

Innovation, Science and Economic Development Canada

RE: Consultation on a Modern Copyright Framework for Artificial Intelligence and the Internet of Things

To Whom It May Concern:

We write in response to the request for comments regarding a modern copyright framework for artificial intelligence and the Internet of Things.

The Information Technology and Innovation Foundation (ITIF) comments draw on prior research about the critical role intellectual property plays in spurring innovation. ITIF is a non-profit, non-partisan public policy think tank focusing on a host of critical issues at the intersection of technological innovation and public policy. Its mission is to formulate and promote policy solutions that accelerate innovation and boost productivity to spur growth, opportunity, and progress.

Sincerely,

Nigel Cory
Associate Director, Trade Policy, Information Technology and Innovation Foundation

CONTENTS

Overview	2
Authorship and Ownership: AI-Based Work Should Receive IP Protection.....	3
AI-Based Copyright Work Should Reside With Its Owning and Controlling Entity.....	5
Text and Data Mining and Copyright: Allow AI to Use Material Just Like Any Other User Who has Legal Access to Copyright Material.....	6
Relevant Use Cases.....	8
Use Case: “Deep Fakes” and the Legitimate Uses of AI-Based Technology for Creative Purposes	8
Use Case: Facial Recognition and Creative Commons Licenses.....	10
Use Case: AI Musicians, Fair Use, and Copyright.....	10
Endnotes	12

OVERVIEW

Intellectual property (IP) is based on the idea that those who combine the spark of imagination with the grit and determination to see their vision become reality in books, film, music, technology, medicines, designs, sculpture, services, and more deserve the opportunity to reap the benefits of their innovation—and that this reward incentivizes more creative output. In the past, IP law operated under the assumption that all creative works would be entirely created by people. However, the advent of artificial intelligence (AI) has raised the prospect that in the future a significant number of works may be created by an autonomous computer system without direct human involvement. Canadian policy should protect the principle on which IP law is based, whether the works are generated by people or computer systems.

ITIF welcomes Canada’s detailed and thoughtful review of copyright implications for AI and the Internet of Things.¹ Central to Canada’s consideration of IP and AI should be a pragmatic and realistic understanding of AI and its capabilities. While AI systems can be increasingly autonomous and creative, they still have a considerable way to go before they achieve the sophistication that many people imagine. However, Canada should not create AI-specific requirements for activity that is already legal and being done through non-technical means.

ITIF’s submission focuses on two key recommendations:

1. Canada should recognize and protect AI-generated IP and the need to assign ownership of AI-generated work to the person or organization that owns the AI system.
2. Canada should allow all users—whether commercial or non-commercial—to use AI for text and data mining (TDM) so long as they have legal access to the material. There should be no additional, special approvals for the use of TDM tools. Canada should avoid making its copyright framework overly complicated in considering TDM-specific requirements and exceptions and instead ensure that the use of TDM is consistent with current IP law, whether this relates to infringement or licensing.

AUTHORSHIP AND OWNERSHIP: AI-BASED WORK SHOULD RECEIVE IP PROTECTION

Work produced by AI should receive IP protections. The innovation-incenting function of IP law should not change based on whether a computer or a person creates the output. While it is true that an AI system does not respond to the financial incentives that are central to IP-based innovation policy, the entities who develop and own AI systems most certainly do. Canada should ensure its laws and policies recognize (and not prohibit) non-human creativity and innovation, while clarifying the criteria, rules, and legal linkages for assigning rights to relevant owners (whether individuals, organizations, or firms).²

As AI becomes ever-more sophisticated, it is moving closer to joining the “creative class.” While this occurrence has been debated within IP circles for decades, we’re finally at a point where it is becoming a common practical issue, rather than a hypothetical scenario or ad hoc and novelty occurrence.³ There are two ways of looking at AI, but both come to the same conclusion about how AI-created works should be treated by IP law.

One way is to consider AI as just another tool, analogous to dice, in providing non-deterministic outputs. In these situations it is still the human creator who determines how to use non-deterministic outputs to generate creative content.⁴ Another is to consider AI as something different than a tool, and more like an independent person. Imagine a parent trains a precocious child on research techniques, and then that child uses this training to discover something new on their own. Generally speaking, this child would get to claim full IP rights for their discovery.

Extending this analogy to AI systems, some may say an AI system is like the child trained by a grown up, and that anything it then creates should be the rightful IP of the AI system. While the human may have trained the AI system, as the parent trained the child, the AI system did all the work. And while an AI system could not exist without its human owner (to pay for hardware and electricity and keep it running), neither would a child survive without the parent (paying for food, clothing, and shelter). Indeed, AI systems are now capable of creating new works and new ideas on an increasingly autonomous basis. If IP law does not recognize works produced by AI systems, then there is a risk that humans will have little incentive from IP law to create and train AI systems that can autonomously create new ideas. So in this situation, as in the first one, the human owner of the AI system should receive the IP rights for any discoveries and creative works produced by an AI system.

Canada needs to ensure its IP law allows humans to be assigned authorship of IP produced by AI systems given that AI will likely be able to produce genuinely new and novel creations that would receive IP protections if produced entirely by a human. Canadian IP law should not be biased against using machines.

For example, in 1998, John Koza developed an algorithm as part of an “invention machine” that created simple circuit designs. In 2005, Mr. Koza’s machine passed one of the first “Turing tests” after a patent examiner (not knowing it was created by a computer) showed demonstrable creativity (a “non-obvious step” in IP law) and granted it a patent, making it one of the first intellectual property protections ever granted to a nonhuman designer.⁵ However, as ITIF argues should be the case, Mr. Koza was the one who actually obtained the patents generated from the machine’s output.

