

An Artist's Reply to Public Consultations on Generative AI Copyright in US & Canada

Written by: [Tim Boucher](#) ([Lost Books](#)), 26 October 2023

Introduction

The following document is a submission to the US Copyright Office's [Notice of Inquiry on Copyright and Artificial Intelligence](#) [Docket No. 2023–6], written by and submitted as a content creator using AI tools as part of the creative process. It is simultaneously being submitted to the [Government of Canada's public consultation](#) on generative AI and copyright. (The document is primarily written within the American context, but has strong applicability to Canada as well.)

The following consists of an artist's description of their multimedia ebooks made in part using generative AI as a case study, and speaks more broadly regarding issues related to copyright and artificial intelligence in the Arts, with some recommendations of potential paths to explore for solutions. A high level summary is included below for convenience.

Key Points

- Artists develop pioneering uses of new technologies, playing a [critical function in the innovation process](#) and the furtherance of science and social progress.
- Artists should consequently have the same incentives to create and legal protections over their creations afforded by copyright, regardless of the technologies used in their production, whether or not they include AI.
- Artists need to be able to analyze and compare past works in the creation of new works, including using AI to do so.

Summary of Recommendations

1. The Copyright Office should affirmatively enshrine authorship rights as belonging to the person who undertakes the arrangements for the making of the work, as the UK does with computer-generated works and Canada does with cinematography.
2. The Office should affirmatively clarify that use of copyrighted works to train AI generally qualifies as fair use, to reduce uncertainty.
3. If necessary due to substantial similarity concerns, the Office should develop a framework that assigns only thin copyright protection for certain categories of AI-generated outputs with low human-involvement, such that only near-identical copies might be considered infringing.

Artist's Statement

Description of Works

The [AI Lore Books](#) are a collection of short fiction ebooks featuring experimental combinations of human and AI-assisted text and image contributions. They use AI to augment human storytelling in a massive world-building sandbox. The books are published by Lost Books of Canada, an AI publisher run by author Tim Boucher, a dual US/Canadian citizen.

The genre of the AI Lore Books is dystopian sci-fi mixed with fantasy and hyperreality elements (where the borders between real and fiction is intentionally blurred to enhance the uncanny valley effect). Thematically the books address risks, fears, and possible futures for humans co-existing alongside ever more sophisticated AI technologies as they spin out of our control.

Drawing on historical contexts such as 'Golden Age' pulp science fiction magazines (where many of the legendary authors of sci fi earned their stripes), and the long tradition of serial fiction from centuries prior, the works number 116 volumes as of this writing. The books form multiple interlocking "networked" narratives, where each volume contains hyper-linked references to other related volumes, creating unique trails for readers to explore based on their interest through the world-building of the stories. In this way, the books also draw from cultural influences like the "Choose Your Own Adventure" genre.

Each volume generally consists of between 2,000 to 5,000 words, and ranges from about 40 to 150 images (occasionally above 200 images). Sometimes the images explicitly tie into and directly illustrate the accompanying text, whereas other times they drift moodily in other directions, resulting in a kind of fragmentary trip through another world entirely. Taken altogether, the visual art and textual contents create an evocative and sometimes almost cinematic vibe.

Structurally and stylistically, the works vary from one volume to another considerably, yet share a number of common elements. Among these is an emphasis on world-building and intricate depth of in-universe lore, which is often told through the form of fictional encyclopedia entries. Artificial intelligence tools excel in this type of fractal fragmentary recursive creative writing exercise, where facts are less important than invention and imagination. Many of the volumes also contain short stories or ultra-short flash fiction slice-of-life vignettes elaborating on a theme or premise.

The books retail [direct to consumers as EPUB & MOBI files](#), ranging in price from \$1.99 to \$4.99 USD. Many readers come back and purchase multiple different volumes (and in some cases dozens), as they follow their own trail through the stories contained in the books.

About the Author

The author of these works, Tim Boucher, has spent the better part of a decade working in online Trust & Safety for the likes of platforms, blockchain protocols, and non-profits. He has worked extensively in content moderation and filtering, counter-disinformation, data protection, platform policy, and product management. With regard to copyright specifically, he has also reviewed countless DMCA copyright infringement claims submitted to platforms, and built a system for managing public records of related copyright claims for blockchain-hosted files.

