



# Creative Writers' Attitudes on Writing as Training Data for Large Language Models

Katy Ilonka Gero  
Department of Computer Science  
Harvard University  
Cambridge, Massachusetts, USA  
katy@g.harvard.edu

Meera Desai  
School of Information  
University of Michigan  
Ann Arbor, Michigan, USA  
madesai@umich.edu

Carly Schnitzler  
University Writing Program  
Johns Hopkins University  
Baltimore, Maryland, USA  
cschnit1@jh.edu

Nayun Eom  
Department of Sociology  
Harvard University  
Cambridge, Massachusetts, USA  
neom@fas.harvard.edu

Jack Cushman  
Harvard Library Innovation Lab  
Harvard University  
Cambridge, Massachusetts, USA  
jcushman@law.harvard.edu

Elena L. Glassman  
SEAS  
Harvard University  
Allston, Massachusetts, USA  
glassman@seas.harvard.edu

## Abstract

The use of creative writing as training data for large language models (LLMs) is highly contentious and many writers have expressed outrage at the use of their work without consent or compensation. In this paper, we seek to understand how creative writers reason about the real or hypothetical use of their writing as training data. We interviewed 33 writers with variation across genre, method of publishing, degree of professionalization, and attitudes toward and engagement with LLMs. We report on core principles that writers express (support of the creative chain, respect for writers and writing, and the human element of creativity) and how these principles can be at odds with their realistic expectations of the world (a lack of control, industry-scale impacts, and interpretation of scale). Collectively these findings demonstrate that writers have a nuanced understanding of LLMs and are more concerned with power imbalances than the technology itself.

## CCS Concepts

• **Human-centered computing** → **Empirical studies in HCI; HCI theory, concepts and models**; • **Computing methodologies** → **Natural language generation**.

## Keywords

Large language models, natural language generation, creative writers, creative writing, writing assistants, data collection, training data, archival practices, grounded theory.

## ACM Reference Format:

Katy Ilonka Gero, Meera Desai, Carly Schnitzler, Nayun Eom, Jack Cushman, and Elena L. Glassman. 2025. Creative Writers' Attitudes on Writing as Training Data for Large Language Models. In *CHI Conference on Human*

*Factors in Computing Systems (CHI '25)*, April 26–May 01, 2025, Yokohama, Japan. ACM, New York, NY, USA, 16 pages. <https://doi.org/10.1145/3706598.3713287>

## 1 Introduction

As large language models have gained wide-spread usership, so have investigations into the text used to train them [10, 17, 27, 33, 37, 61, 65]. There has been a particularly swift and negative response to the use of creative writing, such as novels, memoirs, and longform journalistic books, as training data, as the authors of this writing express outrage at the use of their work without consent or compensation [9]. Such frustration has led to dozens of lawsuits against technology companies like OpenAI, Microsoft, and Google, alleging that this use of writing goes against U.S. copyright law [1].

Not everyone agrees that the use of creative writing for training data is illegal, or even immoral. In comments to the U.S. copyright office, The Author's Alliance, whose mission is to advance the interests of authors who write for the public benefit, argues that using writing as training data, in the majority of cases, is protected by the copyright doctrine of fair use, and that licensing schemes for training data are both logistically infeasible and poor policy [6]. Other writers have expressed distaste, but not concern, for the use of their writing in this new and unexpected way [12].

In this work, we focus on understanding the response of creative writers such as novelists, memoirists, poets, and fan fiction writers, all of whom engage in book-length writing that has been used as LLM training data without their consent. While focusing on creative writing may seem niche in the larger world of textual data collection, book-length writing has been found experimentally to be among the most highly valuable training data for LLMs [63]. And, to the best of our knowledge, while there have been a series of licensing agreements between technology corporations and news publications (e.g., [71]) there is no LLM available that has been trained on the consensually collected writing of creative writers. Information scholars have argued that web-scale data collection practices challenge traditional ethical principles governing the use of human subjects data, and that data should be collected and used in accordance with data subjects' perspectives [83], necessitating

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

CHI '25, Yokohama, Japan

© 2025 Copyright held by the owner/author(s). Publication rights licensed to ACM.  
ACM ISBN 979-8-4007-1394-1/25/04  
<https://doi.org/10.1145/3706598.3713287>

research into the perspectives of those creating valuable data. Additionally, as the web becomes increasingly hostile to scrapers [62] and content is locked down by AI companies [7], the need to understand what motivates—or demotivates—people to participate in consensual sharing becomes important to permitting an open web.

With all this playing out in the public and legal sphere, we see a place for researchers to characterize the concerns of creative writers. We do not investigate these issues from a legal perspective, but rather an ethical and human-centered one, centering the voices of creative writers. In this work, we ask:

- RQ1: How do creative writers reason about the real or hypothetical use of their writing as training data?
- RQ2: Under what conditions, if any, would they consent to their writing being used?

To answer these questions, we conducted a grounded theory investigation by interviewing 33 creative writers. We sought to interview a wide variety of authors, in terms of genre, delivery method (i.e., method of sharing or publishing their work), degree of professionalization, demographics, and engagement with AI in their writing practice. We did not presuppose that writers should, or should not, want their writing used as training data. Instead, we sought to understand how they are reasoning about this issue and the concerns that they raise.

We found that writers reason about the use of their writing as training data according to three interconnected principles. 1) Writers understand that they learn from and influence each other via **the creative chain**, and question what role LLMs play in this process. 2) Writers desire **respect** for their work as time-consuming labor that requires expertise. 3) Writers note that **the human element** of writing, particularly embodiment and emotion, means that AI cannot replicate writing processes.

These principles were often at odds with writers' realistic expectations of the world they live in. 1) Writers experience a **lack of control** over how their writing is used by technology companies, and how LLMs that may be trained on their writing are used more generally. 2) Writers struggle to make sense of their contribution to LLMs given their **interpretation of scale**, when any individual's contribution is so small. 3) Writers have both current and predicted concerns about how this technology will have **industry impacts**, some of which supersede personal impacts.

In the discussion, we consider how writing changes when it becomes data, and propose a colonial lens to understand paths forward. We also discuss the place of LLMs in the **creative chain**, and how they may (or may not) be like libraries. Finally, we outline what kind of research might best support writing communities.

## 2 Background

In 2023, the Writer's Guild of America, representing screenwriters, and the Screen Actors Guild, representing television and radio artists, went on strike over labor disputes with the Alliance of Motion Picture and Television Producers [3, 22]. One of the main focus points of both strikes was the use of generative AI. Although neither directly involved literary writers, it put concerns about the use of AI to displace or devalue artists into the public arena. Then, in August, 2023, The Atlantic published a story about the use of Books3, a corpus of pirated books, to train LLMs [77]. This

precipitated outrage from creative writers at the use of their work without their consent and without any compensation [9]. At the end of 2023, the professional organization The Author's Guild ran a survey with 2,400 of their members, and "found nearly universal opposition among authors to their works being used to train AI systems without permission" [2].

As writers express their discontent, concerns about the legality of using published books as training data are being raised. The U.S. Copyright office started a study in 2023 regarding copyright issues raised by AI, including several requests for public comment [72]. Dozens of lawsuits against technology corporations were filed, disputing the use of published books as training data, for the most part due to violations of copyright [1]. Such investigations and disputes are ongoing. Legal questions will be resolved by the courts, and likely differ across jurisdictions and nations. In our work, we seek to uncover not strong legal arguments, but rather the humanistic concerns and reasoning of writers.

## 3 Related Work

### 3.1 AI Writing Assistance

Computational processes have a long history in creative writing. Early explorations include the surrealist Dadaism movement and the French Oulipo group, both of which engaged in algorithmic processes for writing. A 2024 review of all kinds of writing assistants found that, in the HCI and NLP research communities, papers on writing assistants have been growing dramatically since the mid-2010s [56]. Recently, much of this work investigates how LLMs, particularly large-scale corporate LLMs like GPT-4, might contribute to particular creative writing genres, like playwriting [68], storytelling [44], and cartoon-captioning [50], or to particular parts of the creative writing process, like character development [76], worldbuilding [20], and revision [11]. Such work demonstrates the capabilities of LLMs to contribute to creative writing, but does little to address why such a system may or may not be adopted by writers. Nor does it address the documented negative attitude that many creative writers have towards LLMs more generally.

Another line of research investigates the social and psychological issues that AI writing assistance brings forth. For instance, research has found that increased amounts of AI generated text decreases writers' sense of ownership [26, 57]. Sense of ownership may be modulated by sense of contribution [52], the style of the AI generated text [48], and writers' engagement with AI contributions [41]. Taken together, these findings almost unanimously show that, on average, AI-supported writing decreases but does not eliminate writer's feelings of ownership, underscoring the need for a larger theory of AI participation in the creative process.

### 3.2 Ethical Considerations of Data Usage

Prior to the current interest in generative AI, researchers investigated the ethics of using social media data in HCI and the computational social sciences. Like LLM training data, social media data repurposes people's data in ways they may not have expected and at a scale that challenges traditional principles governing the ethical use of data involving human subjects. Informed consent, for example, is often infeasible on the scale of social media datasets. Instead, information scholars argue that data should be collected

and used in accordance with data subjects' perspectives on the acceptable use of their data [29, 40, 83].

Research on social media data suggests that data subjects' level of comfort with their data's use depends on contextual factors such as the sensitivity of the data, who the researchers are, or what the purpose of the data analysis is [29, 35, 36, 40]. For example, Klassen et al. (2022) investigate the ethics of research on Twitter data posted by Black people, finding that several participants expressed that White people doing research on specifically Black Twitter seemed colonizing, whereas the study of Black Twitter was not inherently problematic on its own. These findings suggest that writers may also be concerned about the intentions of the people collecting training data and training LLMs, as opposed to being against the use of their work as training data in all cases.

Recently, work across a variety of Computer Science subdisciplines has considered the ethical implications of existing training datasets for language models. Research has found that training datasets for popular, commercial models contains copyrighted material [10, 17], contains potentially problematic content such as toxic or pornographic content [10, 33], and has a skewed religious representation [10]. Others have investigated the representation of text from and about minoritized individuals in pretraining datasets [27, 37, 74], finding skewed representation of minoritized groups in these datasets. Researchers have argued that skewed representation in an LLM training dataset may result in disproportionate harm to minoritized groups if these models are used in high stakes contexts [37, 74]. However, little academic work has investigated how textual data creators themselves feel and reason about the use of their writing as training data.

Just as writers are grappling with the use of their writing as training data, visual artists too are expressing concern, distaste, and often outrage at the use of visual art as training data for generative AI, specifically text-to-image models [46]. Artists have launched a number of their own lawsuits against technology companies [1], often analogous to those from the writing community. Research has articulated different kinds of harm AI art is creating for artists, including direct economic harms as well as chilling effects (that is, generative AI is discouraging artists from sharing their work online, making it harder for artists to find their audience) [46]. Lovato et al. (2024), in a survey of over 400 visual artists, found that a majority of artists believe training data for generative models should be disclosed and express concerns about the impact of AI on their industry [64]. We see this work as a primer for our own, as creative writers may or may not reason about generative AI in the same way as visual artists, and individual and industry impacts may differ.

