

# Don't train your model on my novel: AI refusal statements

**Daniel G. Brown**

David R. Cheriton School  
of Computer Science  
University of Waterloo  
dan.brown@uwaterloo.ca

**Carolyn E. Lamb**

School of Computing  
Queen's University  
cel4@queensu.ca

**Lauren R. Byl**

Library  
University of Waterloo  
lrbyl@uwaterloo.ca

## Abstract

We describe the phenomenon of statements forbidding generative AI (genAI) training on copyright pages of novels. These AI refusal statements typically claim that the book's text was not made with genAI tools and forbid use of the book to train genAI models. A sizeable minority of recent books, across many genres, and in both traditionally published and self-published works, have these statements. While probably lacking legal force, they show authors' motivation to keep their works out of AI training and give a space for collective action. We bring these statements to the computational creativity community to reinforce our community's ethical standards and enable better collaboration with creative humans.

## Introduction

One consequence of the rise of generative AI (genAI) is its effect on creative industries. Creative workers in many fields, including creative writing, music, animation, and visual art, have lost work (e.g., (CVL Economics 2024)), pay (Brunotts 2024), and professional identity (Lamb, Brown, and Grossman 2024) due to genAI. Their work has been used (typically without permission) to train models replacing their labour. This shock has led many creators to engage in strategies, both individual and collective, to protect themselves from genAI's harmful effects (The Authors Guild 2023; Lamb, Brown, and Grossman 2024). One strategy, encouraged by the Authors Guild, is the creation of AI-related statements on the copyright pages of books, claiming to prohibit the use of the book for genAI training.

In this short paper, we explore these AI training refusal statements. We investigate their frequency in a variety of popular fiction genres and reflect on their legal validity, scope, and purpose. While unlikely legally enforceable, AI training refusal statements express authors' and publishers' strongly held views about the use of their work and are part of a broader trend of authors asserting ownership over their ideas and products. We bring them to the attention of the computational creativity community as evidence CC researchers must engage seriously with creators whose work we use as data or inspiration.

## Copyright and AI-generated fiction

The status of generative AI with regards to copyright law remains controversial in the US and Canada, and around the

world. This paper focuses on American and Canadian law. Copyright provides for authors to profit from their own creative work. In the United States and Canada, newly written fiction is protected by copyright until seventy years after the author's death, so an author who dies in 2025 will have their works lose copyright on January 1, 2096. In theory, copyright allows authors (and their representatives) to use force of law against third parties who sell or distribute work without authors receiving payment. Famously, the creators of genAI large language models (LLMs) have trained them on large sets of pirated books such as the Books3 dataset (Reisner 2023) and the pirate website LibGen (Knibbs 2025).

Copyright laws in the US and Canada contain exceptions allowing use in some circumstances without permission. Fair use (US) and fair dealing (Canada) allow copyrighted materials to be used in situations such as education, research, or critique, without license. GenAI companies have claimed their training on pirated books is fair use (Belanger 2025), and that their models do not directly compete with materials they are trained on. In the first AI and copyright case decided in the US, Thomspon Reuters successfully argued against this claim for AI-based legal information (Kaufman 2025).

Authors and publishers clearly disagree with model creators. There are 27 copyright related legal challenges against AI providers in the US (Knibbs 2024) (updated 2025); see also (Weisenberger, Milton, and Enright nd). In Canada, there are at least two copyright and AI cases ongoing (CanLII v Caseway AI; Canadian Media Companies v OpenAI). Sam Altman (CEO of OpenAI) lobbied the US federal government to declare AI training from copyrighted inputs fair use (Belanger 2025), Microsoft has similarly lobbied the Canadian federal government for clarity on copyright's interaction with AI services (Microsoft Corporation 2024).

If frameworks are established to allow training on copyrighted work, they might come with consent mechanisms (opt-in or opt-out) or compensation (The Authors Guild 2023). Some publishers of news, non-fiction, and scholarly works have brokered legal deals with genAI companies, allowing for training (The Atlantic 2024; Kwon 2024).

## Copyright pages

The copyright page of a book, whether traditionally published or self-published, typically contains the name of the

copyright holder and publisher, year of publication, ISBN, and miscellaneous legal assertions, such as the disclaimer that the book is fiction and that resemblances to real people or events are coincidental, or a statement banning unauthorized reproduction or uploading.

