



February 14, 2020

Re: DRAFT ISSUES PAPER ON INTELLECTUAL PROPERTY POLICY AND ARTIFICIAL INTELLIGENCE

Comments of the Wikimedia Foundation

Introduction and Background

The Wikimedia Foundation appreciates the opportunity to comment on the Secretariat’s Draft Issues Paper on Intellectual Property Policy and Artificial Intelligence. The Wikimedia Foundation is the non-profit organization that hosts and supports a number of free, online, collectively-produced resources, including Wikipedia. The Foundation’s objective is to help create a world where every human being can share freely in the sum of all knowledge.

Our mission and projects thus intersect with artificial intelligence and machine learning (“AI/ML”) in two major ways. First, we develop AI/ML tools to assist editors in improving our projects, such as finding gaps in Wikipedia’s coverage, or identifying poor quality articles or edits.¹ Second, the Foundation’s projects are designed to be used, reused, and repurposed, and since they are freely licensed, our volunteer editors’ works, organizing, and other contributions are frequently used in training AI/ML systems.

Our experience with these systems highlights an important consideration when discussing authorship in relation to AI/ML systems: such systems as they exist today are tools developed and deployed by humans. While the term “artificial intelligence” has long been associated with the concept

¹ See, e.g., Jonathan T. Morgan, *Designing ethically with AI: How Wikimedia can harness machine learning in a responsible and human-centered way*, Wikimedia Foundation (July 2019), <https://wikimediafoundation.org/news/2019/07/18/designing-ethically-with-ai-how-wikimedia-can-harness-machine-learning-in-a-responsible-and-human-centered-way/>; Miriam Redi et al., *Can machine learning uncover Wikipedia’s missing ‘citation needed’ tags?*, Wikimedia Foundation (April 2019), <https://wikimediafoundation.org/news/2019/04/03/can-machine-learning-uncover-wikipedias-missing-citation-needed-tags/>; Miriam Redi, *How we’re using machine learning to visually enrich Wikidata*, Wikimedia Foundation (March 2018), <https://wikimediafoundation.org/news/2018/03/14/machine-learning-visually-enriching-wikidata/>.

of artificial sentience, as it is used in research and industry today, it most typically means a system that performs tasks that previously required human cognitive processing, such as translating complex and varied audio recordings of speech into text. The most salient aspects of these systems — that they are complex; that their internal decision-making processes can be difficult to predict; or can run autonomously — do not mean that they constitute a cognizable legal entity or an author. Complex algorithms that govern ordinary computer systems do not raise novel questions of authorship; unpredictable processes, such as generating random numbers or simulating chaotic phenomena, are still the tools of those who use them; and autonomous action by machines is at least as old as clockwork.

Questions of machine authorship would likely blur distinctions between many different types of systems, all of which can be labeled as “artificial intelligence.” Attempting to define when a system is sufficiently autonomous or sufficiently “creative” to be an author risks extending authors’ rights to an ill-defined set of systems. Instead, the Organization and its member states would likely be better served by addressing the implications of AI/ML tools that are commonly being used and refined today: automated systems designed to perform specific tasks by one party, and deployed in particular applications of those tasks by another. Within that scope, there are a number of questions relevant to IP development, such as how authorship might be attributed between the tool’s creator and its user, and the rights implicated by the use of copyrighted works in training AI/ML systems.

There are also several critical considerations not only regarding how AI/ML developments affect IP law, but how IP law may affect the impact of AI/ML systems on individuals’ human and legal rights. For instance, copyright policy affects what data sets are used to train AI/ML, frequently leading to skewed and biased systems. Copyright policy should not be so restrictive as to create incentives for AI/ML systems to be trained only on public data, or works from before the early twentieth century. Furthermore, when AI/ML systems involve the personal information of individuals, assertions of trade secrets, copyrights, or other IP rights should not prevent affected individuals from exercising their privacy rights. Also, when AI/ML systems are used to make determinations that affect people’s lives, such as setting credit ratings, bail amounts, or potential for future violations, IP rights should not prevent inspection, auditing, and analysis of systems. Intellectual property laws should not be usable as pretexts to evade accountability.