### 3.3 Addressing Concerns about Data Usage

Technical research attempts to address artists' concerns about training data. To address concerns that generative AI may plagiarize pre-training data contributors, researchers have proposed approaches to minimize model memorization [15, 19, 43, 49, 55, 84] and style mimicry [59, 79, 82]. Another line of work attempts to protect the IP of writers or programmers by watermarking text [53, 75, 85] and attributing generated content to specific training data [23, 39]. To enable individuals to opt out of training data, researchers propose methods for unlearning [18, 45] and modular language model

architectures [38, 58, 67]. Such work presents potential technical solutions, but does not answer questions about whose work should be protected and under what contexts.

Several researchers have applied learnings from the library sciences about how to collect and archive data to machine learning contexts. Jo and Gebru (2020) outline five important aspects of data collection from archival and library sciences, and how they may be applied in the machine learning setting, such as actively collecting underrepresented data and allowing participants to denote sensitivity and access levels for their data. Relatedly, Desai et al. (2024) take an archival perspective on pretraining data explicitly, arguing that "to the extent that people use LLMs as interfaces into history and culture, the selection of data shapes and constrains that experience" just as archival collections do. While some advocate for more attention and care to be put into data collection [32], it has been documented that incentivizing data work is difficult in research and industry communities [34, 78].

## 4 Methodology

We engaged writers through hour-long interviews which were primarily conducted virtually. One of the authors conducted about two thirds of the interviews; another conducted the other third of interviews. Interviews were automatically transcribed, and then listened to by one author in order to clean up errors in the transcript. This meant that at least two authors listened, in full, to each interview (the original conductor of the interview, and the person who cleaned the transcript). Additionally, interview transcripts were often read in whole by a third author. The first four authors participated in the coding, and all authors participated in discussions about the codes.

### 4.1 Procedure and Analysis

We employed Grounded Theory to conduct the interviews and analyze the data. Because the details of how Grounded Theory is employed can vary depending on school of thought [70], rather than describing our approach as belonging to a certain lineage we instead detail our procedure. We began coding and theorizing about the interviews shortly after we began conducting them, which allowed us to constantly compare our theories with incoming interviews and update our interview guideline as the study progressed. Our initial interview guideline was based on concerns we had heard writers bring up in a variety of domains: as reported in the news, as noted in ongoing lawsuits, as represented by professional organizations such as the Author's Guild and the Author's Alliance, and as mentioned to us by writers we knew in our professional and personal networks. Although traditional Grounded Theory eschews existing literature, our research intended to directly investigate a growing issue. Therefore, we could not approach the data with "an empty head." Instead, we approached it with "an open mind" and indeed found that writers' reasoning and concerns about these issues did not always match the way they were represented by the news, lawsuits, or professional organizations.

The interview guideline was, thus, updated as emergent themes developed through coding. For instance, our initial guideline did not include questions about the nature of the institution collecting the training data (for profit, non-profit, research, etc.). However, early

on in the interviews it was clear that some writers cared a great deal about the type of institution, and so we updated our guideline to explicitly ask about this in subsequent interviews. Ultimately our guideline was updated five times, and the final guideline can be found in Appendix B.

Initially we employed open coding, but found that our interviews covered a lot of ground and open coding was, in fact, too open. Drawing on the approach of Deterding and Waters (2021), we then performed “index coding”, where we coded interviews based on the themes that were explicitly asked about in the interview guideline. Our developing interview guideline acted as the first set of codes. Then, individual index codes—such as attitude towards compensation, or concerns about reproduction of writing style—were collated and open coding was done within a single index code. This more narrow open coding allowed us to deeply investigate specific writerly concerns.

Then, axial and selective coding was done on all codes that came out of the open coding performed on each index code. This allowed us to see emergent and general themes that were occurring across the concerns that writers were bringing up. With these new, general codes, we performed a final round of coding in which all excerpts were re-coded with the general codes to both test their performance on the data, and to note any data that did not fit into this framework. These codes are what are reported as our results.

## 4.2 Recruitment

Writers were recruited in a variety of ways: through our professional and personal networks, through cold emailing writers from specific categories (e.g., romance novelists), through posting on writer forums (e.g., forums for writers that use AI in their writing practice), and through snowball sampling where we asked interviewees to suggest other writers to interview. All writers were from North America.

We initially wanted to recruit writers from a variety of genres (fiction, nonfiction, poetry, and fan fiction) and who publish in a variety of methods (big 5 publishing, independent publishing, self-publishing for profit, and self-publishing for free). Throughout the study we added three more types of variation we wanted to seek out: level of professionalization (majority of income comes from creative writing, majority of income comes from other writing-intensive profession, majority of income does not come from writing), self-described demographic information (race, gender, sexuality, and religion), and engagement with AI in writing practice (none, has tried it out, regular use in creative composition, and regular use for paratext e.g., marketing materials). Therefore, we engaged in theoretical sampling where we recruited participants based partially on what was coming up during our analysis. Note that we did not explicitly ask for demographic information (e.g., sexuality); instead a writer might’ve brought up self-described demographic information as being relevant to our questions, for instance noting a lack of writers like them being represented in training data. In Appendix A we give detailed accounts of the distribution of writers.

Although we did not recruit for this specifically, we found that our participants had a wide range of views on language models more generally. Some considered the technology to have no positive applications, others considered it to be neutral, and still others

thought it had pro-social applications, even if they were against the current way in which it was created or used.

## 4.3 Ethics Board Approval, Privacy, and Anonymity

This study was approved by the relevant ethics review board. Interviewees were anonymized, both by the removal of their names as associated with coded excerpts, as well as by removing or altering any details that may be identifying, for instance a description of their published work. Finally, although video and audio were recorded, only anonymized transcripts were retained, with video and audio recordings deleted within one year of the completion of the study.

## 4.4 Positionality Statement

All of the authors of this paper are based in North America, and all of our interviewees were similarly based in North America—either the USA or Canada. Although training data for LLMs can come from across the world, there is some evidence that most of it is written by those in North America [27, 38]. However, we note that the cultural context of North American writers may be different to those from other parts of the world.

The research team contained researchers from Computer Science, Information Science, Sociology, Rhetoric and Composition, and Law. In addition, several members of the research team have their own creative writing practice.

## 5 Findings

All our codes can be found in Table 1. First, we report on the principles that are the foundation for how writers think about writing in general, and that they use to reason about the use of their writing in this new context. Second, we report on their realistic expectations about how the world works. Here, we see writers’ principles can sometimes, but not always, clash against what they think is possible in the world they live in.

We note that there are a variety of ideas that thread throughout multiple codes. For instance, a major discussion point was if and how compensation may work. This idea comes up under the code **respect** (where compensation is often seen as a mechanism of showing respect) as well as under the code **interpretation of scale** (where writers reason about how much compensation could reasonably be expected) and **lack of control** (where writers note that it may be impossible to control how their writing is used, negating the possibility for fair compensation). Our codes outline how writers are reasoning about these issues, rather than presenting a list of concerns.

### 5.1 Principles

**5.1.1 The Creative Chain.** Participants often compared making their writing accessible to train a language model to their principles and expectations around making their writing accessible in other contexts. In general, participants recognize that reading, imitating, referencing and transforming other writing is essential both to learning to write and writing itself. In turn, many participants described a responsibility to support education and the production of art by making their writing accessible and taking a permissive



**Table 1: Overview of the codes that emerged from our grounded theory analysis.**

<b>Principles</b>	
The Creative Chain	<i>"It's just being a link in the creative chain. I made something. Somebody else saw that, and they got a spark of an idea to do something else, and they want to build off of it. We're all building off of each other constantly." (P20)</i>
Respect	<i>"I think it just comes again to respecting the works and respecting the people. Because I think that's probably my main issue with a lot of AI stuff is just the complete lack of respect for the works that they have been using to profit off." (P18)</i>
The Human Element	<i>"Honestly, I think a lot of the reason that people like reading other people's stuff is because there's the human element there. You're looking at one guy's perspective of an insane situation ... It's the project of one person's mind." (P13)</i>
<b>Realistic Expectations</b>	
Lack of Control	<i>"One of the things that we're all really feeling right now is a lack of any options in this. When I found out that my work was in Books3 I emailed my agent right away. And I was like, I know there's like nothing you can do, I just want you to know that this is a thing, and that I don't like it." (P20)</i>
Industry Impacts	<i>"We could easily be putting artists out of work. And what does that mean for them as artists and humans, [who need] to put food on their table, but also like, what does that mean for culture and society?" (P9)</i>
Interpretation of Scale	<i>"Cause I'm trying to think, okay, upsides and downsides. It's hard to pinpoint one, because again, I'm a little drop of water in a huge ocean. I don't have a distinct style, so I don't... I don't really know." (P24)</i>

stance towards work that references theirs. As P12 (professional novelist) put it, *"I think that being very rigid around intellectual property, and the defense of ideas... turns into an awful scenario that I don't really believe in, because I believe that all art inspires other art, and I love that basic premise."* P20 (professional novelist) described this as *"being a link in the creative chain"*: *"We're all building off of each other constantly, you know. My work is just the product of everything I've absorbed through my life... nothing makes me happier than seeing people make other things with my work."*

Making information free and accessible for others to learn from was an essential principle for several participants. Because of this, some writers felt ambivalent or positive about their writing being used to train language models, despite other concerns they might have (P5, P17, P24, P25, P27, P30).

EXPECTATIONS OF CONTROL IN THE CREATIVE CHAIN: WHAT IS THE PROJECT I AM A PART OF? In contrast to other creative projects a writer might contribute their work to, some participants expressed that the overall goal of language models was unclear to them, and it was therefore difficult to assess what it meant to contribute their writing. P3 (poet and essayist) noted that *"it feels similar to what I imagine an artist who's being curated might feel: ... they would want to know about the content or subject, or the nature of the thing that they're being curated into."* Understanding the goal of the project

was important for P14 (poet), who said, *"If I knew the end result or kind of, Here's why we're interested in your stuff, I would be much more interested in opting in."*

DISAGREEMENTS ABOUT THE ARTISTIC POTENTIAL OF LANGUAGE MODELS. Participants disagreed about whether contributing to training datasets could yield benefits such as supporting the production of worthwhile art, helping others to express themselves, or supporting education. Several participants had trouble imagining the artistic potential of language models because they did not like the style of generated text (P21, P25), found it *"generic"* (P6, P13, P24), or had concerns about the political biases of these models (P18, P21). Many others expressed doubts about the creative potential of LLMs, both because they view LLMs as *"regurgitation"* (P18, fan fiction writer) of other work rather than creating new ideas, and because of the fundamental role of humanity in creativity. (See [the human element](#) 5.1.3.)

On the other hand, several writers reported regularly using LLMs in their creative process (P1, P8, P15, P28), or cited literary work that used language models that they liked (P4, P12, P15, P22). Many others saw the value in using LLMs *"as a tool"* (P2) for creative expression, imagining that LLMs might help writers generate initial

ideas (P18, P27), provide writers with feedback or analytic summaries during the writing process (P3, P5, P16, P27), or help writers move past writer's block (P3, P13, P15, P18, P29).