As long as the AI-based creation meets the conditions that statutory subject matter must meet before receiving IP protections, it should receive protection. Policymakers should be consistent in recognizing AI-based creations that would otherwise receive protection if it had a human creator/inventor. Although the AI system itself may be the proximate creator of the work, others, such as the owner/controller of the AI system at whose initiative the work is ultimately created, should be entitled to ownership of the AI’s works.⁶

Nothing in international treaty law explicitly authorizes, or prohibits, protections for such computer-generated works.⁷ The Berne Convention states the Union is created, “for the protection of the rights of authors in their literary and artistic works,” without defining who the author is, which it does due to the fact that, “national laws diverge widely, some recognizing only natural persons as authors, while others treat certain legal entities as copyright owners.” Furthermore, this debate is not new. In 1965, the U.S. Copyright Office’s annual report (in the section “Problems Arising From Computer Technology”) raised concerns about whether computers can author musical works.⁸ In the late 1980s and early 1990s, the World Intellectual Property Organization (WIPO) considered protections of “computer-produced works” in discussions of a possible (but never achieved) model copyright law.⁹ The debate around AI and IP is simply the latest iteration of this IP policy debate.

Canada should build a framework to make two key differentiations: between users who merely push a button on a computer that uses AI and the genuine authors/creators of AI; and between the use of computers and AI as part of the process for creation (as a general tool) and when AI contributes substantially to the creation of IP itself. Virtually all types of copyrightable work is regularly created using computers, whether it’s Microsoft Word for literary works, architectural works using Autodesk’s AutoCAD, or pictorial and graphic works enabled by Adobe Photoshop. There is a clear dividing line between the creative techniques used by digital and analog authors. Cutting and pasting is easier and faster on a computer, but the verbs “cut” and “paste” betray their analog origin.¹⁰ The AI we’re referring to is the segment of technology that can be considered sophisticated and smart, not simply a process for automating a process as a matter of efficiency. But some AI systems go beyond relatively minor increases in efficiency and perform virtually all of the key tasks that lead to the creation of a new work.

If Canada wants to be a world leader in innovative digital technologies, it needs to ensure it doesn’t discriminate against what is likely to be one of the most transformative general-purpose technologies (similar to the microprocessor) of the future. Some advocates argue that computer-generated works should become public property.¹¹ While there is and should not be any restrictions on allowing an individual to use AI to create public domain content, requiring this would remove the incentive for people and organizations to invest the considerable time and effort into using an increasingly important driver of innovation. In the case of AI, the proximate creator may not be human, but the owning and controlling entity will be, whether it is an individual or organization. Despite AI’s increasing sophistication, AI does not operate autonomously, in a vacuum, outside of an individual or organization’s ownership and control—it is another source of creativity people and firms use to drive innovation.

Other countries provide realistic, pragmatic models and reference points for Canada, such as countries that provide copyright protection to computer-generated works, such as in Hong Kong, India, the United Kingdom, New Zealand, and Ireland.¹² In 1988, the United Kingdom became the first country to include provisions on “computer-generated works” as part of its Copyright, Designs and Patents Act (“CDPA”). It states that these are works “generated by computer in circumstances such that there is no human author of the work.”¹³ Furthermore, as it relates to authorship, the CDPA provides that, “[i]n the 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 arrangement necessary for the creation of the work are undertaken.”¹⁴ Of note, this protection only extends to literary, dramatic, musical and artistic works and not to media works (sound recordings, films, cable programs, and published editions), although a similar system also applies with regard to design rights.¹⁵ It’s therefore not a major leap to see how this could be adapted and extended to include AI-created copyright material and patented products. Rather than treating AI-based work as a work “generated by a computer in circumstances such that there is no human author of the work,” a computer-generated work should be a work “generated by a computer in circumstances such that the computer, if a natural person, would be an author.”¹⁶

Such a pragmatic approach stands in contrast to parts of the debate in Europe, which has tended to focus on much broader legal and ethical debates about how to account for and manage AI. For example, a 2017 report from the European Parliament’s Committee on Legal Affairs included a diverse, confusing range of legal, policy, and ethical issues and proposals around AI and robotics.¹⁷ In this same report, the Commission on Civil Law Rules on Robotics called for the European Commission to elaborate on “the criteria for ‘own intellectual creation’ for copyrightable works produced by computers or robots.” Meanwhile, the Committee on Industry, Research and Energy stated that next steps should “respect contractual freedom and leave room for innovative licensing regimes” and “cautions against the introduction of new intellectual property rights in the field of robotics and artificial intelligence that could hamper innovation and the exchange of expertise.”¹⁸

The report also called on the European Commission to create a specific legal status for robots, including bestowing “electronic personality” to robots when considering legal responsibilities in cases where they’ve caused injuries or damages.¹⁹ In reaction, a group of 156 AI experts from 14 European countries sent a letter to the European Commission warning that granting robots legal personhood would be inappropriate from a legal and ethical perspective. (One could easily envision this leading to absurd results, such as a robot having to spend time in jail for injuring a human). The main focus of the opposition to the European Parliamentary proposal was that it would absolve robotics manufacturers of legal responsibilities for the actions of their machines (“I didn’t build a faulty robot, the robot itself was faulty.”)²⁰

This debate in Europe focuses on one avenue for accounting for AI-based legal concerns, such as those related to IP creations: bestowing a degree of personage to the machine. A pop-culture analogy can be found in the “The Measure of a Man” (episode nine of the second season) of *Star Trek, the Next Generation*, which deals with question about Lieutenant Commander Data (a male android with a built-in AI system) and whether he is a piece of Starfleet equipment or a sentient being.²¹ In line with this, Malta is considering a robot citizenship test as part of its national AI strategy.²² Similarly, others have suggested that computers should hold IP rights and that these could be shared under contract.²³ But of course, the big difference between “Mr. Data” and a Roomba robot is that the former had consciousness and free will. Roomba has neither and just bumps into things until it changes direction. And for the foreseeable future, AI will resemble Roomba, not Mr. Data.