Issue 6: Authorship and Ownership

In paragraph 12, the Secretariat asks if copyright should be attributed to works that are autonomously generated, or if a human creator should be required; how any copyright attributed to AI-generated works should be apportioned; and whether a separate system of protection should be considered. In brief, human creators should be required for the foreseeable future; rights should vest either in a human creator or no one, depending upon the amount of creativity and authorship

provided to the work itself by the human authors; and no sui generis systems of protection are necessary.

Presuming the possibility of a machine creator incorporates presumptions about the nature of AI/ML that are not often reflected in reality, and doing so may even distort the proper analysis of other copyright questions involving AI/ML. The current state of machine learning and artificial intelligence encompasses many algorithms and processes that fall far short of a system that could be considered an independent author. The term “artificial intelligence” is commonly used to describe systems that are simply more sophisticated versions of older automated systems that rely upon large sets of data and the repeated application of statistics. They are therefore merely tools that help people work more efficiently. Thus, the concept of a protectable work produced without the involvement of a natural person is far more niche — a natural person will necessarily be involved in any works of authorship in the near future, either by setting up all of the conditions necessary for a machine to output a predicted work, or by creating the instructions necessary to combine aspects of chance and probability that will generate an unpredictable result.

Beginning an inquiry into copyright law and AI/ML with the premise of a machine author with its own authorial rights runs the risk of skewing or needlessly complicating perspectives on authorship disputes involving the types of AI/ML systems currently in use and that will continue to be used into the future. This is not to say that the questions dealing with the exclusive rights of an artificial being in its creative works should never be addressed; it is merely to say that they likely should be addressed with a lower priority than more fundamental rights for such beings (such as rights to existence, liberty, freedom from involuntary servitude, or freedom of expression), and that even those more basic rights are unlikely to become non-academic questions decades after more prosaic matters need to be litigated or ruled upon. Should machine personhood become a practical topic, existing legal structures for the authorship of natural persons would likely be easier to adapt in that environment, rather than creating a system meant for a putative being that may have broad, unintended consequences regarding the use of more commonplace, simpler tools today.

A more likely model for potential authorship disputes concerning AI/ML would involve one individual using an AI/ML tool that was created by a second individual. This model may provide insight into Question 12(ii) without needing to first make determinations about non-human authorship, and suggest further guidance on how AI/ML systems can be accounted for by current legal doctrines. Presupposing that AI/ML systems are, or could be, non-human authors needlessly inserts an uncertain legal construct into established law.

In many jurisdictions, non-natural persons cannot be authors. For instance, the United States Copyright Office has consistently taken the stance that works produced through fully randomized machine processes or natural processes may not be registered because protections only extend to work

produced by a human (e.g., randomized linoleum flooring designs or naturally “carved” driftwood).² This does not prevent works of authorship from containing elements that were created by automated, or even natural, processes. However, in those cases where authorship is found, it is vested in the person who deliberately set the automated process in motion in a creative way, or who contributed human creativity by modifying or combining and arranging the results of natural or automated processes in a sufficiently creative way. A piece of mounted driftwood carved and polished by waves may not be copyrightable, but an artist might incorporate such a piece into a copyrighted work, or might create a stencil to shape the effects of waves on a beach on a wooden sculpture, and copyright the result.

In this case, a natural person is contributing expression by using a process of nature as a tool. Where, instead of a natural process, an artificial, automated process is involved, we can look to the natural person as the potential author, and the separate process, however complicated it may be, as simply a tool that they used. A writer that uses an automated spell-checker, or even an automated thesaurus, is still the author of her written work; we do not, absent some extraordinary circumstance, attribute any authorship to the developer of those automated systems, even if they incorporate AI/ML or otherwise perform their functions without the writer’s explicit direction.

In evaluating whether the tool user is the proper author of the creation, existing law and jurisprudence on copyrightability and traditional standards of copyright can adequately address the matter. For instance, United States law articulates a threshold of creativity necessary for a protected work in *Feist Publications, Inc. v. Rural Telephone Service Co.*³ This case can be directly applied in the United States to works created with the assistance of an AI/ML tool: if the overall work created with the assistance of the tool failed to meet the threshold of creativity, *Feist* would still apply to limit the application of copyright protections.