**REPRESENTATION AND NOURISHING THE FUTURE OF TECHNOLOGY.** In addition to supporting the production of art, writers identified other positive, meaningful aspects of contributing their writing to training data. Some writers expressed excitement about the idea of “contribut[ing] to the future” (P28, poet) or “nourish[ing]” (P15, novelist and journalist) language models with literary writing (also P1, P26). As P11, a professional novelist who had never used an LLM, explained, “*I’ve had an influence on the way this machine thinks, which is incredible. I mean, it’s probably that’s my place in history more than any encyclopedia entry, you know.*”

Some writers with racial identities that are under-represented both in language model outputs and in the literary world more broadly were enthusiastic about contributing their data (P23, P27), as “*there needs to be more diversity in some of these models*” (P27).

**DECONTEXTUALIZATION: BEING INCLUDED COULD MISREPRESENT OR REIFY A WRITER’S STYLE, WORK, AND IDENTITY.** Some participants expressed concerns about how contributing to language models had the potential to decontextualize their work. For P24 (novelist), decontextualization is a threat to their artistic evolution and identity: “*I don’t want my past books to be training... because I want to always evolve and become a better writer, and I feel like that would leave me stuck in that same voice or ability.*” P32 (professional essayist and novelist) also expressed concerns about decontextualization, as they thought through what would come up if someone prompted a model trained on writing they wrote throughout their gender transitions. They were comfortable with someone reading their old work online, where it would be accompanied by contextual information like when it was written (“*I trust readers, I think readers are really smart*”); in an LLM, they feared it could instead be weaponized against their community.

For P19, the way language models racially decontextualize their training data was a big concern:

*Specifically as a Black writer, Black scholar, Black cultural worker and artist, if these models are being trained to give an output and the way that the ingesting happens is taking from a lot of different places, I think it supports this idea that ... we can sort of move towards a universal—either multiracial or deracialized—art. And then that, of course, is going to support people in positions of power, specifically white people. (P19)*

**EXPECTATIONS OF RESPECT AND HUMANITY IN THE CREATIVE CHAIN.** Participants made analogies between contributing their writing to language model training datasets and various ways their work does or could relate to other creative projects: it is like someone turning your book into a movie or play (P3, P12, P20), it is like having your style mimicked by your student or competition (P6, P7, P31), it is like having your words appear uncredited in someone else’s work (P12, P13, P17, P18, P29, P31), it is like someone writing fanfiction about your novel (P1, P18, P20), it is like being curated into someone’s art show (P3). Participants described how these different scenarios present different risks to their careers, reputations, and sense of creative control. As such, in each of these scenarios, participants described how they expect to receive respect through

some combination of credit, compensation, creative control, or at least the sense that a human being who admires their writing engaged with it, expressing a shared value of creative writing.

**5.1.2 Respect.** While writers understood that by publishing their work they are allowing others to use that work in perhaps unexpected ways, they also felt that the use of their work as training data indicated a lack of respect both for writing as an endeavor and writers as people. Although writers did not express that all uses of their writing require consent—for instance, as in the section above about the creative chain, they expect that others will learn from their work or perhaps even try to mimic it—the use of their writing as training data felt different. P6, a professional romance novelist, said that the use of their writing for AI may be more like signing away their foreign subrights, which would require consent and be part of contract negotiations when publishing.

**MECHANISMS OF RESPECT: COMPENSATION IS KEY FOR PROFESSIONALIZED WRITERS.** The idea that writers wanted respect was very much tied to the mechanisms in which respect can be shown. The mechanisms that writers desired was a function of many things: their degree of professionalization, the type of writing they worked on, their sense of precarity in the labor market, and their perception of the organizations that were or could be using their writing. Primarily, writers wanted respect for writing as a laborious act on its own—to write a book takes time, effort, and expertise. As P9 (professional romance novelist) put it,

*The biggest myth about creators, artists, writers, you know, name the person who creates art or entertainment in the world, is that we do it for the muse. Like, I live in this house, and I have to pay a bill every month to pay for this house, you know, I have to put food on my table and clothe my children and myself. Like this is my job... We should pay people for their work, right? (P9)*

The people we talked to who were professional writers—who earned the majority of the income from their creative writing practice—were most likely to cite compensation as the main mechanism of respect they would expect. For instance, P8, a professional freelance writer, noted that, “*And I say this as a, you know, a very pro AI person who thinks it’s the next creative way. But content producers actually need to be compensated on some level.*” P11 (professional journalist and novelist) made the comparison with other kinds of expertise: “*If NASA wanted my expertise to hire me to help in the space program they would actually have to pay me quite a lot of money, and they would do so.*” However, technology companies were not treating writers as if they brought any kind of expertise or as if their labor was valuable.

**POWER DIFFERENTIALS MODULATE DESIRE FOR COMPENSATION.** Writers were also quick to point out the large power differential between the corporations creating these models and the writers trying to earn a living. For instance, P7 (professional journalist) explains that, “*The tech companies are making money hand over fist, and I feel like they can compensate people.*” Professional writers were likely to say they would require compensation regardless of whether or not the model creator was making money, but many writers did make a distinction between organizations they perceived could afford to compensate writers, and non-profit organizations or

applications. We return to the idea of who can afford to compensate writers in the realistic expectations section.

DIFFERENT KINDS OF WRITING REQUIRE DIFFERENT KINDS OF RESPECT. Writers discussed genre distinctions in terms of the level of labor involved. P9, a professional romance novelist, talked about the use of certain kinds of books feeling particularly disturbing, “*like Pulitzer Prize-winning work that’s happening in nonfiction*,” which represents decades of research then being taken without consent or compensation and made “*part of the churn*.” In this case, writers wanted respect for time-consuming labor that writing represents.

Other writers noted that the use of certain kinds of writing felt more disrespectful than others because of the contents of the writing, rather than the labor taken to produce it. P4 (fan fiction writer and poet) talked about some writing being “*deeply personal*” and being “*disturbed by the lack of care*” that is shown by the act of putting such work into a language model. P32, in discussing a colleague who had written a nonfiction book which included accounts of assault, noted that their feelings and words about the person who assaulted them now lives on in the predictive models trained on that book, saying, “*There is an aspect of that that is like literal haunting*.” Here, respect was not about compensation but about understanding the sensitive nature of some kinds of writing.

HOW TO VALUE WRITING. Writers discussed how they might consider whether or not to contribute certain works of theirs, which led them to express how different kinds of writing may be valued differently. P14 (poet and translator) said that if they were asked for consent, “*there are things that I would probably hold back in order to wait until I’ve gotten the value out of them and then at that point I wouldn’t worry about it as much*.” Similarly P13 (fan fiction and science fiction writer) noted that they used to post a lot more of their ideas on social media before they started winning awards for their stories; now, they consider their ideas to be more valuable and are more hesitant about sharing them publicly. Other writers noted that older writing is often less valuable, either, as P14 noted, because they have already gotten value out of it, or because older writing is worse—they have become a better writer and therefore value their more recent writing more (P7, journalist and novelist).

The desire for respect was very much about respecting the humanity of the person who created the work. (As P32 noted, everyone deserves respect.) Implicitly, writers were often complaining that the use of writing as training data without consent or compensation indicated a lack of recognition that people had created the writing. When text can be scraped from the web, it is easy to forget how that text was created, or to recognize the text as more than “just” data. This led writers to articulate what they consider special about their writing compared to other kinds of writing, including text that may be generated by computers, which we report on next.

**5.1.3 The Human Element.** Large language models generate text because they are designed to respond efficiently and accurately to a prompt via token prediction. If a human writer gave this as a reason for why they write, they would be, as Shannon Vallor notes, “doing it wrong” [81]. The writers we spoke with saw their writing as fundamentally different from the writing produced by LLMs, because of their humanity and readerships that similarly value and respect their humanity through their writing. With this principle, we aim to roughly capture how and why writers distinguish their

writing from the text generated by LLMs. Creative writers have many reasons for why they write, but notably, none of them are about efficiency or prediction. Instead, some of the many reasons cited include: creating art (P15, P18, P25), making a living (P8, P9), processing and sharing emotion (P10, P16), expressing one’s experiential, embodied subjectivity (P5, P19, P27), and improving one’s thinking on and understanding of culture, place, society, and oneself (P13, P14, P29). Underlying each of these reasons are three key sub-principles: 1) that writing is a fundamentally human activity distinct from the text generation of LLMs, 2) that the humanity of writing is valuable to both writers and readers, and 3) that the ‘readership’ for a training dataset is not the same human audience that writers wish to reach.

WRITING AND TEXT-GENERATION ARE NOT THE SAME. The parsing and production of text is often enough to garner comparisons between creative writing and text generated by LLMs. However, many writers point out that this is a false equivalence. Motivations and methods for each differ dramatically.

Many writers see the text generation of LLMs as an attempt to mimic human writing, and, in many cases, an author’s unique style or voice. This is tied to discussions of [the creative chain](#) (subsection 5.1.1) and writers view the training of LLMs on a particular writer’s voice as impoverished and decontextualized. When they read writing that inspires them, writers appreciate, contextualize, question, compare, and transform a text within their own social, embodied, and emotional contexts, something mass aggregates of textual data cannot do. P6 (professional /novelist) contrasts these processes: “*My absolute favorite book was instrumental in the way that I wrote the favorite of all of the books I’ve written. But it’s still coming through me. Whereas, I feel like ChatGPT would end up poorly mimicking me, as opposed to being something that came from someone else, but was formed by their appreciation for me*.”

Several writers express that what LLMs do when they generate text in an existing author’s voice is like a “cheat code” for writing and that a person’s experience and perspective enable real writing to happen. P17 says:

*Write your own goddamn book. Like, you gotta go through it. That said, my writing’s very personal, it’s very autobiographical. So, no one’s gonna write my book, I get that. But it does feel kinda like a cheat code, you know? Like, I shouldn’t be able to go in and [ask a model], ‘Hey, rewrite this scene in the voice of Cormac McCarthy.’ Cormac spent his whole life doing that, so why should I be able to just rip it off? (P17)*

Implicit in these distinctions is an understanding that LLMs do not have authorial intent, which is central to writing as a human activity. That LLMs do not have authorial intent when generating text is viewed as an asset to those writers choosing to use them in creative ways. P31, who has used LLMs in their creative practice, notes this explicitly, while problematizing how LLMs are colloquially talked about as “collaborators” in some contexts. They say:

*I’m troubled by the idea that it’s a collaboration or a co-writer. I don’t think like that when I’m doing it. That’s attributing maybe a bit too much power or agency to it. I don’t ascribe [a] pseudo sentence to it or anything*



*like that. ... There's no will on the other side to make something.* (P31)

This is a positive thing for all of the writers we spoke with who use LLMs creatively. Because the writer has human agency and the LLM does not, writers maintain both control and uphold the value of humanity in the creative process.