Under current US and Canadian copyright law, every human-authored work of creative writing is copyright protected on creation, unless it is specifically released by the creator into the public domain. Copyright can be registered in each country with similar benefits, chiefly presumption of the author's ownership. A copyright page is not how a book gains copyright protection, but in both the US and Canada, notices serve a legal purpose. In the US, "a court will not give weight to a defendant's use of an innocent infringement defense. . . if . . . the work in question had a proper copyright notice." (Federal Register 2017). In Canada, the author's name appearing in the work can serve as proof of authorship, in absence of copyright registration (Copyright Act, s34[1-2]). In the US, 17 USC § 401 lists three elements required for this legal protection: the copyright symbol, first year of publication, and the name of the owner of the work.

Further to these legal purposes, publishers and authors often use copyright notices to inform readers how they can use the work. These educational components are not specified in legislation. Mazzone (2006) and Poor (2008) find publishers often use these notices to enforce a maximalist point of view and that these notices are often inaccurate.

## AI refusal statements

Within this existing context of AI training controversy and copyright pages, we have noticed a trend in the past year of fiction authors including a statement about AI on their copyright page. A short example of these AI statements is from Mila Finelli's "Empire of Temptation":

No AI has been used in the creation of this book, nor may it be used to train AI.

In self-published books, these statements are typically written by the author. In traditionally published books, if they appear, they are typically provided by the publisher and appear in the same form for all that publisher's books.

These statements make one or more claims about the book. The two most common, both found in Ms. Finelli's, are that the book was not written using genAI, and that the author refuses for the book to be used to train genAI. Authors may make broader claims: such as that the book may not be used for any other kind of model training, or about their own creative process, or about the importance of their work as fiction writers. As a class, we call these *AI refusal statements*. When they specifically ban use of the work for AI training, we call them *AI training refusal statements*.

Using an AI training refusal statement on a book's copyright page is recommended by the Author's Guild (2023), an advocacy organization for writers, to "opt out" and prevent genAI companies from training using a book. We examine a representative sample of AI training refusal statements that have appeared in recent books and discuss their scope and efficacy. Specifically, our research questions are:

1. How common are these statements? Are they more common in self-published books than traditionally published books, or in marginalized literature forms (like queer romance or disability fiction) than more mainstream books?
2. What is the scope of these statements?
3. Do they make a meaningful legal claim?
4. Do they serve a protective function for authors, rather than a legal function?

## Data

Between December 25, 2024, and January 6, 2025, we examined the top 50 Amazon Canada bestsellers for ten mainstream genres of fiction (mystery, romance, science fiction, fantasy, graphic novels, thriller/suspense, historical fiction, horror, literary fiction and young adult) and five marginalized genres (queer romance, erotica, disability fiction, Black literature and Canadian Indigenous literature). We also examined the top 10 "hot new releases" for each genre.

From this initial data set of 900 books (15 genres  $\times$  (50 + 10) per genre) we removed books published before January 2024 (many current best-sellers were published long before the rise of genAI). We removed duplicates, which were common since many books are listed in more than one genre or format. We removed books not in the English language, books that are not creative (like blank books or calendars), collections of multiple books listed as a single entry, and short-story collections. This resulted in a filtered set of 195 best-sellers and 81 "hot new releases;" removing duplicates between these two sets resulted in a total set of 246 books.

Amazon's classification of books into genres is inexact. Many books did not appear, at a glance, to fit assigned genres; "literary fiction" included Elise Camden's "His Wife's Secretary", which is erotica, while "Canadian Indigenous literature" included non-fiction books about Indigenous experience and polemics disavowing Indigenous mistreatment. We did not remove books due to genre classification.

For each remaining book, we used a variety of sources to read its copyright page: Amazon Previews, Kindle Unlimited, Google Books, Amazon Prime, library e-books, and visits to physical bookstores and libraries. We viewed copyright pages of all but 8 best-sellers, and all but 4 "hot new releases", so we view our data as representative.

## Analysis

**Question 1: How common are these statements?** We found AI training refusal statements in 13 of 77 "hot new releases" (17%), and in 29 of 187 bestsellers (15%). They appeared in 17 of 103 traditionally published books (17%), and in 21 of 131 self-published books (16%). We did not find statistically distinguishable differences in the prevalence of AI refusal statements between books of different publication types. Although marginalized authors are at greater self-perceived risk from the use of genAI (Lamb et al.), we did not conclusively find that there were more AI refusal statements in marginalized genres. Queer romance (with 7 in 31 books, or 23%) and erotica (with 5 in 21 books, or 24%) are the most frequent genres to have AI training refusal statements, but the overall results are not significant,

and other marginalized genres like Black literature don't have any AI training refusal statements in the sample. This may be partly due to the low numbers of books remaining in each genre after filtration. AI refusal statements appear in 12 of the 15 genres; the only ones that don't have any are Black fiction, historical fiction, and horror.