To give another example, a sculptor using a 3D printer provides all of the creative input, even if the printer contains operating software that was created or owned by its manufacturer. Even if the sculptor were to make an error allowing the printer to run longer than intended (thus running autonomously, and without the sculptor’s immediate creative input) and produce an unexpected result, it would still be the sculptor who would have copyright in the work created as a result of their mistake. If all the sculptor did was turn on the printer and it extruded a blob of plastic, the sculptor’s actions would likely fall short of creative expression necessary for authorship, and the resultant blob, no matter how evocative, would simply not be protected by copyright.

As for the contribution of the tool makers — those who designed or contributed to the AI or process, and those who provided or selected inputs to it — they are, in most cases, unlikely to have

² United States Copyright Office, *Compendium of U.S. Copyright Office Practices* §§906.6-906.7 (2017), <https://www.copyright.gov/comp3/chap900/ch900-visual-art.pdf>.

³ 499 U.S. 340 (1991).

contributed sufficient expression to the work in order to be considered authors. For instance, even if an AI tool has already been programmed to create images, a person must decide to use the tool to produce a desired output. Therefore there will nearly always be contribution from a natural person present in the result of the image creation tool. The extent to which that person is the author of a protectable work then relies upon *Feist* and other existing doctrines of authorship.⁴

When an individual person uses an AI/ML tool to make a protectable work, there may be a question as to whether the tool's maker, as well as the tool's user, may have a claim to authorship in the created work. As with traditional questions surrounding potential joint works and contributions, the specific nature of the contribution matters a great deal — far more so than the question of whether any system marketed as AI/ML is involved in the process.

AI/ML systems that add substantial amounts of material with very little input from their end users may create situations where assigning authorship may seem counterintuitive, but these questions do not need to be resolved by designating an author, much less deciding that the creators of the AI/ML system are the authors. Instead, existing doctrines regarding “thin” copyright may apply: if an end user has made sufficient creative decisions in the use of the AI/ML tool, he might have a copyright in the resulting work, even while a separate end user, making similar decisions with the same tool, may also have a copyright in her independently created work, regardless of its substantial similarity to the first user's.

If the end user's inputs are insufficiently creative to merit authorship, however, there is no particular reason that the AI/ML system's developer should be assigned the copyright. If a simpler algorithmic system were merely choosing from a limited set of hand-coded results created by the developer, those specific results, if protectable expression themselves, could be considered works authored by the developer. For instance, if a 3D printer manufacturer created a model of a horse and included it with the printer as a test pattern to be printed at the push of a button, the manufacturer, and not the button-pusher, might have an interest in the resulting sculpture. However, where the results of the process, such as from an unpredictable AI/ML system, could not be said to be sufficiently determined, selected, or fixed by the developer or other system contributor, the answer may simply be that the work has no legal author and thus is not protected by copyright.

Treating such a work similarly to a product of nature avoids several complications and confusions. For instance, if a default rule held that the developers of a system were the authors of every work made via that system, they could become the holders of copyrights in works whose very existence

⁴ Manufacturers or developers of some programs and systems would be considered authors of those programs and systems themselves, insofar as they contain protectable expression. This is distinct from any copyright interest that might exist in outputs made by using those programs. For instance, while Microsoft may have a copyright in its word processing software Word, it has no copyright interest in texts written using Word.

they were unaware of. AI/ML systems developed in collaborative processes as free / open source software could create countless works with a similar multitude of authors, likely opening the door to punishingly complex disputes over ownership of the works, without a clear reason for any of the parties involved to justify such ownership.

Furthermore, it is likely that allowing an algorithm or process to create protectable works of authorship absent expressive input from a human would create significant unintended consequences. As one example, consider ubiquitous programs such as word processors or proofreading programs. As they improve, they are beginning to suggest common language, identify points of ambiguity, and provide automated rephrasing of written works. If their contributions rose to a level that might constitute a derivative work, the programmers of a word processor could potentially assert a blocking ownership right in every work written on their software. Similar AI/ML tools used in other fields could easily create abusive restraints on competition via such rights. Yet if algorithms gain the ability to create protectable works of authorship without proximate human expression, there is a high chance that people far removed from the output of the algorithm will nevertheless try to claim ownership of such outputs.

Question 12(iii) asks whether a separate sui generis system of protection should be envisioned for works autonomously created by AI. Such a system would seem unnecessary, considering the existing history of law dealing with joint and collaborative works, as well as works that incorporate unpredictable elements within their creative expression. To the extent that copyright laws need to better distinguish between reproductions of a computer program and uses of a computer program, the elaboration of such distinctions should extend to computer programs generally, and not be limited to computer programs that are labeled as AI/ML.