WRITING'S HUMANITY MAKES IT VALUABLE. Almost all of the writers we spoke with expressed some variation on the idea that at least part of the value of writing comes from the fact that it was written by a person, to be read by another person. Writers consistently offered things that readers value in writing—a unique style, voice, or perspective, authenticity, and humanity chief among them. What writers think readers of creative writing value is not easily replicable by LLMs. “*I write poetry as a reason to connect with other people,*” P16 says. P13 remarks: “*Honestly, I think a lot of the reason that people like reading other people's stuff is because there's the human element there. You're looking at one guy's perspective of an insane situation ... It's the project of one person's mind.*”

Writers who use LLMs as part of their creative processes also see distinct differences between the motivations and methods of writers and LLMs. Some of these writers use the ‘mimicry’ that LLMs enact as valuable parts of their process. P25 describes how they use the mimicry of LLMs in their work positively, after inputting their writing as source material into an LLM:

*I want the machine to start telling me things that I haven't [thought about], like a little thing for me to come up with a better question. Once I come up with a better question, not for the machine but for myself, then I either go and look it up or I'm like, 'Great, now I'm gonna go think about it.'* (P25)

THE AUDIENCE FOR A TRAINING DATASET OF CREATIVE WORK IS NOT THE AUDIENCE THAT WRITERS WANT OR IMAGINE. Undergirding our conversations with writers was the understanding that the imagined or intended audience for a writer's work is other people, not a dataset using writing to train an LLM. Every direct mention of the words ‘audience’ or ‘reader’ in our interviews was in reference to an explicitly human audience. At least one of these two terms was mentioned in all of the interviews we conducted.

In light of the strong connection between humanity and writers' conceptions of audience, writers were left to conjecture about who the ‘audience’ of training data really was. Algorithms, corporations, and users of LLMs were all variously floated as ideas of what might constitute the readership for training datasets. P26 wants to share their work with people, but is “*not ready to put [even] a little piece of it into the algorithm.*” The scale and combinatorial nature of training datasets provoked skepticism and sometimes fear around this imagined ‘readership,’ quite unlike the readerships writers imagine for themselves.

## 5.2 Realistic Expectations

### 5.2.1 Lack of Control.

*One of the things that we're all really feeling right now is a lack of any options in this. When I found out that my work was in Books3 I emailed my agent right away. And I was like, I know there's like nothing you can do, I just want you to know that this is a thing, and that*

*I don't like it. And if there are any conversations on the publisher side, on the agency side, anything where I can add my voice to that, anything we can take action with, I want to do that. Then he wrote back and he said, I know, I've had this same email dozens of times this week. Everyone hates this, nobody wants their work being used like this, but we have no way to opt out. There is no legal recourse for it.* (P20)

All writers emphasized the lack of control they had in this situation. Lack of control over what is done with their writing more generally, as well as lack of influence over or trust in how technology corporations might make use of their work. While writers acknowledged that the use of their writing in unexpected ways is, in fact, an expected part of publishing—i.e., being part of the creative chain—they also expressed that certain use cases are regulated.

CERTAIN USES OF WRITING ARE RESTRICTED AND THERE ARE AVENUES FOR RECOURSE. Writers often sign contracts that outline a publisher's right to republish a book in a foreign country, or create a movie adaption, indicating that writers do have control over a variety of use cases of their writing. For instance, P11 (a professional novelist and journalist) said, “*There are lots of sources of pirated books on the Internet, but they're illegal, and we try to shut them down. If ever I find a free copy of my book circulated somewhere, I'll ask my publisher to send a cease and desist letter, and it usually gets taken down.*”

Currently, the use of writing as training data for LLMs is both outside of writers' control and they lack any recourse when their writing is used in this way against their wishes. Writers understood that this lack of recourse had several sources. In addition to their lack of control, writers agreed that, even if such use was restricted, it is difficult, if not impossible, to stop models from being trained or remove models that already exist (P2, P6, P12, P32).

IT IS DIFFICULT TO UNDERSTAND HOW TO RESTRICT ONLY CERTAIN USE CASES OF LLMs. Writers did not always disagree with the creation of language models per se, and instead had concerns only about certain use cases. But many writers noted that it seems impossible to prevent certain use cases. For instance, P24 (a fantasy novelist) would consider consenting to the use of their writing if they could ensure that the model would not be used to compete with them, but felt “*like that would be an impossible ask. People are just gonna lie and go 'No, I'm not writing a novel.'*” Others noted that while it may be possible to audit big companies, “*there are all these people putting out small models, anyone can do this now, and it almost feels impossible to rein in*” (P2, a writer experienced with LLMs).

Several writers said that they would be happy for their writing to contribute to prosocial models, for instance models that supported literacy efforts or were specifically designed to help people without access to other forms of education. Yet, they noted that they have no control over who makes the models or what they are used for. In this way, they called upon the idea that they were not being involved in model making, which is one of the ways writers wished to be given respect for their work and their expertise.

DISTRUST IN TECHNOLOGY COMPANIES FURTHERS THE BELIEF THAT RESTRICTION IS IMPOSSIBLE. Others noted that they didn't trust that even if they were promised that certain use cases could and



would be restricted, that those promises would be upheld. *“The way that it’s been working so far is that people take stuff, and then later we’ll be like, Oh, my God, I’m so sorry! It’s too fucking late, it already happened”* (P13). One writer articulated how their understanding of the actions of technology companies in the past influences their expectations with this particular issue:

*This is what tech companies have always done. Uber is like, we’re just gonna make these cars. And we’ll see if cities can catch up with us. Or Google Maps when they’re like, we’re just gonna photograph everybody’s house. And then we’ll see if anyone gets mad. I think that’s totally what AI writing software is doing. They’re waiting to see what people are gonna do.* (P12)

Finally, writers acknowledged the difficulty of predicting future use cases. As technology changes, it can be hard to know what you might be agreeing to in the moment. Some established professional novelists were especially concerned about those at the start of their careers—established writers were concerned about up-and-coming authors (P9, P12, P20) and writers we interviewed who were at the start of their careers noted their particular predicament (P13). As an early stage writer, one has less power to negotiate certain terms out of their contract and may end up agreeing to terms that, years down the line, perhaps when the writer is more famous, are unexpected, unprecedented, and ultimately undesirable.

**5.2.2 Industry Impacts.** Though public discourse about the impact of language models on the creative writing industry often takes place in the future, writers described present day impacts of language models on their careers and fields. P12 described the strain language models have placed on literary editors, who *“spend hours every week”* sifting through large amounts of *“junk pitches and generated stories.”* P12 recalled how this influx of generated work led one prominent literary magazine (Clarkesworld) to *“close submissions entirely,”* which *“really hurts people who are just getting started, people who already have this narrowing slate of options where they can submit their work.”* Other writers expressed similar concern over the impacts of language models on new writers getting started in the industry (P9).

Writers with academic teaching positions described confronting the impacts of language models on creative writing education. These writers explained how their creative writing careers and teaching careers are inextricable, either because they are hired to teach because of their creative writing careers, or because they see teaching as part of their creative practice. These writers note that even if they are morally against LLMs, for instance because of the energy used to train and use them, they must deal with this technology in their classrooms (P12, P14, P16, P18, P32).

Some writers also described positive impacts of language models and AI on their careers. Self-published authors described using language models and other forms of generative AI to help with the work a traditional publisher might do, like creating a cover or producing marketing materials for their work (P24, P27). P31, who makes use of language models in their work, described how their career has benefited from advances in LLMs and the resulting increase in public interest in AI.

**PREDICTED FUTURE IMPACTS.** By and large, most participants found that in their current state, language models are not direct

competition for writers as they do not produce sufficiently sophisticated, high quality writing. When speculating about how future advanced versions of language models might impact their careers and industries, many writers imagined competing with a language model that had been trained to mimic their writing. Other writers speculated more broadly about how language models might compete with or displace writers, i.e., *“an AI that writes better novels than humans”* (P12, professional novelist).

Writers expressed different perspectives about whether or not these potential futures were a threat to the creative writing industry in general and their careers in particular. P6, a romance novelist who makes her living by selling her novels, finds these futures imminent and threatening to her livelihood: *“I asked [my finance guy] how fast we could move up my retirement...I think the day is going to come when it’s harder for me to make a living, because so many people can do exactly what I do.”* Other commercially successful writers speculated that because their readers are interested in their specific style or voice, they *“have very little to worry about”* (P9). As P11 put it, *“I really don’t see why anyone would look to buy an AI generated book by [P11] when mine are readily available.”*

Several authors expressed a lack of concern about the commercial or monetary impact of LLMs on creative writing, because of the value readers place on the humanity behind a creative work. P10 emphasizes the elisions and false equivalencies drawn between writing and text generation in their response, underscoring the value of writing as a distinctly human enterprise.

*I don’t think [LLMs are] particularly useful for the creative industry. I mean, I could see them being useful for, like, ... discovering a new kind of a new antibiotic or something like that... But, human expression? I guess I think of poetry or fiction as a deeply human activity. And it sort of baffles me why people would want to read what has been historically a human activity that has been produced by a machine.* (P10, poet and essayist)

Because of the importance of a human perspective in creative writing, some writers do not believe that language models could ever *“evolve to where ... [it can] come up with new stories that are relevant to the current readership”* (P24). P11 imagined that while people would be interested in, for example, seeing the first movie with an AI-generated script, the novelty would soon wear off, and LLMs would just become *“part of our toolkit to make good art”*.

**BACKDROP OF A CHANGING INDUSTRY.** Some writers acknowledged that their resistance to contribute to training datasets was *“compounded by a sense of precarity”* in the creative writing industry (P3, poet). P20 (professional science fiction novelist) noted that *“it’s a really difficult time for people in our profession and for creative industries in general. The job market is constantly shrinking.”*

Some writers expressed protectiveness towards the creative industry, citing the value of humanity intrinsic to creative work. The questions these writers were asking were existential, implying that the existence of AI-generated art and literature is forcing us to ask, *“What kind of world do we want to live in?”* P9 (professional romance novelist) notes:

*We could easily be putting artists out of work. And what does that mean for them as artists and humans, [who need] to put food on their table, but also like, what does*

*that mean for culture and society? The argument from the AI side is, 'Well, this is culture and society too. This is just advancement.' But, as an artist, I'm always going to say humans make art because it's about feelings.* (P9)

P13, along with several other writers we spoke with, wants to live in a world where people can make a living as writers: *"I don't think it's good to make it cheaper to make art for the fact that people need to be employed and I would prefer to be employed as a writer than as a lawyer."*

Others expressed that since they did not expect commercial success as a creative writer, they felt *"ambivalent"* (P14) about the impacts of language models on their careers. As P14 put it, *"You know, I've made \$20 on poetry over the course of my life. So. It's not, yeah, I'm not on [the commercial] side of [writing] or interested in that."* P30 similarly agreed that for poets, *"the stakes are different": "as a poetry MFA student ... we're not doing it for the career prospects."*

**5.2.3 Interpretation of Scale.** The amount of training data used to train large language models is hard to comprehend, and often, in fact, opaque to those outside of the corporation that trained the model. To briefly contextualize this section, we point out that GPT-3 was trained on 700GB of text [13], and newer models are trained on significantly more, such as Llama 3 which was trained on over 15TB of data [60]. If we assume the average book contains about 100,000 words, then a single book would be the equivalent of 0.1MB, or 0.0000001% of the training data of GPT-3. Such magnitude is difficult to reason about; as P24 said, *"Cause I'm trying to think, okay, upsides and downsides. It's hard to pinpoint one, because again, I'm a little drop of water in a huge ocean."*

COMPENSATION IS HARD TO DETERMINE AND LIKELY TO BE INSIGNIFICANT GIVEN THE SCALE OF LLMs. While most writers objected to their data being collected by corporations, they simultaneously had difficulty comprehending how to value their contribution when considering the scale of training data used. Some noted that their style was not distinct enough to be noticeable in any outputs or to significantly change the model (P2, P24). P10 (poet and essayist) described this as, *"It's hard for me to trace how money being made is based on the training data."* Others felt it was hard to feel personally slighted when their contribution is so very small (P29, fan fiction writer). Still others were unsure exactly how compensation would even work: *"It's not like a Spotify thing where you can be like, Okay, they listen to this track, or they read this chapter"* (P17).

