

U C B E R K E L E Y

C E N T E R F O R L O N G - T E R M C Y B E R S E C U R I T Y



C L T C W H I T E P A P E R S E R I E S

# Resistance to Text-to-Image Generators in Creator Communities

JANIYA PETERS, BONWOO KUH, ISABEL LI



CLTC WHITE PAPER SERIES

# Resistance to Text-to-Image Generators in Creator Communities

JANIYA PETERS, BONWOO KUH\*, ISABEL LI\*

*\*Authors contributed equally.*

October 2024



# Contents

## **EXECUTIVE SUMMARY 1**

## **INTRODUCTION 2**

Precedents in U.S. Copyright Law 3

Commercial Trade Regulations for Generative Artificial Intelligence 6

Centering Resistance in Legal Remedies 9

## **METHODOLOGY 11**

## **INTERVIEW THEMATIC ANALYSIS 13**

Theme 1: Artistic modalities and style preferences shape attitudes towards text-to-image generators. 13

*Creative Workflows* 13

*Creative Friction* 13

*Unfamiliarity with AI* 14

Theme 2: Resistance accesses data narratives and protections to protect work from AI training and mimicry, and align interactions with their preferred use contexts. 15

*Use Contexts* 15

*Artistic Appropriation* 16

*Resistance* 17

*Data Provenance* 19

Theme 3: Resistance invokes data narratives and protections where copyright law fails to enforce itself. 20

*Employment* 20

*Intellectual Property Rights* 21

## **DISCUSSION 24**

## **CONCLUSION 29**

Recommendations 29

## **ACKNOWLEDGMENTS 31**

## **ABOUT THE AUTHORS 31**



## Executive Summary

The commercial development of generative artificial intelligence (Gen AI) systems has caused critical discussions in creator communities. Questions have particularly emerged around how AI models are trained with copyright-protected materials produced by authors, photographers, visual artists, and other creators. Several civil suits in the United States and United Kingdom have been filed claiming copyright infringement against Gen AI developers. Federal and international trade legislation, including the landmark European Union Artificial Intelligence Act, may provide a pathway to audit emergent Gen AI technologies, and provide remedies for misuse of copyrighted work.

This report presents findings from a survey of creative professionals who have resisted the integration of text-to-image generators in their creative practice, with a summary of their mechanisms and aims for resistance. Resistance in this context is defined as the refusal to engage, participate, or contribute to a technological system, and the subversion of that system through use of medium, artistic techniques, and workflows. By studying resistance as a site of dispute and value differentiation, this report cites breakdowns between technological implementations and expectations, and helps illuminate how policy might facilitate alignment between creative practitioners and their value preferences.

This report also makes several recommendations for policymakers and other stakeholders to improve copyright and data management, and to help artists identify misuse of their work in generative AI systems. These include:

- Adoption of H.R. 7913 the Generative AI Copyright Disclosure Act of 2024, a Congressional bill that would obligate AI developers to disclose copyrighted materials used in training sets to the Register of Copyrights;
- A survey addendum by the U.S. Copyright Office and online service providers to take-down notice procedures documenting misuse of text-to-image generators;
- Default opt-out policies on creative software programs, social media, and image distribution platforms that **restrict** the use of artists' works for AI training;
- Enhancement of working groups, research labs, and educational seminars led by artists, sponsored by the Library of Congress; and
- Interface options on creative software programs, social media, and image distribution platforms that delineate and separate generative artificial intelligence features from primary services.

# Introduction

Advancements in generative artificial intelligence systems have placed visual artists, image copyright holders, and image subjects in a vulnerable position. What were once rudimentary computer programs capable only of producing simple black-and-white forms,<sup>1</sup> generative AI systems can now produce rather complex imagery, forms, and styles that human users direct through text-based prompts.

Since OpenAI's 2021 release of DALL-E, competitors such as Midjourney and Adobe have released their own commercial models, and Stable Diffusion has released a free open source model that allows users to modify the source code. These technologies are enabled by deep learning models, including *general adversarial networks (GANs)* and *latent diffusion*. Training images are tagged with CLIP embeddings, or text embeddings, which represent the content of the image. Subsequently, users pose text prompts to the generator to produce context-specific images.