Issue 7: Infringement and Exceptions

The questions in paragraph 13 ask whether the use of data subsisting in copyright works without authorization for machine learning should constitute an infringement of copyright, and if not, whether there should be explicit exceptions created for the use of such data.

Many of these uses should likely not be considered an infringement of copyright, particularly if the uses do not prejudice the interests of authors with respect to the normal exploitation of their works, such as by providing competing means of an audience consuming the works. As the Secretariat notes, this situation is very similar, and in many cases, identical to, text and data mining scenarios, which have existing exceptions. Such exceptions, whether explicit to text and data mining, or contained within broader types of limitations and exceptions like fair use or fair dealing, can be dealt with via existing legal structures.

One major reason for this is that whether or not a particular system is AI/ML has less relevance to whether an author's rights were affected than other criteria. Lawmakers would likely create fewer unintended consequences by ensuring that limitations and exceptions include valuable uses that do not impact the market for a creative work in its role *as* a creative work. The purpose of the use and the end effect on the market for the original work are more likely to be relevant than the specific type of technology (or the marketing associated with it) that is employed by the user of the work. The scope of technologies that can be labeled AI/ML are broad and ill-defined enough that evaluating the category as a whole would require a fact-specific inquiry into the purpose and effect of various systems.

Even if particular instances of AI training require a fact-specific inquiry, we believe that existing law that deals with the processing of multiple copyrighted works⁵ can be applied directly to these issues. As noted above, the relative sophistication or complexity of the system in question should not alter the basic inquiries in copyright law: whether one of an author's exclusive rights has been implicated, and whether that implication occurs within the bounds of a limitation or exception to the scope of the exclusive rights granted to the author.

Issue 8: Deep Fakes

Paragraph 15 asks two separate questions related to “deep fakes”: how to determine copyright ownership in a deep fake created using data that may have been subject to copyright; and whether there should be a system of remuneration for persons whose likenesses are used in a deep fake. Each of these questions deserves a separate analysis.

As for determining copyright ownership, whether or not the resultant work is called a “deep fake” is largely irrelevant to questions of ownership and infringement. If copyrighted works were used in the creation of a deep fake, the ownership analysis should be identical to that of a situation where copyrighted works were used to generate a work that bore no resemblance to a real person.

Questions of the rights of a person depicted in a deep fake will vary, depending upon whether the deep fake merely imitates the person's likeness, or reproduces a creative work that they have made. The rights of a person whose mere likeness is replicated should not inhere in copyright. Rights of publicity or personality may be implicated, but a person's bare likeness is, in the vast majority of cases, not a creative work authored by that person. A person's performance, however, can be protected by copyright, and if a specific, previously fixed, and protectable performance is replicated, that may implicate the performer's rights to that performance work.

⁵ See, e.g., *Authors Guild v. Google, Inc.*, 804 F.3d 202 (2d Cir. 2015); *Authors Guild v. HathiTrust*, 755 F.3d 87 (2d Cir. 2014); *A.V. ex rel. Vanderbye v. iParadigms, LLC*, 562 F.3d 630 (4th Cir. 2009); *Perfect 10, Inc. v. Amazon.com, Inc.*, 508 F.3d 1146 (9th Cir. 2007).

However, whether deep fake technology is used or not is largely beside the point for both of these cases. If a painter or a computer graphics modeler replicates a person's likeness, even in the absence of deep fake or AI/ML technology, the considerations are the same: there is likely no copyright implication, and possible implications regarding rights of publicity. Likewise, if a digital animator manually recreates a specific, protected dance performance, copyrights in the performance may be implicated whether or not AI/ML technology is used.

Issue 9: General Policy Issues

Question 16(i) asks whether there are unforeseen consequences of copyright on bias in AI applications, or if there is a hierarchy of social policies vis-a-vis human creation versus innovation.

We wish to emphasize that copyrights and AI/ML may intersect in more critical ways than how copyright policy balances technological innovation against artistic creation, or how AI/ML systems may complicate liability. That more important intersection deals with ongoing questions of how AI/ML and other black-box systems and algorithms are used to make decisions that have drastic effects on individuals' lives,⁶ and how copyright and other IP policy can hinder accountability for them.