AI-Based Copyright Work Should Reside With Its Owning and Controlling Entity

The human involvement that Canadian policymakers should focus on is on that of the owning and controlling entity using the AI for creative purposes. The involvement of a natural person should not be required for AI-based works to receive copyright protection, as long as it meets all the other criteria. But there needs to be a clear connection to an owner/operator, whether this is a person or corporation. For decades, machines have been autonomously generating works which have traditionally been eligible for copyright and patent protection.²⁴

In doing so, Canada needs to differentiate between computer-generated works and the associated problem of who should be considered the rightsholder of these works. This issue of assigning authorship/ownership for AI-based work is akin to other, analogue copyright issues. The same challenges will arise in determining whether the work is an infringing copy, an unlawful derivative work, a lawful derivative work, a joint work, or a sole-authored work. Yet, this is more a legal question about how IP is generated and shared and used, rather than a question about whether the IP should be recognized at all.

To see this, imagine one scenario where someone buys an “AI machine” and sets it up so it generates songs and automatically posts those songs online with a royalty-free license as free for anyone to use. And it produces and uploads 100 songs a minute in perpetuity. The relevant IP question is not should the purchaser of the AI machine be able to copyright the songs or whether the songs are in the public domain. The relevant

question is what the owner of the machine chooses to do (license the AI's output, or let be in the public domain) and whether the AI machine is producing infringing work.

In line with this, Canada should (conceptually) look at AI as some combination of an autonomous tool and employee/subcontractor—both work within the confines of a legal entity, who benefits from their labor, including their creativity. Outside of computer-generated works, U.S. copyright law already has a mechanism for authorship by artificial “persons.” In the case of a work made-for-hire, the employer for whom the work was prepared is considered the author for legal purposes.²⁵ Regardless of how autonomous AI may be, AI will be operating within the confines of a legal entity's operations.

For example, in the case of drug discovery, researchers at Harvard University, the University of Toronto, and the University of Cambridge created a generative model, which they trained on 250,000 drug-like molecules, that generated plausible new molecular structures.²⁶ The AI can create drug compounds more independently of humans (although it does still suggest structures which are sometimes nonsensical) and without lengthy simulations.²⁷ These institutions controlled the AI and therefore should benefit from their work by being assigned any associated IP.

Questions about whether there should be some threshold or criteria for involvement for someone to be recognized as an author/creator related to AI is similar to debates that afflict analogue IP applications about who should or shouldn't be recognized as an author/creator. Obviously, it would be useful for Canada to use this consultation process as the foundation for further research (as much as it can in the absence of relevant court cases and legal precedents) on where and how to draw the line between human involvement in AI that could allow an individual to qualify for authorship of an AI-based work, such as those on curating training data, observing processes, or setting AI's orientation. However, this will be difficult in the absence of legal changes that first allow AI-based works to receive IP and corresponding court cases that help analyze these contextual factors that will likely go into a decision about where the line should be drawn.

The *Acohs vs. Ucorp* case in Australia highlights some of the potential issues Canada will need to consider as part of its research, especially the potential difficulties in establishing authorship in the context of computer-aided or generated work.²⁸ Both firms in the case used software to help clients produce individualized reports (which each contain unique source code) about hazardous materials. Their systems are interoperable in that Ucorp used reports from Acohs. Acohs initiated a copyright-infringement case against Ucorp on the basis that its system was infringing the copyright of its reports. Acohs was not challenging Ucorp's use of its reports, which is part of the implied license.²⁹ The court rejected Acohs's claim that the source code was an original literary work as there was no evidence that the user, in entering the details of the hazardous material they needed information on, had in mind the resulting source code when they did this. However, the judge did recognize that the reports were original literary works, but Acohs's case for protection failed on the basis that the judge could not identify joint authorship (of both the source code and the report output). The judge rightly rejected Acohs's position on the ground that it was an artificial concept that the computer programmers and the clients (who entered their data) collaborated with each other in the writing of the source code and in creating the reports. This highlights the challenge Canadian policymakers face in not only protecting the source code of AI, but the products they produce, and assigning authorship within the context of both scenarios.

TEXT AND DATA MINING AND COPYRIGHT: ALLOW AI TO USE MATERIAL JUST LIKE ANY OTHER USER WHO HAS LEGAL ACCESS TO COPYRIGHT MATERIAL

Text and data mining (TDM) are powerful tools that allows researchers in a wide range of disciplines, from bioinformatics to digital humanities, to plough through texts and datasets and interpret minute details. In the field of computational linguistics (e.g., human language technology, or natural language processing), some

experts estimate that TDM already accounts for about 25 to 30 percent of all research projects.³⁰ For scholars and scientists, access to the rigorously scrutinized work of their peers, such as academic journals and databases, has always been a vital resource. Researchers who subscribe to these sources can explore them using traditional keyword searches and meta-tags predefined by publishers, but that has serious limitations. However, manually reviewing all of these sources is a slow and tedious process, the results of which are often inaccurate and incomplete. TDM gives researchers the ability to not only find a needle in a haystack, but to quickly find and categorize all manner of small objects hidden in many hundreds or thousands of haystacks.

For example, medical researchers can use technologies like natural language processing to quickly analyze the outcomes of thousands of clinical trials. Similarly, TDM technologies can be used to process the data contained in a large collection of scientific papers in a particular medical field to suggest a possible association between a gene and a disease. In line with this, Elsevier—a global information analytics business—allows subscribers to carry out independent text mining research and offers customized text mining services in the fields of life sciences and pharmaceutical research.³¹ Indicative of its management of IP, Elsevier can add (upon request, no doubt after payment, licensing, and rightsholder recognition issues are agreed) other licensed content to the database that its clients want to use for TDM.³² This type of analysis supports efforts to develop data-driven precision medicine initiatives that use the latest evidence to deliver personalized treatments. Data mining cannot provide all of the insights gained from human experts closely studying texts, but it does allow researchers to use rapidly developing tools to draw on a much larger pool of literature and data to support their work.