In addition to his creative and artistic projects, Boucher has a clear-eyed operational understanding - based on hands-on experience - of how the best intentions of technology's creators can go astray when confronted with simple human nature. His creative work and dystopian multi-modal storytelling with the books are embedded in and inseparable from the lived personal experience of having spent years handling complaints of real humans confronted with problems caused by technology.

Motivations

Whereas for other types of writing, AI's known limitations around misinformation might be a drawback, the author makes use of AI writing tools partly to exploit their tendencies to "hallucinate" non-existent or flat out wrong "facts." It is incorporated as a "feature not a bug" in this fictional context. Casting AI tools thus into the literary role of unreliable narrators in the books helps amplify the uncanniness and artificiality of the texts, as well as situate them in an old literary tradition which feels fitting given their current state of AI sophistication. The effect creates a strangely enjoyable puzzle for readers to try and solve as they piece together how the story elements fit, what it all means, and which passages might have been written by AI or by a human, and how much that really matters in our blended hyperreal future.

Tools Used

At 116 volumes, the AI Lore Books have been developed using many different AI text and image generator tools over time. It would be difficult to go back and generate a full list of every tool used, due to the hundreds of hours spent experimenting with them over nearly two years across a multitude of different services. But some of the notable ones include:

- Midjourney
- ChatGPT V 3.5 & 4
- Claude
- Dall-E V 2 & 3
- [TextSynth](#) (multiple open-source LLM models)
- Stable Diffusion (via multiple service providers)
- Character.ai
- *Many others*

The AI Lore Books also serve as a sort of historical record and commentary, documenting the state-of-the-art capacities of these tools (for the good and the bad) as viewed through the twin lenses of art and fiction at different points in the development of these models. Within a few years, what is contained within these books will look quaint and vintage in comparison as these technologies progress.

Informing Readers About Use of AI

Lost Books promotes itself to prospective readers as an “AI Publisher” and bills the books as “Illustrated AI Mini-Novels” to help set reader expectations and establish genre. Many of the books do contain a great deal of original human-written text and images.

The books individually do not list which specific AI models or services were used in their production, but they all contain a text notice on their copyright pages that they may contain elements generated by artificial intelligence. Many of the later ones also include an expanded disclaimer for greater clarity that they are also subject to human review and editing. A few of the newest books jokingly invert the need for disclaimers in the first place (and their ultimate utility), warning the potential reader that the document may include contributions from a human.

Record Keeping

With the current state of technology, it is not yet practical to effectively annotate a given text (as in an ebook format or an online article, for example) to indicate which passages were generated by a human, by an AI, or in some more blurry combination of the two. Being able as a creator to turn on (or off) this kind of meta-data would probably add a new and interesting element of analysis and enjoyment of the stories and their contents, but it does not yet exist.

We believe that development of systems like this would be empowering to readers and end users of platforms and reading apps and devices who could customize their feed or store settings based on personal preference for inclusion of human versus AI generated content and sources.

However, development of products and supporting systems like that to accurately capture at time of creation very granular provenance metadata for micro-elements of a work is going to take time and effort to build, not to mention widespread adoption across industry to make them useful. It is an effort worth pursuing.

Until such a time where much of that secondary provenance and attribution work might be reliably automated and included at a granular level within a work, there are many modes of primary artistic creation where it wouldn't be desirable (or perhaps always even possible) as an artist to have to be concerned with manually keeping line by line or image by image records of exactly how something was generated, where, when, using what prompt, etc.

To be able to create using these tools as an artist relies very much on being able to get into a “flow state” with them, so your ideas flow out of you and come to life using the tool seamlessly through a process of iterative inspiration and direction. It would turn the pleasurable expression of creating something deeply interesting and meaningful and beautiful in the moment into a kind of bureaucratic task of keeping minute elements of paperwork up to date. The complexity of assembling those records with any completeness using current technologies would make it prohibitively difficult to do so in many cases, perhaps impossible in others.