One substantial problem is bias in AI/ML systems.⁷ When selecting training data, developers of AI/ML systems will often select public domain or freely licensed works, both for ease of access and to avoid potential infringement liability exposure. This creates a particular form of selection bias in training data, which can often lead to other forms of more harmful bias in results. For instance, a system ingesting published books in English from before 1923 will likely exclude a disproportionate number of authors of color; systems using biographies and news stories of people from before this cutoff date to correlate them with occupations would fail to associate women with jobs that were denied to them in the early 20th century or before, and carry those systemic biases into the future.⁸

Copyright, along with other types of IP, should also not be misused to create barriers to algorithmic transparency. Already, we are seeing confidentiality being used as an excuse for companies not to make their algorithmic decision making processes available for audit. This is especially concerning in cases where algorithmic decision making is used to make decisions about people's lives,

⁶ See, e.g., Frank Pasquale, *The Black Box Society: the Secret Algorithms that Control Money and Information* (2015); Virginia Eubanks, *Automating Inequality: How High-Tech Tools Profile, Police, and Punish the Poor* (2017); Cathy O'Neil, *Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy* (2017).

⁷ See Aaron Halfaker, *Mitigating biases in artificial intelligences - the Wikipedian way*, Wikimedia Foundation (October 2018), <https://wikimediafoundation.org/news/2018/10/10/mitigating-biases-artificial-intelligences-wikipedian-way/>.

⁸ See, e.g., Amanda Levendowski, *How Copyright Law can Fix Artificial Intelligence's Implicit Bias Problem*, 93 Wash L. Rev. 579, 593 (2018), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3024938.

or their rights.⁹ Algorithms today are being developed or are already used to determine prison sentences, loan terms, credit ratings, and bail amounts. Online, algorithms are used to determine what content to surface to a wider audience on social media, and in some cases, what content to remove or block. Individuals should be able to trust that automated systems can and have been audited by someone with their interests in mind, and should be able to understand on some level how a decision was made.

While companies may rely primarily upon trade secret to obscure the inner working of automated decision making, the copyrightability of code has the potential to allow for misuse of copyright to deny algorithmic transparency. For instance, external audits are the primary way that biases are identified in algorithmic systems. Public interest organizations or other auditors of such systems will often need to make reproductions of copyrighted code in order to analyze it. Even if they do not explicitly copy source code onto their own devices, these researchers will inevitably make temporary copies in RAM as they interact with or merely use the AI/ML systems. Doing so may expose them to threats of liability for copyright infringement if they are using the software without the permission of the algorithm's developer, or potentially if they use the software contrary to the terms of the software license.¹⁰ Researchers investigating potential bias also frequently need to compile and share repositories of algorithmic results, for analysis and comparison; claims of copyright or other restrictions on these data outputs can prevent the public from learning of bias or discrimination present in the algorithms. Researchers investigating potential discrimination may also be exposed to liability for violating anti-circumvention measures.¹¹ These are not theoretical threats: recent history contains several examples of companies misusing copyright provisions to prevent transparency and accountability.¹²

Policymakers in the Member States and elsewhere should strive to ensure that copyright exceptions and limitations remain robust to allow for beneficial uses like news reporting, research, auditing, and public accountability, and must work with companies and NGOs to build processes by which IP rights may be respected, but not misused to prevent accountability as AI/ML systems affect the full scope of fundamental human rights. With respect to the concerns raised regarding the misuse of copyright and other IP rights to prevent algorithmic accountability, we refer the Secretariat to the Toronto Declaration,¹³ which outlines a human rights framework to be applied to the use of AI/ML

⁹ *Id.* at 597-606.

¹⁰ *See, e.g., MDY Indus., LLC v. Blizzard Entm't, Inc.*, 629 F.3d 928, 939-942 (9th Cir. 2010) (discussing possibility that violations of license conditions in use of otherwise legally-obtained software may result in a finding of copyright infringement).

¹¹ *See MDY*, 629 F.3d at 943-955.

¹² *See, e.g., Online Policy Group v. Diebold, Inc.*, 337 F.Supp.2d 1195 (N.D. Cal. 2004); Edward Felten, *The Chilling Effects of the DMCA*, Slate.com (March 29, 2013), <https://slate.com/technology/2013/03/dmca-chilling-effects-how-copyright-law-hurts-security-research.html>; Levendowski at 602-610 (noting attempts to use copyright law to prevent reverse engineering).