Despite many writers wanting compensation for the use of their work, they acknowledged that such compensation is likely to be small given the amount of data needed to train LLMs. P15, who has worked with LLMs in their creative practice, reflecting that they may be one of one million authors, noted that *"I don't actually think that the companies who have created these models owe me a large amount of money"* even if the amount of money needed to compensate all authors would be significant. Other writers noted that they wouldn't even necessarily want a small amount of compensation: *"If it was gonna be 23 cents, because that's really how much it's worth, in some sense like... I think I should care more. But I don't actually want a check for 23 cents"* (P7). P32 noted that a very small amount of compensation in fact shouldn't be considered compensation at all—compensation necessitates that it compensates the writer, and a small amount wouldn't actually be compensation for their work.

ALTERNATIVES TO COMPENSATION MAY ALSO BE IMPRACTICAL. Some writers, either in addition to or in lieu of compensation, wanted to be credited for their contribution to an LLM (P14, P17, P18, P23, P24, P27). But P24 noted that while being credited in a model was an interesting idea, it ultimately seemed futile:

*I guess, because no one will go look at [a list of contributing authors]. They won't. They won't care. If I'm going to write something, I just want to use the model. I just want to chat, write in my prompt and get my results. I'm not going to take the time to go, Gee! I wonder what 864,000 people were used to train this model.* (P24)

A few writers did consider ways to deal with the scale of LLM training data. P18 (fan fiction writer) wanted credit to be based on the output, such that outputs that were *"more influenced"* by some writing could highlight the particularly relevant contributions. P15 suggested that rather than consider compensation at the individual level, *"the moral thing to do would be to figure out some way to offset the disruption you're causing. The way that you would do that isn't necessarily just by paying everybody whose work has been fed into you know... it would be some other grant or taxation."* Similarly P33 noted that industry disruption felt more salient than disruption to their individual career. P14 suggested that something like stock options would allow them to be compensated in line with the value that was created.

THE IMPRACTICALITY OF RESPECT MADE SOME DISENGAGE ENTIRELY. Writers' understanding of the scale of LLM training data made it difficult for them to square their desire for respect as laborers and creatives. Writers wanted respect for their work, whether this be through consent, compensation, credit, or involvement, but felt that all and any of those were impractical. It would be impractical to get everyone's consent (P11), compensation would be too small (P7, P15, P32), credit would be futile (P24), and no one expected to be invited into the upper ranks of technology corporations. This made many writers turn away from LLMs altogether; several of our interviewees actively disengaged with LLMs because of these issues, including never having tried them out themselves or avoiding learning more about them, even if they thought LLMs were theoretically or morally fine.

## 6 Discussion

*There is a deep sort of sadness to this notion that language can be mimicked in this way and that it can be sort of produced and really have no essential meaning or a central humanity, or at least not a detectable one. And then the idea that eventually these models and these applications will be able to simulate essential humanity. Or actually write moving, tear-jerking, compelling, bestselling novels. I mean, that's a pretty interesting idea. And I'm sure it's a place that we're going, and I'm not sure that we're prepared for it.* (P12)

As seen in the above quote, writers often traversed a range of emotions in response to the implications of LLMs, and many writers noted that they were not sure what should be done. Overall, while writers had concerns about the implications of language models abstractly, most writers expressed distaste for large corporations

using their work for dubious or unclear ends, and making a profit along the way.

An implicit theme in many of our interviews was the way in which LLMs turned writing into data. Writers consider their work—writing—to be valued in a variety of ways (e.g. appreciation, compensation, prestige) by a human readership. LLMs, in defiance of these norms, subsume writing into the larger and more contentious category of data, where value becomes murky and the audience becomes algorithmic rather than human. Data has been analogized to labor [8], property [73], and natural resource (e.g. "data mining"). Such analogies, although never perfect, attempt to sort through how data should be valued, especially data generated by, and often about, people. But is writing data?

We argue that writing becomes data when it is used to train generative models, and this transformation creates the issues writers raise, many of which stem from a reorientation of their audience from people to algorithms or perhaps companies. Writers noted the power differential between themselves and the corporations making use of their work. This can be understood through the frame of precarious work [5]; as writers feel that their work is increasingly precarious, the power differential between themselves and the organizations seeking to train LLMs grows larger.<sup>1</sup>

We can also see this dynamic as a form of extraction, in which writing is "taken" from writers for the profit of others. Couldry and Mejias (2019) propose 'data colonialism' as a new form of colonialism to make sense of the use of large amounts of data by a small group of corporate and government actors. If data colonialism is a useful framework to understand data collection for language models, which we think it may be, then we can turn to theories of decolonialism to understand how to resist it. Mohamed et al. (2020) put forth the idea of dismantling power asymmetries to resist data colonialism. This moves beyond our initial exploration of how data may or may not be collected in this context.

While our findings often align with those studying visual artist communities—work by Jiang et al. (2023) finds that AI image generators create a chilling effect on cultural production, similar to our findings about industry impacts on writers—the direct economic impacts seem to be distinct (e.g. few of our interviewees expected to see personal economic impacts). However, the experience of precarious work and a lack of power remain the same. We believe that decolonial thinking can address a multitude of concerns (e.g. through pro-worker AI policies [4] or public AI initiatives [66]) even as they differ in their specifics across communities.

When asked who they expect to best advocate for their interests, writers did not think governments, publishers, universities, and certainly not technology companies, had their interests at heart. Some writers relied on their agents, but most writers would only trust other writers to safely and respectfully steer LLM creation. We recommend that writers form political communities which can take the helm in determining how their work is used.

As we reported in [the creative chain](#) (5.1.1), while writers expect their writing to be part of a larger cultural project over which

they have little control, LLMs challenged their notion of what it meant to participate in the creative chain. Some writers envisioned how LLMs could be more aligned with the kinds of cultural projects that they wanted to contribute to. For instance, several writers discussed the positive values of libraries, and considered a world where LLMs were more like libraries. Yet, libraries and language models have important differences. The understood intent of libraries is that a wider variety of people can have direct access to your writing, whereas the final audience of an LLM may be ignorant of individual author contributions. Libraries are non-profit institutions that serve the general public. They are also physical locations that serve their local community in a variety of ways.

Some of these differences could be addressed: LLMs could be created by non-profit institutions to serve the general public, as recommended in [66] and seen in projects like [54] and [80]. The contributions of individual authors could be made more clear, for instance through work on attributing generated text to specific training data [23, 39]. However, it is likely difficult to completely remove the obscurity that lies between training data and generated text, and this disjuncture may continue to cause unique problems.

Although our work primarily investigated training data for *large* language models, writers did express interests in smaller models that were either private or local to a small community. As P2 put it, *"I really love the idea of personal, more customized models for your friends and groups ... it just feels like fun exploration."* Such models, with smaller amounts of creative writing in the training data, would make their place in the creative chain more clear, and make it easier to engage in respect for those who contributed to the training data.

## 7 Future Work

We outline two avenues of future work. The first is how to prevent the use of writing as training data without consent. The second is how to create models that writers would want to contribute to. We see these as the negative and positive sides of futuring: how to discourage a future we don't want, and create a future we desire.

**Preventing nonconsensual usage.** While there is preliminary work on detecting if certain text has been used to train a model [39], or applying watermarks to text [53, 75, 85], more robust and easy-to-access methods would allow writers to determine whether their writing has been used without their consent, which is the first step in recourse. Legal protections could enforce this, but social norms could also discourage institutions from nonconsensual usage.<sup>2</sup> Research into the economic impacts of LLMs on writers could also support this agenda.

**Creating writer-controlled language models.** Community datasets could allow writers to contribute only to projects they believed in. Such work would treat writers as experts rather than resources, and could draw on research on community archives [16, 30, 31, 42] and archival practices for machine learning [24, 47]. Technical work that supports smaller models that can be run locally, or models that can more easily interact with sensitive training data [67], could support community models which may not have the compute power of corporate ones.

<sup>1</sup>Such power differentials have been explored in fan fiction communities. For instance, Busse and Farley (2013) report on fan fiction writers who note that the balance of power between two fan writers and between a fan writer and a professional writer is completely different, resulting in different social norms for remixing professional work than fan work.

<sup>2</sup>For example, fan fiction communities have been found to have highly consistent social norms which are policed internally [28].

## 8 Conclusion

Most of the writers we talked to were not against the use of creative writing to train LLMs in theory, but rather against the way in which LLMs have been steered and created thus far. Even those who could not imagine why someone would want to use an LLM in their writing practice were open to the idea that people should be allowed to experiment with this technology. Still, most writers were either exasperated with the situation, seeing no way to influence it, or they simply accepted it, often due to their interpretation of their small contribution given the scale of LLM training data. In this work, we attempt to represent how creative writers are reasoning about the use of their writing as training data for LLMs, and articulate a path forward that gives writers the respect they deserve.