Copyright law can stifle the use of text and data mining on resources they can legally access and analyze with non-automated means. The use of data mining on copyrighted material often falls foul of existing intellectual property laws because the technical process involves accessing and extracting data from its original source and copying it into another database for analysis. For example, even if individuals or firms can lawfully access and read the material, such as through a university library subscription, copying a substantial part of works may infringe copyright in those works (what is ‘substantial’ depends on the context and circumstances). Canadian copyright law should allow publishers to set the subscription fees for access to their content, prohibit unauthorized reproductions of their content, and receive appropriate compensation. But it should not require people with lawful access to content, such as paid subscribers, to seek approval from publishers for using automated research methods.

However, combining copyrighted works (assuming the user has legal rights to use the material) into a searchable database that yields useful information is a permissible fair use. There is case law that suggests that acts of incidental or intermediate copying which do not ultimately result in the external re-use of protectable (expressive) parts of a copyright work should not be considered infringing.³³ After all, there is nothing illegal about “mining” databases manually; AI-based technology only automates the process. For example, the “Google Book’s” project provides a searchable index of quotes and snippets of text from 25-plus million books it scanned (under fair use).³⁴ (The main legal issue was whether the quotes and snippets go beyond fair use.) Canada should ensure that its IP frameworks allow people/firms to do this, while respecting the rightsholder and ensuring that proper licensing arrangements and payments are made and followed.

At the heart of the issue with TDM and copyright is the need to strike the right balance between access and use and respecting and supporting the role of IP and rightsholders’ ability to benefit from it. TDM has become a copyright issue around the world due to various ways to set this balance.³⁵ In the United States, TDM is governed by “fair use,” especially as it relates to “non-consumptive uses.” The U.S. copyright regime is considered more favorable to TDM practices than what appears to be the case under European laws, in part, because of the inherent flexibility of the U.S. fair use doctrine.³⁶

Canada should keep in mind a central point as it considers TDM and copyright—that AI should use (as part of computational analysis) copyrighted material like everyone else. In many cases AI simply automates processes and uses that were already (legally) taking place. If users already have legal access to the material they are using, then this shouldn't be an issue. This highlights the fact that there is no need to rewrite the rules of copyright simply because AI systems are now one of the potential users of publicly shared content. In line with this, Canada should not create an additional permission to access data that they already have legal access to use. In practice, copyright holders (or those that manage their portfolios) will well find ways to “upsell” on TDM access. For example, by creating APIs that allow for easier data requests or multiple copies or licenses that allow for sharing the data with other third-party data processors. That is fine, and useful, but doesn't require new IP rights.

TDM mining should not be based on getting special approval, require any specific misuse provisions that are already covered by IP law, and force the single use of data. The consultation details three potential safeguards and limitations that are misguided and redundant. Canada should not make a TDM-specific conditionality as it does in its listing of potential safeguard and limitations (point 1). Whether it's via a license or contract, or via creative commons material online, if the access is legal, TDM users should not need an additional approval. Similarly, a user should not have to take additional measures to prevent the distribution of reproductions of works or other subject matter if using TDM as if this is done illegally, it would already be covered by existing IP law (point 2). Furthermore, Canada should not require a person using TDM to not further reproduce any works as this would prevent re-use of data that a user has already legally collected and used (point 3).

Overall, Canada should be ambitious and allow all users—whether for commercial or non-commercial, such as academia—to use TDM on resources they are legally accessing and analyzing with non-automated means. In 2014, the United Kingdom enacted legal changes that provide an exception for non-commercial uses.³⁷ Similarly, the European Union considered a specific exception for research institutions to carry out text and data mining on lawfully accessed, copyright-protected works.³⁸ This exemption is reasonable because it creates a special dispensation for data mining and does not alter other laws that prohibit the unauthorized extraction or reproduction of copyrighted works. However, Canada, the United Kingdom, and the EU should allow everyone to take advantage of these more efficient and effective data-driven research methods.³⁹

RELEVANT USE CASES

The consultation asks how individuals and organizations are using AI to produce or to assist in the production of works or other copyright subject matter. This section details: the use of the same AI technology used in “deep fakes” for genuine, creative purposes; the use of creative commons licenses for facial recognition; and the rise of AI musicians and implications for fair use/fair dealing and copyright.

Use Case: “Deep Fakes” and the Legitimate Uses of AI-Based Technology for Creative Purposes

The use of AI-based technology for creative purposes, and how this relates to concerns about “deep fakes,” is indicative of how organizations are using AI to produce or to assist in the production of works or other copyright subject matter.

The increasingly realistic AI-based technology that allows content creators to tell the stories of real-life people and events in TV shows and movies—such as through docu-dramas—can also be used maliciously to insert a person's likeness into objectionable content, especially fake pornographic content featuring famous actresses and falsified speeches by politicians. Policymakers around the world are only just beginning to consider the impact of digital replicas and voice and image manipulation software on entertainers and expressive works.

And while it is not yet a prevalent issue, policymakers have only begun to grapple with the risk that deep fakes present for the spread of disinformation, including political disinformation and election misinformation.⁴⁰

The underlying technology has many legitimate creative uses. Movie/TV producers have used similar, but less advanced, technology for some time.⁴¹ More recently, major digital effects studios have used AI to convincingly map a famous actor's likeness onto another performer's to add value to their stories, as well cut down production time and costs. For example, 2018's *Guardians of the Galaxy 2* showcased a de-aged, 1980s version of star Kurt Russell, while 2016's *Rogue One* re-created Peter Cushing's Grand Moff Tarkin character from 1977's *Star Wars: A New Hope*, despite the fact that Cushing died in 1994.⁴²

While “deep fakes” are a troubling and rapidly growing problem, like many other technology issues, policymakers shouldn't ban the underlying technology given its many legitimate uses, but target cases where it is used for specific malicious purposes.⁴³ For example, without explicit exceptions for expressive works, some U.S. states' proposed bills (such as New York) targeting the use of deep fakes for malicious purposes would preclude movie/TV show creators from telling the stories of real-life people and events.⁴⁴ The First Amendment of the U.S. Constitution protects creators working on fiction, non-fiction, or some hybrid of the two. This is why any new policy related to deep fakes should explicitly exempt first amendment-protected uses, including news reporting, commentary, and analysis.