Overall, a substantial minority of books in most genres, whether self-published or traditionally published, have adopted AI refusal statements. This is a trend extending broadly across many facets of the publishing world.

**Question 2: What is the scope of these statements?** We assume that some of these authors added AI refusal statements following the advice of the Author's Guild, but we did not find any books in our data that closely recreated the exact wording provided by the Authors' Guild. Instead, authors and publishers write these statements in their own words.

In terms of scope, authors' AI refusal statements take several forms. First, authors may state that their work is not produced by genAI. Second, authors may state that their work must not be used to train genAI, perhaps offering a legal claim that such training would be unlawful. Third, authors may include additional refusals, forbidding various other uses of their text besides genAI training. Finally, authors may include additional statements about their values or opinions about genAI. The examples discussed in the remainder of this section illustrate each of these variations.

From Mila Finelli's "Empire of Temptation": "No AI has been used in the creation of this book, nor may it be used to train AI." This is perhaps the simplest example of an AI refusal statement, satisfying the first two forms described.

From Jagger Cole's "Emperor of Lust": "The unauthorized reproduction, transmission, distribution of, or use of this copyrighted work in the training of AI is illegal and a violation of US copyright law." This statement makes a legal claim about AI training.

From Elliott Rose's "Chasing the Wild": "No part of this work was written with artificial intelligence. I support human creativity, not generative forms of AI. No part of this work may be used to create, feed or refine artificial intelligence models, for any purpose." This statement satisfies many forms listed above, and includes a values statement.

From Dav Pilkey's "Dog Man: Big Jim Begins": "No part of this publication may be reproduced, transmitted, downloaded, decompiled, reverse engineered, used to train any artificial intelligence technologies . . . now known or hereafter invented, without the express written permission of the publisher." This book, published by Scholastic, rejects many different technologies, and shows that a major publisher is using AI refusal statements.

From Shannon Mae's "How to Hack a Hellhound": "NO AI/NO BOT. We do not consent to any Artificial Intelligence (AI), generative AI, large language model, machine learning, chatbot, or other automated analysis, generative process, or replication program to reproduce, remix, summarize, or otherwise replicate any part of this creative work, via any means: print, graphic, sculpture, multimedia, audio, or other medium. We support the right of humans to control their artistic works." This long statement forbids style

transfer of a work to new media formats, a common CC task (Horn et al. 2015). It also makes a cultural statement (using the first person plural) about the importance of creators' control over their work.

We first became aware of AI refusal statements via Layla Dorine's "A Daddy for Christmas 2: River", published in December 2024. This book is not in our data sample. This book uses the same AI refusal statement as Mae (substituting "the author" for "we"), but adds, "No part of this book has been created using AI-generated images or narrative, as known by the author, who sources her cover designs from a real human digital artist who pours his soul into the things he creates." The use of "soul" is interesting, though it is not found in any of the AI refusal statements in our data sample. Previous studies of responses to AI by fiction authors have discussed a similar concept, which is associated with human writers' conscious intentionality in communicating affective experience to a reader (Lamb, Brown, and Grossman 2024). We highlight Dorine's refusal statement because it references this quality.

**Question 3: Are AI training refusal statements legally meaningful?** An obvious question is if AI training refusal statements have legal validity. Can an author restrict readers more than copyright law allows by adding a note to a copyright page?

They would probably not have any greater effect than a plain copyright statement. As Stim states, "Placing a copyright notice . . . is like sticking a "Keep off the grass" sign on your front lawn. People who respect your property rights will stay off, and those who don't . . . you'll just have to chase them away" (2023). Some AI training refusal statements directly claim that training genAI models on unlicensed copyright data is unlawful, but this is unresolved. Fair dealing and fair use exceptions are applied on a case-by-case basis, and AI training can take on several different forms, so it may be that use of a work for AI training is lawful in some cases and not in others. That said, Mazzone (2006) and Poor (2008) argue that these kinds of copyright notices can have a chilling effect on end users.

The Author's Guild guidance around AI training refusal statements (2023) makes misleading claims about their legal efficacy. First, that "copyright law gives authors exclusive rights to decide if others can use their works"; it does, but there are exceptions, and genAI companies are leaning heavily on these exceptions in their legal defense. Second, that genAI companies will "observe opt-outs when pressed." They generally haven't (Belanger 2025). Perhaps more importantly, these statements can't help at all if AI companies are using pirated sources for training (Knibbs 2025).

