deepfake” imagery are derived from the factual relationships observed between features of the still images comprising a video recording. These feature attributes are extracted and reconstructed into new images generated from predictions about how those features might be arranged in a new context. Although some deepfakes may involve the reproduction of portions of copyrighted works, the generated images that overlay or replace portions of those images are new works. This is no different than an artist creating a series of sketches depicting the likeness of another person based on a visual observation; copyright vests in the author, not the subject. Therefore, to imply that the copyright for images depicting a person’s likeness should somehow vest in the subject, rather than the author of those images is to invert a fundamental aspect of copyright and would have far-reaching impacts on all forms of visual expression.¹⁶ We encourage WIPO

¹⁶ As noted above, tort and privacy law provide more appropriate legal remedy for unwanted depictions of a person’s likeness.

to take a cautious approach if it contemplates amending the interests protected by copyright and to be as explicit as possible in addressing the impacts of new technologies with respect to those interests.

Issue 10: Further Rights in relation to Data

Given the non-expressive, non-exploitive nature of using copyrighted works in the development of many algorithmic models and other machine learning tools, it is unclear why or what kind of additional recognition authors need when their works are included in data sets used for machine learning. As noted above, the most commonly used datasets involving copyrighted works are composed of works for which authors have granted some type of general license.¹⁷ For many of these licenses, no additional recognition is required. Others, such as the CC-BY license, require attribution when works are reproduced, but such attribution makes little sense when reproductions are not made publicly available, but are instead made for the convenience of the researchers and developers working with the dataset. To require public attribution for each author of a work in a dataset from which a model was derived would be like asking a novelist to list each and every author whose works they read at any time before drafting, for each would represent a datapoint in the author's mental model of the appropriate content and structure of a novel. CDT sees no reason to pursue this result in the machine learning context.

Individual works within data sets are already protected against unauthorized, exploitative uses. Those protections address the incentives a human-centric IP system intends—motivating the creation of new works by protecting authors' ability to monetize them. Existing copyright law already inhibits the use of works for machine learning; adding additional protections for particular arrangements of the data, on top of the existing rights for the authors of individual copyrighted elements, would further hinder, rather than promote, the progress of science and the useful arts.¹⁸

First, to the extent that researchers and developers of automated systems find value in particular arrangements or curations of databases or data sets, those arrangements and curation efforts are best made by the researchers and developers to suit their particular needs. But the first party to arrange or curate a data compilation should not be able to prohibit others from rearranging or differently curating that database, nor should they be able to prevent others from arranging the same data in the same fashion. As discussed above, jurisdictions with additional protection for databases are now facing difficulty reconciling their *sui generis* rights with the need to access and use data. CDT strongly suggests that copyright policy in the United States should avoid this problem by not extending additional protections to databases.

¹⁷ See, *Common License Types for Datasets*, <https://help.data.world/hc/en-us/articles/115006114287-Common-license-types-for-datasets>, last visited Feb. 14, 2020.

¹⁸ See Arjan Wijnveen, *How copyright is causing a decay in public data sets*, LinkedIn, Nov. 28, 2016, <https://www.linkedin.com/pulse/how-copyright-causing-decay-public-datasets-arjan-wijnveen>.

Second, further restricting access to and use of databases and data sets will limit researchers' ability to reduce bias in algorithmic models. The lack of access to sufficiently representative data is a common contributor to biased algorithms.¹⁹ Conversely, access to more diverse data sets helps researchers test for and mitigate bias in automated systems.²⁰ Even under existing copyright law, the "friction" associated with using some of the largest, most comprehensive data sets steers researchers and developers toward data sets that are easier to access and involve less legal risk.²¹ This also shrinks the world of data sets available to researchers, which is a bias in of itself, regardless of the relative bias expressed in the datasets themselves. Additional protections for data sets would only increase that friction, making it even more difficult to combat bias in algorithmic models.

Finally, although there are other concerns with access to and use of data sets, such as preserving the privacy and confidentiality of certain data, those are outside the scope of copyright. CDT acknowledges that certain use restrictions may be appropriate for some kinds of data to ensure that the rights of data subjects are not infringed. However, those use restrictions need not be based in copyright. In fact, copyright is a poor tool to protect the rights of people whose personal information may be reflected in data sets because in many cases, the data subject and the copyright holder (if any) are not the same entity. Although we disagree that additional copyright protections are warranted for data sets (beyond any protections granted to the works therein), we strongly encourage WIPO to limit any proposed expansions of intellectual property protections to only those which directly advance the purpose of copyright.

Respectfully submitted,

Stan Adams
Deputy General Counsel & Policy Counsel
Center for Democracy & Technology
1401 K St. NW, Suite 200
Washington, DC 20005

¹⁹ See, generally, Amanda Levendowski, *How Copyright Law Can Fix Artificial Intelligence's Implicit Bias Problem*, *Washington Law Review*, Vol. 93: 579 (2018).

²⁰ *Id.*

²¹ *Id.* at 593.