## References

- [1] [n.d.]. Generative AI – Intellectual property cases and policy tracker. <https://www.mishcon.com/generative-ai-intellectual-property-cases-and-policy-tracker>. Accessed: 2024-09-11.
- [2] [n.d.]. New Authors Guild AI Survey Reveals That Authors Overwhelmingly Want Consent and Compensation for Use of Their Works. <https://authorsguild.org/news/ag-ai-survey-reveals-authors-overwhelmingly-want-consent-and-compensation-for-use-of-their-works/>. Accessed: 2024-09-11.
- [3] 2023. WGA on Strike. <https://web.archive.org/web/20230502190624/https://www.wgacontract2023.org/announcements/wga-on-strike>.
- [4] Daron Acemoglu, David Autor, and Simon Johnson. 2023. Can We Have Pro-Worker AI? *CEPR Policy Insight* 123 (2023).
- [5] Blake A. Allan, Kelsey L. Autin, and Kerrie G. Wilkins-Yel. 2021. Precarious Work in the 21st Century: A Psychological Perspective. *Journal of Vocational Behavior* 126 (April 2021), 103491. <https://doi.org/10.1016/j.jvb.2020.103491>
- [6] Authors Alliance. 2023. Comment from Authors Alliance, Posted by the U.S. Copyright Office. <https://www.regulations.gov/comment/COLC-2023-0006-8976>. Comment ID: COLC-2023-0006-8976.
- [7] Echo Wang Anna Tong and Martin Coulter. 2024. Exclusive: Reddit in AI content licensing deal with Google. *Reuters* (Feb 2024). <https://www.reuters.com/technology/reddit-ai-content-licensing-deal-with-google-sources-say-2024-02-22/>
- [8] Imanol Arrieta-Ibarra, Leonard Goff, Diego Jiménez-Hernández, Jaron Lanier, and E Glen Weyl. 2018. Should we treat data as labor? Moving beyond “free”. In *aea Papers and Proceedings*, Vol. 108. American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203, 38–42.
- [9] Leah Asmelash. 2023. These books are being used to train AI. No one told the authors. *CNN* (Oct 2023). <https://www.cnn.com/2023/10/08/style/ai-books3-authors-nora-roberts-cec/index.html>
- [10] Jack Bandy and Nicholas Vincent. 2021. Addressing “documentation debt” in machine learning: A retrospective datasheet for BookCorpus. In *Thirty-fifth Conference on Neural Information Processing Systems, Datasets and Benchmarks Track*.
- [11] Karim Benharrah, Tim Zindulka, Florian Lehmann, Hendrik Heuer, and Daniel Buschek. 2024. Writer-Defined AI Personas for On-Demand Feedback Generation. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*. 1–18. <https://doi.org/10.1145/3613904.3642406> arXiv:2309.10433 [cs]
- [12] Ian Bogost. 2023. My Books Were Used to Train Meta’s Generative AI. *Good. The Atlantic* (Sep 2023). <https://www.theatlantic.com/technology/archive/2023/09/books3-database-meta-training-ai/675461/>
- [13] Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. Language Models Are Few-Shot Learners. *arXiv:2005.14165 [cs]* (July 2020). arXiv:2005.14165 [cs]
- [14] Kristina Busse and Shannon Farley. 2013. Remixing the remix: Fannish appropriation and the limits of unauthorised use. *M/C Journal* 16, 4 (2013).
- [15] Nicholas Carlini, Daphne Ippolito, Matthew Jagielski, Katherine Lee, Florian Tramèr, and Chiyuan Zhang. 2023. Quantifying Memorization Across Neural Language Models. In *The Eleventh International Conference on Learning Representations*.
- [16] Michelle Caswell, Marika Cifor, and Mario H Ramirez. 2016. “To suddenly discover yourself existing”: uncovering the impact of community archives. *The American Archivist* 79, 1 (2016), 56–81. <https://doi.org/10.17723/0360-9081.79.1.56>
- [17] Kent Chang, Mackenzie Cramer, Sandeep Soni, and David Bamman. 2023. Speak, Memory: An Archaeology of Books Known to ChatGPT/GPT-4. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, Houda Bouamor, Juan Pino, and Kalika Bali (Eds.). Association for Computational Linguistics, 7312–7327. <https://doi.org/10.18653/v1/2023.emnlp-main.453>
- [18] Jiaao Chen and Diyi Yang. 2023. Unlearn What You Want to Forget: Efficient Unlearning for LLMs. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, Houda Bouamor, Juan Pino, and Kalika Bali (Eds.). Association for Computational Linguistics, Singapore, 12041–12052. <https://doi.org/10.18653/v1/2023.emnlp-main.738>
- [19] Timothy Chu, Zhao Song, and Chiwun Yang. 2024. How to Protect Copyright Data in Optimization of Large Language Models?. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 38. 17871–17879.
- [20] John Joon Young Chung and Max Kreminski. 2024. Patchview: LLM-Powered Worldbuilding with Generative Dust and Magnet Visualization. <https://doi.org/10.1145/3654777.3676352> arXiv:2408.04112 [cs]
- [21] Nick Couldry and Ulises A. Mejias. 2019. Data Colonialism: Rethinking Big Data’s Relation to the Contemporary Subject. *Television & New Media* 20, 4 (May 2019), 336–349. <https://doi.org/10.1177/1527476418796632>
- [22] Andrew Dalton. 2023. Hollywood’s actors strike is nearing its 100th day. Why hasn’t a deal been reached and what’s next? *ABC News* (Oct 2023). <https://web.archive.org/web/20231020002355/https://abcnews.go.com/Business/wireStory/hollywoods-actors-strike-nearing-100th-day-deal-reached-104136181>
- [23] Junwei Deng, Shiyuan Zhang, and Jiaqi Ma. 2024. Computational Copyright: Towards A Royalty Model for Music Generative AI. *ICML Workshop on Generative AI + Law* (2024).
- [24] Meera A Desai, Irene V Pasquetto, Abigail Z Jacobs, and Dallas Card. 2024. An archival perspective on pretraining data. *Patterns* 5, 4 (2024).
- [25] Nicole M Deterding and Mary C Waters. 2021. Flexible coding of in-depth interviews: A twenty-first-century approach. *Sociological methods & research* 50, 2 (2021), 708–739.
- [26] Paramveer S. Dhillon, Somayeh Molaei, Jiaqi Li, Maximilian Golub, Shaochun Zheng, and Lionel Peter Robert. 2024. Shaping Human-AI Collaboration: Varied Scaffolding Levels in Co-writing with Language Models. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*. ACM, Honolulu HI USA, 1–18. <https://doi.org/10.1145/3613904.3642134>
- [27] Jesse Dodge, Maarten Sap, Ana Marasović, William Agnew, Gabriel Ilharco, Dirk Groeneveld, Margaret Mitchell, and Matt Gardner. 2021. Documenting Large Webtext Corpora: A Case Study on the Colossal Clean Crawled Corpus. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, 1286–1305. <https://doi.org/10.18653/v1/2021.emnlp-main.98>
- [28] Casey Fiesler and Amy S. Bruckman. 2019. Creativity, Copyright, and Close-Knit Communities: A Case Study of Social Norm Formation and Enforcement. *Proceedings of the ACM on Human-Computer Interaction* 3, GROUP (Dec. 2019), 1–24. <https://doi.org/10.1145/3361122>
- [29] Casey Fiesler and Nicholas Proferes. 2018. “Participant” perceptions of Twitter research ethics. *Social Media+ Society* 4, 1 (2018). <https://doi.org/10.1177/2056305118763366>
- [30] Andrew Flinn. 2007. Community histories, community archives: Some opportunities and challenges. *Journal of the Society of Archivists* 28, 2 (2007), 151–176.
- [31] Andrew Flinn, Mary Stevens, and Elizabeth Shepherd. 2009. Whose memories, whose archives? Independent community archives, autonomy and the mainstream. *Archival Science* 9, 1-2 (2009), 71–86. <https://doi.org/10.1007/s10502-009-9105-2>
- [32] Timnit Gebru, Jamie Morgenstern, Briana Vecchione, Jennifer Wortman Vaughan, Hanna Wallach, Hal Daumé Iii, and Kate Crawford. 2021. Datasheets for datasets. *Commun. ACM* 64, 12 (2021), 86–92. <https://doi.org/10.1145/3458723>
- [33] Samuel Gehman, Suchin Gururangan, Maarten Sap, Yejin Choi, and Noah A. Smith. 2020. RealToxicityPrompts: Evaluating Neural Toxic Degeneration in Language Models. In *Findings of the Association for Computational Linguistics: EMNLP 2020*. Association for Computational Linguistics, 3356–3369. <https://doi.org/10.18653/v1/2020.findings-emnlp.301>
- [34] Katy Ilonka Gero, Payel Das, Pierre Dognin, Inkit Padhi, Prasanna Sattigeri, and Kush R Varshney. 2023. The incentive gap in data work in the era of large models. *Nature Machine Intelligence* 5, 6 (2023), 565–567. <https://doi.org/10.1038/s42256-023-00673-x>
- [35] Sarah Gilbert, Katie Shilton, and Jessica Vitak. 2023. When research is the context: Cross-platform user expectations for social media data reuse. *Big Data & Society* 10, 1 (2023), 20539517231164108.
- [36] Sarah Gilbert, Jessica Vitak, and Katie Shilton. 2021. Measuring Americans’ comfort with research uses of their social media data. *Social Media+ Society* 7, 3 (2021), 20563051211033824.
- [37] Suchin Gururangan, Dallas Card, Sarah Dreier, Emily Gade, Leroy Wang, Zeyu Wang, Luke Zettlemoyer, and Noah A. Smith. 2022. Whose Language Counts as High Quality? Measuring Language Ideologies in Text Data Selection. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language*