It can be more straightforward for a person to be protected from deep fakes through privacy laws when it involves non-consensual sex scenes, deep-fakes porn, and non-consensual nude performances.⁴⁵ Obviously, the challenge then becomes identifying the individuals involved in creating the deep fakes. Many U.S. states and countries also have “revenge porn” laws that criminalize the publication of an intimate image without consent (although it's a matter of seeing whether courts apply this law to the case of deep fake pornography videos).⁴⁶ There's the potential for these deep fake videos to constitute criminal harassment (depending on whether a state/country's laws include provisions about misusing a person's image/likeness).⁴⁷ Beyond using existing and new legal tools to address the malicious use of this technology, companies like Google and Facebook are actively developing tools to help identify deep fakes.⁴⁸

This AI-based technology is also relevant to IP issues when it is used to infringe publicity and likeness/personality rights. The right of publicity prevents the unauthorized commercial use of an individual's name, likeness, or other recognizable aspects of one's persona. It gives an individual the exclusive right to license the use of their identity for commercial promotion. In the United States, the right of publicity is largely protected by state law (where it varies from state to state). Only around half of the states recognize it.⁴⁹ However, federal unfair competition law provides a related statutory right to protection against false endorsement, association, or affiliation. Internationally, the rights analogous to the right of publicity are sometimes recognized as “personality rights,” “rights of persona,” or other similar terminology.

However, intellectual property law already provides a framework for the legitimate use of this AI-based technology. Musicians, actors/actresses, or their respective estates can give permission for others to use their likeness/publicity/personality rights as part of some new creation.⁵⁰ For example, hologram-based concerts featuring dead musicians may become more popular as the technology gets better. There have already been sold-out shows (with \$200 tickets) for a holographic performance by Roy Orbison, but again, this is done with the consent of the performer's estate.⁵¹ This leads to a key question for policymakers as to the intellectual property connection to cases where a “deep fake” is derived from copyright-protected material, as it would be infringing these rights given it's an unauthorized modification and republication.

Similarly, performers may not want their likeness to enter the public domain upon death. In the United States, the right of publicity and rights to digital replicas have long been a priority for the SAG-AFTRA union and its members.⁵² If there is value to a person's likeness after their death, their family or whoever makes up their

estate (whether this is friends or a designated charity) should receive the fruits of their labors. Some people would also no doubt want to be excluded from commercial use of their likeness all together (such that a person's.

Use Case: Facial Recognition and Creative Commons Licenses

The consultation asks how individuals and organizations are using AI to produce or to assist in the production of works or other copyright subject matter. The recent backlash to the idea of using publicly available photos to train facial recognition systems highlights some misunderstandings in how copyright law permits the use of copyrighted works for computational purposes, such as training an AI or machine learning systems.⁵³

In March 2019, IBM created the “Diversity in Faces” dataset to provide a set of photos of peoples’ faces of various ages and ethnicities to help reduce bias in facial recognition systems.⁵⁴ IBM compiled the dataset from photos people shared online with a license which allows others to use the images for any purpose (under a creative commons license).⁵⁵ The widespread use of these licenses has been a tremendous boon to society and the economy by creating a wealth of valuable content that others can freely use and adapt for their own purposes, and the “Diversity in Faces” dataset is a perfect example of how openly licensed works generate valuable benefits.

It is clear that IBM can lawfully distribute images with Creative Commons licenses. While some people may be opposed to facial recognition technology, and not like that their images were used to train some company's algorithms, that does not mean copyright law is broken or needs to be changed, or that IBM did anything wrong. As Ryan Merkley, the chief executive officer of Creative Commons, notes, “copyright is not a good tool to protect individual privacy, to address research ethics in AI development, or to regulate the use of surveillance tools employed online.”⁵⁶ It would be unfortunate if general public angst about AI led to the popularization of licensing agreements that explicitly prevent the computational use of data. Platforms like Flickr should resist any pressure and continue to offer technology neutral licenses, ensuring that any data a human can access, a computer can also access. And companies like IBM should be encouraged to continue to package datasets for public use.

Unfortunately, misunderstandings about how open licenses work are commonplace and responsible for new waves of outrage. More recently, an October *New York Times* article called attention to another facial recognition training dataset compiled from openly licensed Flickr photos called MegaFace.⁵⁷ People included in the database expressed similar distaste and frustration, failing to acknowledge that their (or, for people whose childhood photos were included, their parents’) use of open licenses explicitly allows for this kind of use. Ultimately, there are limits to the amount of control individuals have over content they share publicly—whether they do this online, in print, or in person. There is no need to rewrite the rules of copyright simply because AI systems are now one of the potential users of this publicly shared content.

Use Case: AI Musicians, Fair Use, and Copyright

The case of Avia Technologies and its AI performer—called Avia, which stands for “Artificial Intelligence Virtual Artist”—highlights how AI can be applied creatively and legally.

AVIA composes classical music, which is used in soundtracks for film directors, advertising agencies, and game studios. All of AVIA's music is copyright protected. AVIA was the first AI to be official recognized as a composer, having been registered under France and Luxemburg’ authors’ right society (where all of its copyright resides under its own name).⁵⁸ However, while AVIA's music is largely indistinguishable from the work of human musicians, its compositions still require human input with regards to orchestration and

musical production, which no doubt features in its copyright applications. But this points toward a broader potential model for revised intellectual property protections for AI creations, given it recognizes the role of owners/creators in managing their AI, much like current law would with regard to a human composer that worked as an employee or contractor.

AVIA uses machine learning to understand and model high-level abstractions in data, such as the patterns in a melody. On top of this, AVIA uses reinforcement learning, which teaches the software to decide what action to take next (which means it does not need labelled inputs or outputs of data). In this way, AVIA improves its performance without needing explicit instructions. Avia Technologies taught AVIA by feeding it large databases of classical music, which allows it to capture concepts of music theory in order to develop its own model of music theory, allowing it to compose its own music.