Once broader legal questions around genAI training are settled, a regime of licensing and payments, possibly including opt-in or opt-out models, might be enacted, with AI training refusal statements demonstrating opting out, but this would depend on the details of the regulations.

One interesting exception to this legal limbo is the European Union's 2019 Digital Single Market Directive, a directive from the pre-genAI era which allows for text and data mining unless a rightsholder explicitly opts out. One book

in our sample – the mystery novel “Echoes on the Fens” by Joy Ellis – explicitly references the DSMD:

No part of this book may be used or reproduced in any manner for the purpose of training artificial intelligence technologies or systems. In accordance with Article 4(3) of the Digital Single Market Directive 2019/790, Joffe Books expressly reserves this work from the text and data mining exception.”

Referencing the DSMD is resourceful, but its legal enforceability is questionable, as “Echoes on the Fens” was published in the United Kingdom, where the DSMD does not apply due to Brexit (Simmons and Simmons 2021).

**Question 4: Do AI training refusal statements serve a protective function?** If AI training refusal statements are not directly legally enforceable, they might protect or benefit authors indirectly. Lamb *et al.* (2024) show that fiction authors are engaged in informal strategies, personal and collective, to defend themselves and their field from generative AI. An AI training refusal statement can serve these goals by allowing authors to identify their values (against AI training on their work); it might also establish a standard of behavior (“I take a stand against AI training, and you should too.”).

In our data, many copyright statements do explicitly identify an author or publisher’s values. For example, Hachette does not put AI refusal statements on its books, but includes the note, “Hachette Book Group supports the right to free expression and the value of copyright. The purpose of copyright is to encourage writers and artists to produce the creative works that enrich our culture.” This is a clear statement of the publisher’s values, possibly to encourage writers to choose to publish with them because of their stance. Similarly, the Authors Guild recommendation that authors include AI refusal statements in their copyright pages can allow authors to demonstrate a collective will to be heard on this issue, even if they are individually unenforceable.

## Discussion

### Legality vs morality

While some problems posed by the rise of genAI are novel, the concept of an author’s legal and moral ownership of their work – and, consequently, their right to dictate what readers can and can’t do with it – is not. We find loose analogues to several other existing situations in art and publishing.

One analogous situation is the controversy in publishing over cultural appropriation – the practice of a white author, or an author from a majority culture, incorporating minority cultural products into their work. What constitutes cultural appropriation, and whether and to what extent it should ever be allowed, is a topic of intense debate among authors (Young 2010). However, the idea underlying cultural appropriation is that members of a culture, especially a minoritized or colonized culture, “own” their lived cultural experience and that this experience should not be taken outside its original context and used for profit by someone without such legitimate ownership.

Legally speaking, it is not possible to copyright a lived experience. However, the more typical argument against cul-

tural appropriation is not legal but moral: white and majority culture authors should not culturally appropriate.

Similarly, authors have been publicly shamed for writing stories that copy too closely from someone else’s lived experience, even if there is no racial or cultural difference involved (Nowicki 2021).

Moral arguments around appropriation, and around uses of text “owned” by another person, are also common in social media discourse. For instance, in the 2010s, some Black activists on Twitter argued that non-Black participants should not participate in online discussions among Black people, nor should they comment on their tweets when retweeting or use their content in journalism (Freelon *et al.* 2018). Nothing in the Twitter terms of service required this behaviour; the goal was to keep the primary Black Twitter conversation among Black people, and informal community standards were created to uphold this expectation.

These examples show that even if there is not a strict legal prohibition against using an author’s published words in a certain way, community standards can still be created which serve to discourage such use on moral grounds. While the legal question of fair use/fair dealing is important, the fact that something is or may be fair at law does not resolve the question of whether it is moral.

### How should CC researchers respond to AI training refusal statements?

Most computational creativity research in North America is subject to the research fair use/fair dealing exception; exceptions might arise when researchers build systems whose outputs can economically compete with their copyrighted training inputs. (See Brown *et al.* (2021) for more analysis of the Canadian context. Far more has been written about the US situation; see Lemley and Casey (2021) for a particularly strong statement.) However, despite likely not being legally enforceable, AI training refusal statements make authors’ and publishers’ wishes clear regarding how their work is and is not to be used. It therefore presents ethical difficulties for CC researchers to ignore these statements, which are grounded in authors’ reasonable concerns.