¹³ The Toronto Declaration, <https://www.amnesty.org/download/Documents/POL3084472018ENGLISH.PDF>

systems. Ensuring that AI/ML applications are designed and implemented within a framework of human rights, and not merely according to a *sui generis* set of ethical principles, can more properly ensure that various aspects of the law, including IP law, are coordinated in promoting justice and equality, in addition to science and the useful arts.¹⁴

Issue 10: Data

Issue 10 generally asks whether IP policy should consider the creation of new rights in data. Paragraphs 21-23 frame this question by noting that the several types of recognized IP protection are forms of protections of data, such as “data that represent inventions...” or “data that represent original literary or artistic works.” However, the category of data is far too broad for the frameworks of intellectual property to encompass adequately, and in many cases attempting to do so would be actively harmful for human rights.

Certain types of data are explicitly excluded from the realm of intellectual property. In fact, considering that facts cannot be protected by intellectual property in all but the most niche cases, the vast majority of data is properly not protected by intellectual property, whether that consists of fundamental properties of the universe, the facts of current or historical events, the location of objects in the world, or raw scientific data.

On a more conceptual level, attempting to protect “data” in an intellectual property framework faces fundamental flaws. Intellectual property generates legal protections by creating legal structures that allow certain types of ideas to be treated like tangible property that can be owned and exchanged. However, exclusive ownership is explicitly antithetical to many types of data, and alienability is antithetical with regard to others.

For instance, data includes news reporting and scientific facts; it includes opinions about leading political figures and members of one’s community. Granting exclusive ownership in such data violates human rights to freedom of expression. Data includes scientific and historical facts; creating exclusivities and ownership for these violates rights to education and learning.

Even when ownership seems like a useful characteristic for certain types of data, a property model for that data can result in violations of human rights. For example, rights to privacy, when treated as property, are subject to being traded away, or treated as a bargaining chip in interactions between parties of vastly unequal bargaining power.¹⁵ While many types of IP do create certain

¹⁴ See Anna Bacciarelli, *Ethical AI principles won’t solve a human rights crisis*, Amnesty International (June 21, 2019), <https://www.amnesty.org/en/latest/research/2019/06/ethical-ai-principles-wont-solve-a-human-rights-crisis/>

¹⁵ The debate around treating private data as intellectual property has been active for at least twenty years. See, e.g., Pam Samuelson, *Privacy and Intellectual Property*, 52 *Stanford Law Review* 1125 (1999).

protections for creators by allowing their creations to be treated as property, it does not follow that creating a commodity out of a concept results in its protection.

Rather than attempting to cabin all types of data into an intellectual property framework (with or without regard to developments in AI/ML), the Secretariat and Member States should examine the current successes and failures of existing forms of propertizing different types of intangible creations. Too frequently, value is measured by quantifying revenues and profits collected by large industries, without regard to how intangible value is taken from large groups of individuals. In other cases, measurements of value fail to account for effects on human rights that are difficult to quantify, that resist commodification, or that should not be alienated at any price.

The creation of new IP rights is both far broader than the scope of this paper, and far narrower than the universe of data, information, and knowledge that the questions in the paper appear to attempt to address. The Wikimedia Foundation would recommend that, should these questions be addressed in future, they be approached in a more focused manner: in that they: (1) be based upon specific types of intangible creations that seem likely to benefit from commodification; and (2) include an evaluation of the need for existing categories or modes of intellectual commodification.

Conclusion

The Wikimedia Foundation believes that developments in artificial intelligence and machine learning have comparatively modest effects upon the state of intellectual property law. Many existing doctrines already account for authors using automation to create works without controversy, and robust doctrines already exist to account for the more remote possibilities of non-human creation. Existing law has demonstrated an ability to account for mass processing of copyrighted works; the particular types of technologies associated with that processing should have less effect on the outcome of such cases than the specific facts of the individual cases themselves. Finally, it is critical that stakeholders in IP law recognize the potential for it to inadvertently perpetuate bias, or be abused by those who would seek to avoid accountability in their algorithmic systems. Stakeholders including policymakers, industry, and civil society must ensure that IP law remains flexible enough not to be misused in this way.