As a result, Lost Books has not retained any such records which would be easy with a reasonable amount of effort to put together as a comprehensive supporting document for the purposes of filing our works with the US Copyright Office, should we desire to do that (we are in Canada, so we will not). We imagine that we could produce for example partial transcripts from some tools, but they wouldn't easily paint the true picture of the creative work which went into it, and any such records are likely to be mingled with private personal data. The difficulty of record-keeping makes it difficult to envision obtaining copyright protection from the USCO in even the human-generated portions of the text.

Our published volumes together contain approximately 9,000 AI-generated images, and approximately 400,000 words. It is important to understand that prompt data is spread across many services over time, and this is all unstructured data in multiple different formats, much of which is unsearchable. Additionally, sometimes services one used to create something in the past shut down, or one may delete one's account because they change their policies. Each system and product works differently, and in many cases if you stop paying, you may lose access to certain features, like usage history (or you might simply have no access to it to begin with).

It is therefore highly unlikely that even a conscientious creator trying to go back after the fact and carefully document which parts of a multi-media submission to the Copyright Office were created by human or by AI and in what precise combination would be able to faithfully do so with any degree of completeness. It is consequently suggested that better more practical paths forward as to “proofs” of creativity and authorship be considered for potential copyright holders. What those might look like will be considered again later in this document.

Workflows Used

As the works extend across 116 volumes (and countless other image & text sets which were not published in ebook format), many different variations of workflows have been experimented with over time. Below are some of the more common variations of how different AI chatbot tools such as ChatGPT and Claude (the only two we work in anymore) were employed throughout the creative process.

- Brainstorm and conversationally explore a given premise or idea
- Perform basic background research on a topic (which one verifies from outside sources)
- Write lists of story ideas around a given theme or premise

- Expand an idea into a short flash fiction story with custom instructions
- Iteratively edit a piece through conversation as with a human writing partner
- Create a fictional encyclopedia entry on a given topic
- Input a long or short format human-written text as the basis for an AI-generated continuation, edit, or brainstorming session
- Perform text completions and recursively feed back in select AI-generated results and new human elements to continue a text
- Generate descriptive image prompts from a given text to use in separate image generator AIs
- Generate book titles, descriptions, and marketing copy
- Write press releases, media pitches, and other types of structured expository writing to support the works

On the image generation side, the workflow options are somewhat narrower, and chat-based image-generation options is a relatively new option, via the new ChatGPT Plus with Dall-E 3 image generator integration. We recently switched away from using Midjourney as a result of Dall-E 3's release. Example tasks we do across these various image generator systems include:

- Write in plain language instructions for a chat-based AI image generator to follow, and, based on the results, give continuing iterative refinement and direction to narrow the resulting outputs until they meet my specific requirements
- Input human generated text descriptions on the fly and explore by changing or adding to a prompt (many of those explorations become the basis for new books)
- Use quotes from an existing text as the basis for image prompts
- Use AI-generated prompts to create images
- Apply custom image parameters to image prompts, where available (as in Midjourney)
- Upscaling images to a larger size
- Upload images to use as samples or the basis for further stylistic image explorations

Concept of the Hypercanvas

In our reading of the Copyright Office's decision regarding *Zarya of the Dawn*, it appears that the Office takes a narrow view of what constitutes the "art object" within the emerging context of AI-assisted art and literary production. We would like to offer the concept of the hypercanvas, where individual generated images are themselves only brushstrokes in a larger work as a potentially more extensive alternative framework to analyze works created using these tools.

Traditionally, outside of AI-assisted media, when one looks at an individual piece of visual art, one might look at the brushstrokes on canvas, and see how together they form the finished piece. The creativity and actual labor which went into producing the work are readily apparent.

Locus of the Creative Act

However, with AI-assisted tools, the locus of the creative act and the subsequent labor which goes into its production is shifted - but by no means diminished. Instead of many individual

brushstrokes composing a work within the frame of a single physical canvas, visual or other art created using AI tools is composed from many text prompts and their graphical outputs, which iteratively create a larger meta-work of art within the latent space of AI models on what we call a “hypercanvas.”

Put simply: each prompt, each image result, and each subsequent iteration along the way constitute in a very real sense the equivalent of a brushstroke within the context of AI art.