Developers of generative AI tools have received criticism for openly scraping the Internet to train their deep learning models without regard to copyright or content filtering. Web crawlers robustly and indiscriminately collect images to compose reference databases for researchers and commercial developers; the most notable of these databases is the LAION-5B, comprising 5.8 billion text-image pairs sourced from Common Crawl technology. Alternatively, some platforms have sourced images from their own repositories, but this also provokes concern. Adobe customers claim the company trained its generative image program, Adobe Firefly, on images in the Adobe Stock image library without clarifying this usage to Stock contributors.

Sourcing images from both free and open platforms, as well as paid services, poses conflicts to several U.S. intellectual property and user-consumer guidelines. Despite the potential benefits of generative image programs in creative production, including quick iteration and modification of ideas, the training history and economic impacts of these tools remain a critical divider in creator communities. This paper offers a thematic analysis of resistance to text-to-image generators in creator communities that may help shape protocols for U.S. copyright law and set standards for creative production.

<sup>1</sup> Chris Garcia, "Harold Cohen and AARON—A 40-Year Collaboration," Computer History Museum, August 23, 2016, <https://computerhistory.org/blog/harold-cohen-and-aaron-a-40-year-collaboration/>.



## PRECEDENTS IN U.S. COPYRIGHT LAW

The U.S. Copyright Office defines “human authorship” as creative literary, artistic, or musical works that have been significantly developed by a human originator. Human authorship is a necessary condition for retaining “copyright,” or the exclusive right to copy, reproduce, and sell one’s works within a protected time span. In the United States, this protection lasts for the duration of an author’s life, plus 70 years. Indeed, many individuals have authored works with the assistance of technical mechanisms — for example, typewriters, word processors, and drawing software — and were able to retain copyright ownership over their unique contributions.

However, the introduction of generative AI technologies such as text-to-image generators has challenged delineations between human-authored and computer-generated contributions. The U.S. Copyright Office has concluded that:

1. **Outputs derived primarily from computer programs do not qualify as “work made for hire.”**<sup>2</sup> In 2022, the Office determined that a computer algorithm that autonomously produced a 2D composition *did not* qualify as a “worker” for the copyright filer, and therefore its outputs could not be protected as “work made for hire.”<sup>3</sup> “Work made for hire” includes labor conducted under contract between an employee and employer. The decision reiterated that the computer algorithm was not an employee, nor could it enter into a contract with the copyright filer.
2. **Individuals may be able to copyright curations of human-authored and computer-generated works.** In 2023, the Office determined that Kristina Kashtanova could not copyright the AI-generated images in her graphic novel.<sup>4</sup> However, the Office protected Kashtanova’s authorship of the “overall selection, coordination, and arrangement of the text and visual elements” within the novel. The decision differentiated between the work’s curation (human-authored) and the images themselves (produced by Midjourney).

2 U.S. Copyright Law §101 defines a “work made for hire” as “work prepared by an employee within the scope of his or her employment” and/or “a work specially ordered or commissioned for use as a contribution to a collective work, as a part of a motion picture or other audiovisual work, as a translation, as a supplementary work, as a compilation, as an instructional text, as a test, as answer material for a test, or as an atlas, if the parties expressly agree in a written instrument signed by them that the work shall be considered a work made for hire.” This exception entails a written contract that work produced is owned by an employer or hiring body.

3 “Re: Second Request for Reconsideration for Refusal to Register A Recent Entrance to Paradise,” February 14, 2022, Copyright Review Board, <https://www.copyright.gov/rulings-filings/review-board/docs/a-recent-entrance-to-paradise.pdf>.

4 “Re: Zarya of the Dawn,” February 21, 2023, United States Copyright Office, <https://www.copyright.gov/docs/zarya-of-the-dawn.pdf>.

RESISTANCE TO TEXT-TO-IMAGE GENERATORS  
IN CREATOR COMMUNITIES

3. **More than minimal AI-generated contributions should be disclosed or excluded from registration.** The Office advises that works comprising more than *de minimis* [insignificant, minor] AI-generated content be disclosed,<sup>5</sup> or risk being excluded<sup>6</sup> from registration. The rule requires the registrant’s discretion in distinguishing human contributions, which are copyrightable, from AI-generated contributions, which are not. It also requires the registrant to determine if an AI-generated contribution is significant enough to be disclosed. For example, an entirely computer-generated 2D composition requires disclosure, but singular elements (e.g., a Photoshopped background object) might not require mentioning. Application of *de minimis* will vary on a case-by-case basis, according to the interpretations of the registrant and the Office.