The founders of Avia Technologies focused on classical music for two reasons: it's the predominate style used in the background of movies, games, etc.; and all the training data—the classical music by Mozart, Beethoven, etc.—are in the public domain (as their copyright protection has expired).⁵⁹ This highlights a critical point—that just because digital content may be available, does not mean that the legal framework should disregard existing IP protections if it used as an input to AI creations.

In the case of copyright-protected material, Canadian law should look at AI-created music, and related concerns about infringement, in much the same way it already does in cases where another human artist who creates a song that sounds like (or samples directly from) an existing one. Much as human artists are influenced by other musicians, AI-based systems should also be able to listen to music in creating their own, original compositions. Obviously, if it creates undistinguishable copies or samples of the music without permission, that would be an infringement, just as George Harrison's composition of "My Sweet Lord" was ruled an infringement on the rights of Bright Tunes, a New York publisher, for the song for "He's So Fine" written by Ronnie Mack and part of its catalog. Similarly, if AI-based musicians are marketed as sounding like a particular musician (after being trained on their music), that raises the prospect of infringing publicity and likeness rights. Inevitably, there will be scenarios where assessing infringement by AI will be difficult, but this is not unique to AI. Just as the current debate about infringement between human artists and whether one uses notes, lyrics, or a melody of a song from another, Canadian courts would have to analyze the gray zone to determine when a case involves AI, music, and potential infringement.

ENDNOTES

1. Innovation, Science and Economic Development Canada, “A Consultation on a Modern Copyright Framework for Artificial Intelligence and the Internet of Things,” <https://www.ic.gc.ca/eic/site/693.nsf/eng/00316.html#s223>.
2. In line with: Ryan Abbott, “Artificial Intelligence, Big Data and Intellectual Property: Protecting Computer-Generated Works in the United Kingdom” (Research Handbook on Intellectual Property and Digital Technologies (Forthcoming), November 2, 2017), <https://ssrn.com/abstract=3064213>; Ryan Abbott, “Hal the Inventor: Big Data and Its Use by Artificial Intelligence” (Big Data Is Not a Monolith, MIT Press, February 19, 2015), <https://ssrn.com/abstract=2565950>.
3. World Intellectual Property Organization (WIPO), *Monthly Review of the World Intellectual Property Office* (Geneva: WIPO, monthly report no. 9, September 1990), https://www.wipo.int/edocs/pubdocs/en/copyright/120/wipo_pub_120_1990_09.pdf.
4. Some academics see AI as another tool, analogous to dice, in providing non-deterministic outputs that should not receive IP protections. James Grimmelmann, “There’s No Such Thing as a Computer-Authored Work - And It’s a Good Thing, Too,” *Columbia Journal of Law and the Arts*, 2016, volume 39: 403, <https://scholarship.law.cornell.edu/cgi/viewcontent.cgi?article=2617&context=facpub>.
5. Jonathon Keats, “John Koza Has Built an Invention Machine,” *Popular Science*, April 19, 2006, <https://www.popsci.com/scitech/article/2006-04/john-koza-has-built-invention-machine/>; See: U.S. Patent No. 7,117,186, U.S. Patent No. 6,532,453, and Patent No. 6,360,191.
6. Shlomit Yanisky-Ravid, “Yanisky-Ravid, Shlomit. “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*, 2017: 659: <https://digitalcommons.law.msu.edu/cgi/viewcontent.cgi?article=1199&context=lr>.
7. The Berne Convention and the World Trade Organization Agreement on Trade Related Aspects of Intellectual Property Rights (TRIPS).
8. The U.S. Copyright Office, *Sixty-Eith Annual Report of the Register of Copyrights (1965)*, (Washington, D.C.: U.S. Copyright Office, 1966), <https://www.copyright.gov/reports/annual/archive/ar-1965.pdf>.
9. The discussants defined a computer-produced work as one generated by a computer where identification of authors is impossible because of the indirect nature of individual contributions. The original owner of the moral and economic rights in such a work would be either the entity “by whom or by which the arrangements necessary for the creation of the work are undertaken,” or the entity “at the initiative and under the responsibility of whom or of which the work is created and disclosed.” WIPO’s Committee of Experts eventually concluded further study was needed, and the model law was never adopted. World Intellectual Property Organization (WIPO), *Monthly Review of the World Intellectual Property Office*.
10. Grimmelmann, “There’s No Such Thing as a Computer-Authored Work - And It’s a Good Thing, Too.”
11. Ralph D. Clifford, “Intellectual Property in the Era of the Creative Computer Program,” *Tulane Law Review*, 1997, https://scholarship.law.umassd.edu/fac_pubs/77/.
12. United Kingdom Copyright, Designs and Patents Act, 1988, c. 48, § 9(3); New Zealand Copyright Act of 1994, § 5; Ireland Copyright and Related Rights Act 2000, Part I, § 2 (Act. No. 28/2000).
13. World Intellectual Property Organization (WIPO), *Monthly Review of the World Intellectual Property Office*; “United Kingdom Copyright, Designs and Patents Act 1988,” <http://www.legislation.gov.uk/ukpga/1988/48/section/9>.
14. Robert Hart, “Copyright and computer generated works,” *Aslib Proceedings*, volume 40, no. 6, (1988): 173–181, <https://www.emerald.com/insight/content/doi/10.1108/eb051098/full/html>.
15. “United Kingdom Copyright, Designs and Patents Act 1988 §214.” <http://www.legislation.gov.uk/ukpga/1988/48/section/9>.
16. Abbott, “Artificial Intelligence, Big Data and Intellectual Property: Protecting Computer-Generated Works in the United Kingdom.”
17. “Report with recommendations to the Commission on Civil Law Rules on Robotics,” European Parliament Committee on Legal Affairs, January 27, 2017, https://www.europarl.europa.eu/doceo/document/A-8-2017-0005_EN.html.
18. Ibid.
19. Ibid.