The resulting hypercanvas work is neither restricted to nor solely contained within the frame of any single image or text output associated with it. An AI work on a so-called hypercanvas contains and extends beyond any of its individual resulting outputs. It is a multi-modal higher-dimensional exploration of the latent space made accessible by AI generators, which is then winnowed down, curated, edited, arranged, and presented to the viewer as a subset of the larger exploration. When we as artists create art or cause art to be created on our behalf in a fixed form based on our (intangible) hypercavases, we carve out only a slice of what this rich and larger multi-dimensional context contains.

New Artistic Medium

Hypercavases, as an exploration of AI’s latent spaces, could be seen as a new type of multi-modal immersive artistic medium that artists work within, and deserving of their own much deeper considerations and eventual protections as new modes of creative expression that further the arts and sciences. Just as a traditional painter works on the two dimensional canvas with paints, an AI artist navigates and creates on this larger high-dimensional hypercanvas.

The specific path taken through latent space is guided by the artist’s vision and reactions iteratively to each output from the AI. It’s a journey of aesthetic exploration and discovery which will be different for each artist who undertakes such a journey, and which is highly dependent on the creative, social, cultural, polictical, historical and other context(s) within which the artist works. The cultural impact of AI art comes, then, from how artists embed aspects of the hypercanvas explorations into specific fixed artifacts, narratives, and meanings. In this way, the hypercanvas becomes actualized in ways that speak to the human experience, and give birth to copyrightable artifacts.

Modicum of Creativity & Creative Spark

In the *Zarya* decision, it appears as a lay-person that there are three levels of potentially copyrightable works under consideration: 1) the individual images that compose the comic book (e.g., art used in the panels); 2) the text included in the book (exclusive of the individual pieces of art used in the panels); and 3) as a compilation, consisting of the “overall selection, coordination, and arrangement of the text and visual elements that make up the Work.” It’s our understanding that the text and compilation was deemed copyrightable, but not the individual art used in panels.

We believe this does not recognize the considerable creative efforts used to “paint upon the hypercanvas,” the highly iterative and intertwined nature of inputs and outputs, and the tangible work of selecting, editing, and arranging the final results into fixed manifestation(s).

To quote photographer and AI artist Phillip Toledano in a [recent interview](#) about his work:

“The funny thing about AI I’ve realized is that, in some ways, you have to think about it more consciously than you do when you’re making a photograph. For instance, if I’m making a picture with AI, I have to think about who’s in the picture. What do they look like? What are their expressions? What ethnicity are they? What’s the weather like? What’s the vantage point of the camera? What lens am I thinking about using? Is it black and white? Is the color correct for this particular era?”

We believe therefore that the minimum threshold of a “modicum of creativity” can be easily proven to have been surpassed in the context of a great deal of AI-assisted artworks.

Likewise, regarding presence of a “creative spark,” if one considers that the locus of the creative act when working in concert with AI tools has simply shifted (in some cases upstream, in others, diffusely), to being that of the “weaver” so to speak, then we see that the creative spark is still very much alive and present within the context of the hypercanvas.

Predictability of Outcomes

Regarding the [Copyright Office’s reply to Zarya](#), one possible test that seems to be proposed regarding the requisite creativity for creating a copyrighted work has to do with predictability of outcomes of generative AI tools. Quoting from the reply:

“...the process is not controlled by the user because it is not possible to predict what Midjourney will create ahead of time.”

We believe this to be an unrealistic benchmark against which to measure human creativity. One need only think, in the visual arts, of the works of creative expressionists like Jackson Pollack, whose massive canvases were covered by paint spatters which would be impossible to predict ahead of time before someone undertook the act of painting them. Similarly, one might also consider musical works of composers like John Cage, which explicitly incorporate random and spontaneous elements like rolling dice that are filled in by performers in each performance of the work. These types of works would fail the Zarya test.

Likewise, in novel-writing, for example, if one sits down and sets out to write a complete 80K word human-generated work over the course of a year, even the best planners and outliners do not happen upon all the particulars of detail, form, character, or sequence ahead of time. The “actual work” of writing consist of capturing those discoveries along the way, of painting or sculpting the larger hypercanvas in a particular creative direction.