- Processing. Association for Computational Linguistics, 2562–2580. <https://doi.org/10.18653/v1/2022.emnlp-main.165>
- [38] Suchin Gururangan, Mike Lewis, Ari Holtzman, Noah A. Smith, and Luke Zettlemoyer. 2022. DEMix Layers: Disentangling Domains for Modular Language Modeling. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. Association for Computational Linguistics, 5557–5576. <https://doi.org/10.18653/v1/2022.naacl-main.407>
- [39] Zayd Hammoudeh and Daniel Lowd. 2024. Training data influence analysis and estimation: A survey. *Machine Learning* 113, 5 (2024), 2351–2403.
- [40] Libby Hemphill, Angela Schöpke-Gonzalez, and Anmol Panda. 2022. Comparative sensitivity of social media data and their acceptable use in research. *Scientific Data* 9, 1 (2022), 643. <https://doi.org/10.1038/s41597-022-01773-w>
- [41] Md Naimul Hoque, Tasfia Mashiat, Bhavya Ghai, Cecilia D. Shelton, Fanny Chevalier, Kari Kraus, and Niklas Elmqvist. 2024. The HaLLMark Effect: Supporting Provenance and Transparent Use of Large Language Models in Writing with Interactive Visualization. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*. ACM, Honolulu HI USA, 1–15. <https://doi.org/10.1145/3613904.3641895>
- [42] Isto Huvila. 2008. Participatory archive: Towards decentralised curation, radical user orientation, and broader contextualisation of records management. *Archival Science* 8 (2008), 15–36. <https://doi.org/10.1007/s10502-008-9071-0>
- [43] Daphne Ippolito, Florian Tramer, Milad Nasr, Chiyuan Zhang, Matthew Jagielski, Katherine Lee, Christopher Choquette Choo, and Nicholas Carlini. 2023. Preventing Generation of Verbatim Memorization in Language Models Gives a False Sense of Privacy. In *Proceedings of the 16th International Natural Language Generation Conference*, C. Maria Keet, Hung-Yi Lee, and Sina Zarrieß (Eds.). Association for Computational Linguistics, Prague, Czechia, 28–53. <https://doi.org/10.18653/v1/2023.inlg-main.3>
- [44] Daphne Ippolito, Ann Yuan, Andy Coenen, and Sehmon Burnam. 2022. Creative Writing with an AI-Powered Writing Assistant: Perspectives from Professional Writers. *arXiv:2211.05030* [cs]
- [45] Joel Jang, Dongkeun Yoon, Sohee Yang, Sungmin Cha, Moontae Lee, Lajanugen Logeswaran, and Minjoon Seo. 2023. Knowledge Unlearning for Mitigating Privacy Risks in Language Models. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Anna Rogers, Jordan Boyd-Graber, and Naoaki Okazaki (Eds.). Association for Computational Linguistics, Toronto, Canada, 14389–14408. <https://doi.org/10.18653/v1/2023.acl-long.805>
- [46] Harry H. Jiang, Lauren Brown, Jessica Cheng, Mehtab Khan, Abhishek Gupta, Deja Workman, Alex Hanna, Johnathan Flowers, and Timnit Gebru. 2023. AI Art and Its Impact on Artists. In *Proceedings of the 2023 AAAI/ACM Conference on AI, Ethics, and Society*. ACM, Montréal QC Canada, 363–374. <https://doi.org/10.1145/3600211.3604681>
- [47] Eun Seo Jo and Timnit Gebru. 2020. Lessons from archives: strategies for collecting sociocultural data in machine learning. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency (Barcelona, Spain) (FAT\* '20)*. Association for Computing Machinery, 306–316. <https://doi.org/10.1145/3351095.3372829>
- [48] Kowe Kadoma, Marianne Aubin Le Quere, Xiyu Jenny Fu, Christin Munsch, Danaë Metaxa, and Mor Naaman. 2024. The Role of Inclusion, Control, and Ownership in Workplace AI-Mediated Communication. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*. ACM, Honolulu HI USA, 1–10. <https://doi.org/10.1145/3613904.3642650>
- [49] Nikhil Kandpal, Eric Wallace, and Colin Raffel. 2022. Deduplicating training data mitigates privacy risks in language models. In *International Conference on Machine Learning*. PMLR, 10697–10707. <https://doi.org/10.48550/arxiv.2202.06539>
- [50] Hasindu Kariyawasam, Amashi Niwarthana, Alister Palmer, Judy Kay, and Anusha Withana. 2024. Appropriate Incongruity Driven Human-AI Collaborative Tool to Assist Novices in Humorous Content Generation. In *Proceedings of the 29th International Conference on Intelligent User Interfaces*. ACM, Greenville SC USA, 650–659. <https://doi.org/10.1145/3640543.3645161>
- [51] Shamika Klassen, Joanna Judith Elizabeth Mendy, Mikayla Buford, and Casey Fiesler. 2024. Black to the Future-The Power of Designing Afrofuturist Technology with Black Women, Femmes, and Non-Binary People. In *Proceedings of the 2024 ACM Designing Interactive Systems Conference*. 2156–2172.
- [52] Charlotte Kobiella, Yarhy Said Flores López, Franz Waltenberger, Fiona Draxler, and Albrecht Schmidt. 2024. "If the Machine Is As Good As Me, Then What Use Am I?" – How the Use of ChatGPT Changes Young Professionals' Perception of Productivity and Accomplishment. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*. ACM, Honolulu HI USA, 1–16. <https://doi.org/10.1145/3613904.3641964>
- [53] Gregory Kang Ruey Lau, Xinyuan Niu, Hieu Dao, Jiangwei Chen, Chuan-Sheng Foo, and Bryan Kian Hsiang Low. 2024. Protecting Text IP in the Era of LLMs with Robust and Scalable Watermarking. *ICML Workshop on Generative AI + Law* (2024).
- [54] Teven Le Scao, Angela Fan, Christopher Akiki, Ellie Pavlick, Suzana Ilić, Daniel Hesslow, Roman Castagné, Alexandra Sasha Luccioni, François Yvon, Matthias Gallé, et al. 2023. Bloom: A 176b-parameter open-access multilingual language model. (Nov. 2023). <https://inria.hal.science/hal-03850124> working paper or preprint.
- [55] Katherine Lee, Daphne Ippolito, Andrew Nystrom, Chiyuan Zhang, Douglas Eck, Chris Callison-Burch, and Nicholas Carlini. 2022. Deduplicating Training Data Makes Language Models Better. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Association for Computational Linguistics, 8424–8445. <https://doi.org/10.18653/v1/2022.acl-long.577>
- [56] Mina Lee, Katy Ilonka Gero, John Joon Young Chung, Simon Buckingham Shum, Vipul Raheja, Hua Shen, Subhashini Venugopalan, Thiemo Wambsganss, David Zhou, Emad A. Alghamdi, Tal August, Avinash Bhat, Madiha Zahrah Choksi, Senjuti Dutta, Jin L.C. Guo, Md Naimul Hoque, Yewon Kim, Simon Knight, Seyed Parsa Neshaei, Antonette Shibani, Disha Shrivastava, Lila Shroff, Agnia Sergeyuk, Jessi Stark, Sarah Sterman, Sitong Wang, Antoine Bosselut, Daniel Buschek, Joseph Chee Chang, Sherol Chen, Max Kreminski, Joonsuk Park, Roy Pea, Eugenia Ha Rim Rho, Zejiang Shen, and Pao Siangliulue. 2024. A Design Space for Intelligent and Interactive Writing Assistants (*CHI '24*). Association for Computing Machinery, New York, NY, USA, Article 1054, 35 pages. <https://doi.org/10.1145/3613904.3642697>
- [57] Mina Lee, Percy Liang, and Qian Yang. 2022. CoAuthor: Designing a Human-AI Collaborative Writing Dataset for Exploring Language Model Capabilities. In *CHI Conference on Human Factors in Computing Systems*. ACM, New Orleans LA USA, 1–19. <https://doi.org/10.1145/3491102.3502030>
- [58] Margaret Li, Suchin Gururangan, Tim Dettmers, Mike Lewis, Tim Althoff, Noah A. Smith, and Luke Zettlemoyer. 2022. Branch-train-merge: Embarrassingly parallel training of expert language models. *arXiv preprint arXiv:2208.03306* (2022).
- [59] Chumeng Liang, Xiaoyu Wu, Yang Hua, Jiaru Zhang, Yiming Xue, Tao Song, Zhengui Xue, Ruhui Ma, and Haibing Guan. 2023. Adversarial example does good: Preventing painting imitation from diffusion models via adversarial examples. *arXiv preprint arXiv:2302.04578* (2023).
- [60] AI@Meta Llama Team. 2024. The Llama 3 Herd of Models. *arXiv:2407.21783* [cs.AI] <https://arxiv.org/abs/2407.21783>
- [61] Shayne Longpre, Robert Mahari, Anthony Chen, Naana Obeng-Marnu, Damien Sileo, William Brannon, Niklas Muennighoff, Nathan Khazam, Jad Kabbara, Kartik Perisetla, et al. 2024. A large-scale audit of dataset licensing and attribution in AI. *Nature Machine Intelligence* 6, 8 (2024), 975–987.
- [62] Shayne Longpre, Robert Mahari, Ariel Lee, Campbell Lund, Hamidah Odeinwale, William Brannon, Nayan Saxena, Naana Obeng-Marnu, Tobin South, Cole Hunter, Kevin Klyman, Christopher Klam, Hailey Schoelkopf, Nikhil Singh, Manuel Cherep, Ahmad Anis, An Dinh, Caroline Chitongo, Da Yin, Damien Sileo, Deividás Matciunas, Diganta Misra, Emad Alghamdi, Enrico Shippole, Jianguo Zhang, Joanna Materzynska, Kun Qian, Kush Tiwary, Lester Miranda, Manan Dey, Minnie Liang, Mohammed Hamdy, Niklas Muennighoff, Seonghyeon Ye, Seungone Kim, Shrestha Mohanty, Vipul Gupta, Vivek Sharma, Vu Minh Chien, Xuhui Zhou, Yizhi Li, Caiming Xiong, Luis Villa, Stella Biderman, Hanlin Li, Daphne Ippolito, Sara Hooker, Jad Kabbara, and Sandy Pentland. 2024. Consent in Crisis: The Rapid Decline of the AI Data Commons. *arXiv:2407.14933* [cs.CL] <https://arxiv.org/abs/2407.14933>
- [63] Shayne Longpre, Gregory Yauney, Emily Reif, Katherine Lee, Adam Roberts, Barret Zoph, Denny Zhou, Jason Wei, Kevin Robinson, David Mimmo, and Daphne Ippolito. 2024. A Pretrainer's Guide to Training Data: Measuring the Effects of Data Age, Domain Coverage, Quality, & Toxicity. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, Kevin Duh, Helena Gomez, and Steven Bethard (Eds.). Association for Computational Linguistics, Mexico City, Mexico, 3245–3276. <https://doi.org/10.18653/v1/2024.naacl-long.179>
- [64] Juniper Lovato, Julia Witte Zimmerman, Isabelle Smith, Peter Dodds, and Jennifer L. Karson. 2024. Foregrounding Artist Opinions: A Survey Study on Transparency, Ownership, and Fairness in AI Generative Art. *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society* 7, 1 (Oct. 2024), 905–916. <https://doi.org/10.1609/aiies.v7i1.31691>
- [65] Alexandra Luccioni and Joseph Viviano. 2021. What's in the Box? An Analysis of Undesirable Content in the Common Crawl Corpus. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*. Association for Computational Linguistics, 182–189. <https://doi.org/10.18653/v1/2021.acl-short.24>
- [66] Nik Marda, Jasmine Sun, and Mark Surman. 2024. *Public AI: Making AI Work for Everyone, by Everyone*. Technical Report. Mozilla.
- [67] Sewon Min, Suchin Gururangan, Eric Wallace, Hannaneh Hajishirzi, Noah A. Smith, and Luke Zettlemoyer. 2023. SILO Language Models: Isolating Legal Risk in a Nonparametric Datastore. *arXiv:2308.04430* [cs]
- [68] Piotr Mirowski, Kory W. Mathewson, Jaylen Pittman, and Richard Evans. 2023. Co-Writing Screenplays and Theatre Scripts with Language Models: Evaluation by Industry Professionals. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (Hamburg, Germany) (*CHI '23*). Association for Computing Machinery, New York, NY, USA, Article 355, 34 pages. <https://doi.org/10.1145/3555455.3555455>