CC researchers can use AI training refusal preferences to guide respectful data gathering and to avoid performing research on the work of authors who have gone out of their way to opt out of such research. There is a long-standing tradition of ethical research in CC (see, *e.g.*, (Pérez y Pérez and Ackerman 2020)), and this new practice enables us to take a stand in favour of creators having autonomy with the uses of their own work.

Furthermore, the presence of AI training refusal statements is itself a possible site of future study. It would be useful, in future work, to correspond with authors who have not provided such statements (especially in self-publishing, where the choice to include one is up to the author) and discern in more detail what differences of attitude, knowledgeability, or experience underly this choice. It will also be interesting to see how common AI training refusal statements are as genAI continues to develop and as the legal questions around the use of copyrighted work for AI training become more settled.

## Acknowledgments

The work of D.G.B. is supported by a Discovery Grant from the Natural Sciences and Engineering Research Council of Canada. We appreciate helpful discussions with Jordan Hale and Su Penn.

## References

- Belanger, A. 2025. OpenAI declares AI race 'over' if training on copyrighted works isn't fair use. *Ars Technica*.
- Brown, D.; Byl, L.; and Grossman, M. 2021. Are machine learning corpora 'fair dealing' under Canadian law? In *Proceedings of ICCA*, 158–162.
- Brunotts, K. 2024. AI will not replace artists. It will devalue them. *Skeptic*.
- CVL Economics. 2024. Future unscripted: the impact of generative artificial intelligence on entertainment industry jobs. Technical report.
- Federal Register. 2017. Affixation and position of copyright notice. 82 F.R. 42735.
- Freelon, D.; Lopez, L.; Cleark, M.; and Jackson, S. 2018. How Black Twitter and other social media communities interact with mainstream news. Technical report, Knight Foundation.
- Horn, B.; Smith, G.; Masri, R.; and Stone, J. 2015. Visual information vases: Towards a framework for transmedia creative inspiration. In *Proceedings of ICCA*, 182–188.
- Kaufman, R. 2025. Copyright's big win in the first decided US artificial intelligence case. *The Scholarly Kitchen*.
- Knibbs, K. 2024. Every AI copyright lawsuit in the US, visualized. *Wired*.
- Knibbs, K. 2025. Meta secretly trained its AI on a notorious piracy database, newly unredacted court docs reveal. *Wired*.
- Kwon, D. 2024. Publishers are selling papers to train AIs – and making millions of dollars. *Nature* 636:529–530.
- Lamb, C. E.; Brown, D. G.; and Grossman, M. R. 2024. Precarity and solidarity: Preliminary results on a study of queer and disabled fiction writers' experiences with generative AI. arXiv preprint 2412.04575.
- Lemley, M., and Casey, B. 2021. Fair learning. *Texas Law Review* 99:743–785.
- Mazzone, J. 2006. Copyfraud. *New York University Law Review* 81:1027–1100.
- Microsoft Corporation. 2024. Consultation on copyright in the age of generative artificial intelligence. Submission to Government of Canada.
- Nowicki, A. 2021. "Cat Person" and me. *Slate*.
- Pérez y Pérez, R., and Ackerman, M. 2020. Towards a methodology for field work in computational creativity. *New Gen. Comput.* 38(4):713–737.
- Poor, N. 2008. Copyright notices in traditional and new media journals: Lies, damned lies, and copyright notices. *Journal of Computer-Mediated Communication* 14:101–126.
- Reisner, A. 2023. Revealed: The authors whose pirated books are powering generative AI. *The Atlantic*.
- Simmons and Simmons. 2021. The DSM directive two years on: is it in force? <https://www.simmons-simmons.com/en/publications/ckctct3gfi1t1s0b02r8s9vgfb/the-dsm-directive-two-years-on-is-it-in-force->.
- Stim, R. 2023. Dear Rich: An intellectual property blog. does a copyright notice protect a website? <https://dearrichblog.blogspot.com/2023/09/does-copyright-notice-protect-website.html>.
- The Atlantic. 2024. The Atlantic announces product and content partnership with OpenAI. *The Atlantic*.
- The Authors Guild. 2023. Practical tips for authors to protect their works from AI use.
- Weisenberger, T.; Milton, D.; and Enright, H. n.d. Case tracker: Artificial intelligence, copyrights and class actions. Technical report, BakerHostetler.
- Young, J. 2010. *Cultural appropriation and the arts*. John Wiley and Sons.