Expanded Notions of Authorship (as in UK)

It is our understanding that the legal status of AI-generated works in the United States is different from that of the United Kingdom, which we believe to be much more favorable to innovation on the part of artists using cutting edge AI tools, as it grants certain automatic copyright protections to computer-generated works. From the [UK Intellectual Property Office](#):

“The “author” of a “computer-generated work” (CGW) is defined as “the person by whom the arrangements necessary for the creation of the work are undertaken”. Protection lasts for 50 years from the date the work is made.”

This law has been in effect in the UK since 1988, and it seems like it may be worth exploring as a way to clarify copyright for outputs for the generative AI era in a way that respects the true creativity and authorship of those who produce such works. By affirmatively declaring this in law, the UK makes the position of artists using this technology within their jurisdiction much more clear and favorable, boosting arts and innovations within the creative sector.

A [paper from the Canadian Bar Association](#) sets out similar recommendations within the Canadian context, which mirrors the UK's position somewhat:

“The complexity and collaborative nature of creating a cinematographic work compares well with the challenges posed by AI created works. For cinematographic works, the Canadian Copyright Act states that copyright subsists in the work's “maker” – which can even be a corporation. In relation to cinematographic works, the Act defines a maker as “the person by whom the arrangements necessary for the making of the work are undertaken”.

We believe that by default, barring any agreements to the contrary, the author of the work ought to be the one who undertakes the processes required for the work to come into being, and who selects the tools, and executes the decision-making processes by which the work comes into being, regardless of the medium or tools used, AI or otherwise. AI may be our paintbrush, but we're still the artist!

AI Art as Work for Hire

Regarding the Zarya opinion letter, there is a statement by the Copyright Office that when artists are using generative AI tools:

“...Prompts function closer to suggestions than orders, similar to the situation of a client who hires an artist to create an image with general directions as to its contents. [...] Absent the legal requirements for the work to qualify as a work made for hire, the author would be the visual artist who received those instructions and determined how best to express them.”

Given that within the context of generative AI tools, the “visual artist who received those instructions and determined how best to express them” is obviously not a human, the parallel breaks down, because the user expectations are decidedly different. Users are explicitly *not* hiring human artists, but are paying a technology service for outputs, often per credit or for a monthly capped usage fee. It's worth noting that most generative AI companies currently in their policies do not claim ownership of outputs - leaving out and open questions of ultimate copyrightability in their user agreements.

Users who are paying for a system to produce AI outputs, barring any agreement or restriction to the contrary, should reasonably expect some stake in ownership of those results. The exact nature of and amount of that stake in ownership should be more clearly and transparently expressed to end users of services, including as to whether the outputs are even copyrightable at all.

Our expectation and ground assumption (recognizing the law is different federally in Canada, and provincially in Quebec where we produce our works) is absolutely that the images and texts which I cause to be created through AI generators are all owned by me (individually and in aggregate), unless otherwise explicitly stated to not be the case.

The US Copyright Office might consider requiring AI generator services make clear in their user agreements that the resulting outputs are not copyrightable within the United States, if that is determined to be so, since commercial use of outputs and productivity is a big part of the value draw of these tools. However, it must also be considered whether making all AI-generated or assisted outputs uncopyrightable might unintentionally serve to inhibit the progress of the arts and sciences in the US. It seems in both the interests of the AI services providers, and end users that the resulting outputs may be copyrightable, provided the criteria are made more clear and predictable for all parties as to how ownership and authorship are assigned.

Updating the Substantial Similarity Test

It is our understanding that ideas or concepts such as, for example, “dog on a skateboard,” are not in themselves copyrightable, only specific fixed expressions of those ideas that meet other parameters set out in the law (idea-expression dichotomy).

One set of present concerns in AI-generated art seems to stem from the relative ease of creating works via generative AI, and that these speed and scale might ultimately endanger the ability for others to create similar works (other dogs on skateboards, so to speak) due to risks or uncertainties around substantial similarity and potential infringements in works that include AI elements.