-
20. Letter, “Open Letter to the European Commission: Artificial Intelligence and Robots,” <https://g8fip1kplyr33r3krz5b97d1-wpengine.netdna-ssl.com/wp-content/uploads/2018/04/RoboticsOpenLetter.pdf>; Janosch Delcker, “Europe divided over robot ‘personhood,’” *Politico*, November 4, 2018, <https://www.politico.eu/article/europe-divided-over-robot-ai-artificial-intelligence-personhood/>.
 21. Mikael Pulliainen, “The character of Data and his humanity in Star Trek: The Next Generation” (Bachelor’s seminar and thesis, 2018), <http://jultika.oulu.fi/files/nbnfioulu-201811103021.pdf>.
 22. Arif Khan, “SingularityNET to Collaborate with the Government of Malta on their National AI Strategy,” *SingularityNET*, November 1, 2018, <https://blog.singularitynet.io/singularitynet-to-collaborate-with-the-government-of-malta-on-their-national-ai-strategy-6a813ffb3987>.
 23. Colin Davies, “An evolutionary step in intellectual property rights - Artificial intelligence and intellectual property,” *Computer Law and Security Review*, volume 27, no. 6, (2011) 601-619, <https://www.sciencedirect.com/science/article/pii/S0267364911001518>.
 24. Ryan Abbott, “I Think, Therefore I Invent: Creative Computers and the Future of Patent Law,” *BC Law Review*, volume 57, (2016): 1079–1126.
 25. 17 U.S.C. § 201(b) (2011); Ryan Abbott, “Artificial Intelligence, Big Data and Intellectual Property: Protecting Computer-Generated Works in the United Kingdom.”
 26. Rafael Gómez-Bombarelli, Jennifer N. Wei, David Duvenaud, José Miguel Hernández-Lobato, Benjamín Sánchez-Lengeling, Dennis Sheberla, Jorge Aguilera-Iparraguirre, Timothy D. Hirzel, Ryan P. Adams, and Alán Aspuru-Guzik, “Automatic chemical design using a data-driven continuous representation of molecules,” *ACS central science*, vol. 4, no. 2 (2018): 268-276, <https://arxiv.org/abs/1610.02415>.
 27. Tim Simonite, “Software Dreams Up New Molecules in Quest for Wonder Drugs,” *Technology Review*, November 3, 2016, <https://www.technologyreview.com/s/602756/software-dreams-up-new-molecules-in-quest-for-wonder-drugs/>.
 28. “Copyright in source code and digital products,” Lavan website, August 3, 2010, https://www.lavan.com.au/advice/intellectual_property_technology/copyright_in_source_code_and_digital_products; “Acohs Pty Ltd v Ucorp Pty Ltd [2012] FCAFC 16,” Lexology, August 16, 2012, <https://www.lexology.com/library/detail.aspx?g=ec18c61f-683a-47db-a5e7-b2f6a3ec7e19>.
 29. “Acohs v Ucorp or the limits of implied licences,” Ipwars.com blog post, April 10, 2012, <http://ipwars.com/2012/04/10/acohs-v-ucorp-or-the-limits-of-implied-licences/>.
 30. Sergey Filippov, “Mapping Text and Data Mining in Academic and Research Communities in Europe” (The Lisbon Council, May 27, 2014), <https://lisboncouncil.net/component/publication/publication/109-mapping-text-and-data-mining-in-academic-and-research-communities-in-europe.html>.
 31. “Elsevier’s Text Mining,” <https://www.elsevier.com/solutions/professional-services/text-mining>; “Elsevier’s text and data mining policy,” <https://www.elsevier.com/about/policies/text-and-data-mining>.
 32. “Rapid Access to Relevant Answers: Elsevier Text Mining,” Elsevier R&D Solutions for Pharma and Life Sciences, https://www.elsevier.com/__data/assets/pdf_file/0019/164512/R_D-Solutions_Brochure_Text-Mining_DIGITAL.pdf.
 33. Matthew Jockers and Matthew Sag and Jason Schultz, “Brief of Digital Humanities and Law Scholars as Amici Curiae in Authors Guild v. Hathitrust,” (amicus briefing, June 4, 2013), <https://ssrn.com/abstract=2274832>.
 34. Scott Rosenberg, “How Google Book Search Got Lost,” *Wired*, April 11, 2017, <https://www.wired.com/2017/04/how-google-book-search-got-lost/>.
 35. Matthew Sag, “The New Legal Landscape for Text Mining and Machine Learning,” *Journal of the Copyright Society of the USA*, Vol 66 (2019). <https://ssrn.com/abstract=3331606>.
 36. “The Exception for Text and Data Mining (TDM) in the Proposed Directive on Copyright in the Digital Single Market- Technical Aspects” (briefing, European Parliament), [https://www.europarl.europa.eu/RegData/etudes/BRIE/2018/604942/IPOL_BRI\(2018\)604942_EN.pdf](https://www.europarl.europa.eu/RegData/etudes/BRIE/2018/604942/IPOL_BRI(2018)604942_EN.pdf).
 37. “Researchers given ‘data mining’ right under new UK copyright laws,” Pinsent Masons website, June 2, 2014, <https://www.pinsentmasons.com/out-law/news/researchers-given-data-mining-right-under-new-uk-copyright-laws>.
 38. European Commission, *Directive on Copyright in the Digital Single Market*, (Brussels: European Commission, 2016), <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A52016PC0593>.