- org/10.1145/3544548.3581225
- [69] Shakir Mohamed, Marie-Therese Png, and William Isaac. 2020. Decolonial AI: Decolonial Theory as Sociotechnical Foresight in Artificial Intelligence. *Philosophy & Technology* 33, 4 (Dec. 2020), 659–684. <https://doi.org/10.1007/s13347-020-00405-8>
- [70] Michael J Muller and Sandra Kogan. 2012. Grounded Theory Method in Human-Computer Interaction and Computer-Supported Cooperative Work. In *Human Computer Interaction Handbook*. Routledge, 1003–1023.
- [71] Matt O'Brien. 2023. ChatGPT-maker OpenAI signs deal with AP to license news stories. *The Associated Press* (Jul 2023). <https://apnews.com/article/openai-chatgpt-associated-press-ap-f86f84c5bcc2f3b98074b38521f5f75a/>
- [72] U.S. Copyright Office. [n. d.]. Artificial Intelligence Study. <https://www.copyright.gov/policy/artificial-intelligence/>. Accessed: 2024-09-11.
- [73] Barbara Prainsack. 2019. Logged out: Ownership, exclusion and public value in the digital data and information commons. *Big Data & Society* 6, 1 (2019), 2053951719829773.
- [74] Heila Preceel, Allison McDonald, Brent Hecht, and Nicholas Vincent. 2024. A Canary in the AI Coal Mine: American Jews May Be Disproportionately Harmed by Intellectual Property Dispossession in Large Language Model Training. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*. ACM, Honolulu HI USA, 1–19. <https://doi.org/10.1145/3613904.3642749>
- [75] Jipeng Qiang, Shiyu Zhu, Yun Li, Yi Zhu, Yunhao Yuan, and Xindong Wu. 2023. Natural language watermarking via paraphraser-based lexical substitution. *Artificial Intelligence* 317 (2023), 103859.
- [76] Hua Xuan Qin, Shan Jin, Ze Gao, Mingming Fan, and Pan Hui. 2024. CharacterMeet: Supporting Creative Writers' Entire Story Character Construction Processes Through Conversation with LLM-Powered Chatbot Avatars. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*. ACM, Honolulu HI USA, 1–19. <https://doi.org/10.1145/3613904.3642105>
- [77] Alex Reisner. 2023. Revealed: The Authors Whose Pirated Books Are Powering Generative AI. *The Atlantic* (Aug 2023). <https://www.theatlantic.com/technology/archive/2023/08/books3-ai-meta-llama-pirated-books/675063/>
- [78] Nithya Sambasivan, Shivani Kapania, Hannah Highfill, Diana Akrong, Praveen Paritosh, and Lora M Aroyo. 2021. "Everyone wants to do the model work, not the data work": Data Cascades in High-Stakes AI. In *proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. 1–15. <https://doi.org/10.1145/3411764.3445518>
- [79] Shawn Shan, Jenna Cryan, Emily Wenger, Haitao Zheng, Rana Hanocka, and Ben Y Zhao. 2023. Glaze: Protecting artists from style mimicry by {Text-to-Image} models. In *32nd USENIX Security Symposium (USENIX Security 23)*. 2187–2204.
- [80] Luca Soldaini, Rodney Kinney, Akshita Bhagia, Dustin Schwenk, David Atkinson, Russell Authur, Ben Bogin, Khyathi Chandu, Jennifer Dumas, Yanai Elazar, et al. 2024. Dolma: An open corpus of three trillion tokens for language model pretraining research. *arXiv preprint arXiv:2402.00159* (2024).
- [81] Shannon Vallor. 2024. The Danger Of Superhuman AI Is Not What You Think. *Noema Magazine* (May 2024). <https://www.noemamag.com/the-danger-of-superhuman-ai-is-not-what-you-think/>
- [82] Thanh Van Le, Hao Phung, Thuan Hoang Nguyen, Quan Dao, Ngoc N Tran, and Anh Tran. 2023. Anti-dreambooth: Protecting users from personalized text-to-image synthesis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 2116–2127.
- [83] Jessica Vitak, Katie Shilton, and Zahra Ashktorab. 2016. Beyond the Belmont principles: Ethical challenges, practices, and beliefs in the online data research community. In *Proceedings of the 19th ACM conference on computer-supported cooperative work & social computing*. 941–953.
- [84] Nikhil Vyas, Sham M Kakade, and Boaz Barak. 2023. On provable copyright protection for generative models. In *International Conference on Machine Learning*. PMLR, 35277–35299.
- [85] KiYoon Yoo, Wonhyuk Ahn, Jiho Jang, and Nojun Kwak. 2023. Robust Multi-bit Natural Language Watermarking through Invariant Features. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Anna Rogers, Jordan Boyd-Graber, and Naoaki Okazaki (Eds.). Association for Computational Linguistics, Toronto, Canada, 2092–2115. <https://doi.org/10.18653/v1/2023.acl-long.117>

## A Self-Reported Demographic Data from Interviewed Writers

We spoke with 33 writers who write across a variety of genres. The majority of writers we spoke with (22 out of 33) write in multiple genres, as indicated in Table 2.

The writers we spoke with also share their work via a variety of publishing or delivery methods, detailed in Table 3. As with genre,

**Table 2: Writing Genre of Interviewees**

Genre	
<i>Fiction</i>	22
Literary Fiction	10
Genre Fiction	9
Sci-Fi and Fantasy	5
Romance	3
Fan Fiction	5
Young Adult (YA) Fiction	3
<i>Non-Fiction</i>	14
Personal Essay and Memoir	11
Journalism	5
<i>Poetry</i>	15

the majority of writers we spoke with (18 out of 33) also share their writing through multiple venues and means.

**Table 3: Method of Delivery for Interviewees**

Publishing or Delivery Method	
<i>Traditional</i>	26
Major Publisher ("Big 5")	9
Independent Publishers, Literary Magazines, and Journals	21
<i>Self-Published</i>	15
Blogs and Forums	9
Kindle Direct or equivalent	6
<i>Performance</i>	7
Readings	5
Installations and Exhibits	3

The degree to which writers were compensated for their creative work varied significantly, detailed in Table 4.

**Table 4: Degree of Professionalization for Interviewees**

Compensation for Creative Work	
Most income comes from creative writing practice	11
Most income comes from writing-related profession	10
Most income does not come from writing	8
Unknown	4

We also noted how writers self-reported their use of LLMs; see Table 5. The vast majority (23 out of 33) were "dabblers"—those who had tried out LLMs, but do not use them in a regular or intensive manner. A significant minority (8 out of 33) use LLMs regularly in either their paratextual or central creative writing practices.

**Table 5: Engagement with LLMs in Writing Practice**

Engagement with LLMs in Writing Practice	
None	3
Some (e.g., "trying it out" or "dabbling")	23
Regular use in peripheral writing and publishing activities (e.g., marketing, paratextual use)	3
Regular use in creative composition (e.g., brainstorming, narrative outline, turns of phrase, etc.)	5

## B Interview Guideline

The following guideline represents our final interview guideline. It includes core questions (solid bullet points) and potential follow-up questions (dashed bullet points). The guideline acted as a structure for the interview; not all interviews went through all follow-up questions or went through the questions in this precise order, depending on how the conversation was progressing. Note that two questions use a stimulus, which was either verbally or visually presented to interviewees to aid in the conversation.

### B.1 General writing questions

- Could you briefly describe your writing education or journey?
  - What initially inspired you to become a writer? Can you recall a specific moment or influence that ignited your passion for writing?
  - How many years have you been writing as a job?
  - How has your writing style or focus evolved over the course of your career?
- Could you briefly describe your current writing practice?
  - What is your typical writing process?
  - What genre or type of writer do you identify with most?
  - How do you seek and incorporate feedback on your work? Are there specific individuals or groups you rely on for constructive criticism?
  - Is writing your main profession, or do you have other jobs? How do you divide your time?
- Do you share or publish your writing?
  - What is your approach to sharing or publishing your work? Who is involved in the process?
  - Has it changed over time? How?

### B.2 Knowledge of and attitudes towards LLMs

- How much do you know about large language models like ChatGPT or Gemini (formerly Bard)?
  - How would you describe LLMs to someone who doesn't know what it is?
- Have you ever used these models?
  - What did you do or try?
  - What did you like or dislike about them?
  - Are there particular applications (perhaps that you see others engaged in) you like or dislike?

- Have you ever used these in your writing process? In what ways?
  - If yes: Has your opinion on LLMs now changed since you first used one?
  - What do you like or not like about using them? Why?
  - Either yes/no: Are there specific aspects of your writing practice where you see these models being relevant or useful?
- Are you interested in using them in the future?
  - If yes: in what stage of the writing process would you use them?
  - Are there applications that you particularly like or dislike?
  - If no: why not?
- Have you thought much about the training data for these models?
  - What are your thoughts on the source of the training data for these models?

### B.3 Knowledge of and attitudes towards language model training data

- Do you know if any of your writing was used to train a large language model?
  - If yes: How did you find out?
  - If don't know: Have you wondered? Why or why not?
- How do you feel about the actual, potential, or hypothetical inclusion of your writing in an existing language model?
  - Are there details about the use of the model that matter to you?
  - **Note: Use stimulus 1**
- Under what conditions, if any, would you want or consent to your own writing be included in an LLM training dataset?
  - Why do you prefer [x] over [y]?
  - **Note: Use stimulus 2**
  - Can you walk me through your thought process in your reactions?
  - If all your conditions were met, would you want your writing included? Why or why not?
  - If compensation is desired: What if compensation is very low? What about other forms of compensation, like stock options, or money going towards professional organizations or scholarships?
- How might these conditions differ for different kinds of writing? e.g. Published v. unpublished work, fiction v. nonfiction. Why?
- What material ways can you imagine your writing career being affected by LLMs?
  - Is there anything about the use of your writing data in particular that contributes to these material effects? (Versus LLMs generally.)
- Are there any upsides of your work being included? e.g. For people to search, summarize, get recommendations.
- What institutions do you think should be collecting and protecting writing data? Or collecting opt-in/opt-out information? e.g. Government labs, universities, professional organizations like the Author's Guild, non-profits.

- Who do you believe would advocate best for your interests? e.g. agents, publishers, lawyers, non-profits.
  - Who do you trust to advocate for your interests?
  - Who do you not trust?
- If you were to sign a contract with LLM provisions, how would you negotiate?

#### B.4 Stimulus 1: Hypothetical usage scenarios

Commercial interests of the organization that developed the model:

- You are notified that several of your books were used by an LLM developed by a for-profit company. Your books are unavailable online for free. The company is charging a monthly subscription for access to their model.
- Contrast A (research purposes): The LLM was developed by a university or government lab for research purposes only (i.e. non-commercial).
- Contrast B (non-profit): The LLM was developed by a non-profit dedicated to providing LLM access globally for free.
- Contrast C (writing was available online): Instead of your books, the LLM used a series of short stories you wrote which are available online for free, although you were paid for these stories by the publisher and signed a contract saying you still own the rights to the stories.

How people are using the model:

- You are notified that several of your books were used to train an LLM. People are using this LLM primarily to help with business-focused writing tasks like writing cover letters, business emails, or press releases.
- Contrast A (creative writing): People are using the LLM for creative writing like writing stories, personal essays, or poems.
- Contrast B (pro-social applications): People are using the LLM in professional contexts like lawyers using it to streamline support in pro bono cases, doctors using it to support diagnosis of rare diseases, teachers using it to help teach difficult concepts (like chemistry) or scientists to make new discoveries.

#### B.5 Stimulus 2: Conditions of consent

Compensation:

- fixed, one-off payment for a specific contribution
- ongoing compensation or royalties entail recurring payments based on the use of the model (i.e. like royalties for book sales)
- ongoing compensation based on usage of the model related to your work; e.g. if someone uses a model for their math homework, you don't get compensated, but if it's used for creative writing, you do get compensated.

Credit:

- public record of contributors listed on the LLM platform
- Incorporate a mechanism for the model to acknowledge the source of its training data when generating outputs. (i.e. citation)
- List contributors as collaborative authors on associated publications or reports

Use of model:

- usage only in education settings
- disallow usage that would compete with your interests (e.g. disallow use to write novels that could theoretically compete with your novel)

Commercial:

- usage only in non-commercial settings (e.g. research, art, or non-profit)
- require compensation in commercial applications

Personal access to model:

- Exclusive access to the model
- free access to the model