As [Lemley argues here](#), substantial similarity tests need to be updated for the AI era. Especially since it is not always possible to determine whether or not a potential infringer had access to the original in order to make allegedly infringing copies.

Thin Copyright for Works with Low Human-Involvement

One approach might be developing a framework for certain categories of AI-assisted or generated works which effectively narrows copyrightability for those works (primarily those which might be considered low involvement by a human creator). It is our understanding that in the case of two photographers who photograph the same underlying real object, substantial similarity has a much narrower end utility, and legal outcomes are restricted to protecting against nearly identical copies.

It seems like considering narrower applications for copyrights of works involving some types of low human-involvement AI-generated elements might be a way to allay some fears about overbroad applications of the similarity test which would overly restrict other authors' use of these concepts. The Copyright Office might consider restricting copy protections for a class of AI elements formally to a "thin" or narrow scope against nearly identical copies of AI-generated or assisted works.

Fair Use in Training Data

We believe in the importance of being able to mutually build on human knowledge and creativity for the betterment of the lives of all peoples. While copyright should protect the ability for people to be rewarded for their works, we should be careful not to unduly hinder the free flow of information and development of new technologies as a byproduct. As stated in the [Artists Using Generative AI - Submission to Copyright Office](#):

"Copyright law should continue to leave room for people to study and analyze existing works in order to craft new ones, including through the use of automated means like those used to create AI models."

We (artists, not lawyers) believe in good faith that under US law, including copyrighted works in AI training sets constitutes Fair Use, and is not infringing. The purpose of including items in AI data sets is not to copy or store them for retrieval. Its aim is to analyze, measure, and compare their properties in aggregate in order to transformatively create new works which are not merely derivative of works in the training data but entirely new creations with new meaning and message. We believe the Copyright Office ought to affirmatively clarify the same in order to dispel legal confusion.

Leaving questions of Fair Use in AI data training sets up to numerous court cases seems likely to yield scattershot inconsistent decisions that will ultimately create a lot of confusion and risk for people involved with developing and using these services. In this regard, Japan's approach to declaring that they [will not enforce copyrights against AI training data](#) is an interesting one. Whether or not this approach translates to US law and cultural values, a clarification would provide a measure of legal risk reduction for diverse groups making use of these technologies.

Opt-In Data Sets for Non-Public Works

For non-publicly available works which might not fall under Fair Use, we support the development of high-quality sustainable AI training data sets that are entirely opt-in, and which fairly compensate contributors at agreed-upon rates for use of their works, where appropriate. Contributors might include not only creators of copyrightable works, but also unseen participants like content moderators, trainers, and others who play crucial roles in collecting, cleaning, and screening included data. We believe that creators and the many other invisible workers affected by these technologies should always be consulted to find the best paths forward.

Alternative AI Options

We strongly support free access for all people to all human knowledge, and firmly believe the notion that ideas freely shared grow stronger and more resilient, giving birth to new and better ones more suited to the times, and that this is an unending process in which all humans - not just content creators - participate in all the time, whether or not they use AI. We need to cherish and protect that millenia old flow and not let copyright unduly restrict it, or allow corporate interests to be the dominant and driving force and value decider behind all human interactions.

We believe there is a strong benefit to having many different types of AI technologies available to the general public and for business purposes. Each system should have diverse methods and capabilities, inputs and outputs. In some cases, certain uses of generative AI technology will need to be able to show a chain of licensing and provenance of information. To serve those needs, having known and well-vetted data sets available as training described will be highly desirable.

At the same time, there is a very real risk that due to regulation, AI technology will become increasingly controlled by the few large corporations who can afford compliance programs, and who implement excessive "safety" measures without any public oversight or accountability. We need to take strong steps now to ensure the long-term viability of alternative, open, and public options for transparently training and developing AI services in ways that are still respectful of human rights. The standards we deploy in these areas should not be so difficult and expensive that smaller players should be denied access to the markets, and their innovations stifled, nor should they shut down the free flow of human knowledge that mutually enriches all our lives.

Thank You for Reading!

We thank the US Copyright Office, as well as Innovation, Science and Economic Development Canada for their time and are happy to participate in further discussions to imagine new possibilities for copyright in this new era of generative AI.