-
39. Nick Wallace, "EU Bill On Data Mining Lacks Ambition," *Center for Data Innovation*, blog post, October 17, 2016, <https://www.datainnovation.org/2016/10/eu-bill-on-data-mining-lacks-ambition/>.
 40. Cade Metz, "Internet Companies Prepare to Fight the 'Deepfake' Future," *The New York Times*, November 24, 2019, <https://www.nytimes.com/2019/11/24/technology/tech-companies-deepfakes.html>.
 41. Patrick Shanley and Katie Kilkeny, "Deepfake Tech Eyed by Hollywood VFX Studios," *The Hollywood Reporter*, May 4, 2018, <https://www.hollywoodreporter.com/news/deepfake-tech-eyed-by-hollywood-vfx-studios-1087075>.
 42. Carolyn Giardina, "How the 'Guardians of the Galaxy' VFX Team Made Kurt Russell Bleed Blue," *The Hollywood Reporter*, February 6, 2018, <https://www.hollywoodreporter.com/behind-screen/how-guardians-galaxy-vfx-team-made-kurt-russell-bleed-blue-1081894>; Carolyn Giardina, "Rogue One: How Visual Effects Made the Return of Some Iconic 'Star Wars' Characters Possible," *The Hollywood Reporter*, December 16, 2016, <https://www.hollywoodreporter.com/heat-vision/rogue-one-how-grand-moff-tarkin-peter-cushing-returned-957258>.
 43. Such as in the case of facial recognition: "The Value of Facial Recognition in Law Enforcement," The Information Technology and Innovation Foundation event, July 24, 2019, <https://itif.org/events/2019/07/24/value-facial-recognition-law-enforcement>; "Banning Facial Recognition Technology in Public Housing Would Be Misguided," The Information Technology and Innovation Foundation media release, July 23, 2019, <https://itif.org/publications/2019/07/23/banning-facial-recognition-technology-public-housing-would-be-misguided-says>; Daniel Castro Michael McLaughlin, "Banning Facial Recognition in Police Body Cameras Will Make Californians Less Safe," *Innovation Files blog post*, September 10, 2019, <https://itif.org/publications/2019/09/10/banning-facial-recognition-police-body-cameras-will-make-californians-less>.
 44. Eriq Gardner, "Disney Comes Out Against New York's Proposal to Curb Pornographic 'Deepfakes,'" *The Hollywood Reporter*, June 11, 2018, <https://www.hollywoodreporter.com/thr-esq/disney-new-yorks-proposal-curb-pornographic-deepfakes-1119170>.
 45. Ashley Cullins, "FX Wins Appellate Court 'Feud' With Olivia de Havilland," *The Hollywood Reporter*, March 26, 2018, <https://www.hollywoodreporter.com/thr-esq/fx-wins-appellate-court-feud-olivia-de-havilland-1097518>.
 46. ITIF has called for a federal law to outlaw revenge porn. "ITIF Praises Proposed Federal Bill That Would Outlaw Revenge Porn," media release, July 14, 2016, <https://itif.org/publications/2016/07/14/itif-praises-proposed-federal-bill-would-outlaw-revenge-porn>; Daniel Castro Alan McQuinn, "Why and How Congress Should Outlaw Revenge Porn" (The Information Technology and Innovation Foundation, July 15, 2015), <https://itif.org/publications/2015/07/15/why-and-how-congress-should-outlaw-revenge-porn>; Ruobing Su, Tom Porter, Michelle Mark, "Here's a map showing which US states have passed laws against revenge porn — and those where it's still legal," *Business Insider*, October 3, 2019, <https://www.businessinsider.com/map-states-where-revenge-porn-banned-2019-10>.
 47. Ryan J. Black, Pablo Tseng and Sally Wong, "Canada: What Can The Law Do About 'Deepfake'?", McMillan law firm website, March 29, 2018, https://www.mcmillan.ca/What-Can-The-Law-Do-About-Deepfake?utm_source=Mondaq&utm_medium=syndication&utm_campaign=View-Original
 48. Cade Metz, "Internet Companies Prepare to Fight the 'Deepfake' Future," *New York Times*, November 24, 2019, <https://www.nytimes.com/2019/11/24/technology/tech-companies-deepfakes.html>.
 49. David Evans, "Can You Protect Your image Like Your Brand?," *WIPO Magazine*, May, 2015, https://www.wipo.int/wipo_magazine/en/2015/02/article_0008.html.
 50. Peter Allen, "A genuinely believable CGI actor? It won't be long," *The Conversation*, January 19, 2017, <http://theconversation.com/a-genuinely-believable-cgi-actor-it-wont-be-long-71410>.
 51. Wendy Lee, "Must Reads: Roy Orbison hologram concert in L.A. invites awe and debate," *LA Times*, October 6, 2018, <https://www.latimes.com/business/hollywood/la-fi-ct-orbison-hologram-20181006-story.html>.
 52. Sarah Alex Howes, "Digital Replicas, Performers' Livelihoods, and Sex Scenes: Likeness Rights for the 21st Century," *Columbia Journal of Law and Arts*, 42 (2018): 345, https://journals.cdrc.columbia.edu/wp-content/uploads/sites/14/2019/04/07.-Howes_FINAL.pdf.
 53. Joshua New, "Copyright Law Should Not Restrict AI Systems From Using Public Data," *Center for Data Innovation*, blog post, October 14, 2019, <https://www.datainnovation.org/2019/10/copyright-law-should-not-restrict-ai-systems-from-using-public-data/>.
 54. John Smith, "IBM Research Releases 'Diversity in Faces' Dataset to Advance Study of Fairness in Facial Recognition Systems," IBM website, January 29, 2019, <https://www.ibm.com/blogs/research/2019/01/diversity-in-faces/>.

-
55. Ryan Merkle, "Use and Fair Use: Statement on shared images in facial recognition AI," *Creative Commons*, blog post, March 13, 2019, <https://creativecommons.org/2019/03/13/statement-on-shared-images-in-facial-recognition-ai/>.
 56. Ibid.
 57. Kashmir Hill and Aaron Krolak, "How Photos of Your Kids Are Powering Surveillance Technology," *New York Times*, October 11, 2019, <https://www.nytimes.com/interactive/2019/10/11/technology/flickr-facial-recognition.html>.
 58. Bartu Kaleagasi, "A New AI Can Write Music as Well as a Human Composer," *Futurism*, March 9, 2017, <https://futurism.com/a-new-ai-can-write-music-as-well-as-a-human-composer>.
 59. Ibid.