In addition to the precedents already established by the U.S. Office of Copyright, several lawsuits have interrogated principles of authorship, similarity, and fair use in relation to AI-generated works. In January 2023, a coalition of professional visual artists sued DeviantArt, Midjourney, and Stability AI for the *scraping* of uploaded works for training purposes without express consent; *duplication* of uploaded works; *deriving* AI outputs from uploaded works; and establishment of *unfair competition* between plaintiffs and users of DeviantArt’s image generator, DreamUp. Plaintiffs identified that their works were used in defendants’ training databases without permission, and requested that their names be removed as potential keywords for image prompts. However, the case was dismissed due to lack of “substantial similarity” evidence between the Plaintiffs’ works and outputs of the Defendants’ products.<sup>7</sup> The outcome indicates a barrier for artists seeking to prove copyright infringement, given the few auditing tools available and limited access to the back-end operations of generative image models. Moreover, identifying individual cases of infringement can be daunting and costly, especially if the enabling technology and distribution are faster than the rate at which a user can identify and report each misuse.

In 2023, stock image vendor Getty Images pursued civil action against Stability AI in the United States and United Kingdom for scraping its 12 million licensed, trademarked photos. The development of these cases challenges copyright law in multiple jurisdictions,<sup>8</sup> but Stability AI must

5 “Artificial Intelligence and Copyright,” 59942 Federal Register, Vol. 88, No. 167, August 30, 2023, United States Copyright Office, <https://www.copyright.gov/ai/docs/Federal-Register-Documents-Artificial-Intelligence-and-Copyright-NOI.pdf>.

6 “Copyright Registration Guidance: Works Containing Material Generated by Artificial Intelligence,” 16190 Federal Register, Vol. 88, No. 51, March 16, 2023, United States Copyright Office, [https://copyright.gov/ai/ai\\_policy\\_guidance.pdf](https://copyright.gov/ai/ai_policy_guidance.pdf).

7 Owen Wolfe, Lauren Gregory Leipold, Grayson Moronta. “Some Stability For AI Defendants: Judge Dismisses All But One Claim in Andersen et. al., v. Stability AI LTD., et. al.,” Seyfarth Shaw LLP, November 9, 2023, <https://www.lexology.com/library/detail.aspx?g=fb25f1cc-6879-489f-bf4c-c46f80134eac>.

8 Copyright, Designs, and Patents Act 1988, Section 9(3).

RESISTANCE TO TEXT-TO-IMAGE GENERATORS  
IN CREATOR COMMUNITIES

still address alleged violations of the UK’s Copyright, Designs, and Patent Acts (CDPA) statutes, including restricted copying (Section 17) and secondary infringement (Section 22, 23), if such activity occurred during training within the jurisdiction. Getty Images has strong evidence of infringement, including outputs produced by Stable Diffusion technology that contain Getty Images’ trademarked logo. Even so, the media company might not represent the interests of smaller creative professionals and collectives. In October 2022, Getty Images partnered with BRIA to create its own generative AI image tools.<sup>9</sup> It may be concluded that Getty Images’ complaint lies primarily on the Defendant’s failure to adhere to any licensing agreement, *not* the application of image generation alone. Individual creative professionals who resist generative image models might not identify with the ultimate goals of Getty Images, and may have fewer avenues for compensation if their work was used to train these models.

**Table 1: Status of Civil Suits against Generative Artificial Intelligence Image Companies, as of September 2024**

	<b>Andersen et al v. Stability AI, Ltd. et al</b> <sup>10</sup>	<b>Getty Images v. Stability AI, Ltd. (filed in U.S.)</b>	<b>Getty Images v. Stability AI (filed in UK)</b>
<b>Court</b>	United States District Court for the Northern District of California	United States District Court for the District of Delaware	High Court of Justice in London
<b>Complaint Filed</b>	January 13, 2023	February 3, 2023	January 12, 2023 (date of Approved Judgment)
<b>Claim(s)</b>	17 U.S.C. § 501; 17 U.S.C. §§ 1201–1205; Cal. Civ. Code section 3344; Cal. Bus. & Prof. Code §§ 17200, et seq.	17 U.S.C. §101, et seq; 17 U.S.C. § 1202(a) ; 17 U.S.C. § 1202(b) ; 15 U.S.C. § 1114(1); 15 U.S.C. § 1125(a); 15 U.S.C. § 1125(c); Delaware’s Uniform Deceptive Trade Practices Act; Delaware Trademark Act, Section 3313.	CDPA, Section 16(2); CDPA, Section 17; CDPA, Section 22; CDPA, Section 23; CDPA, Section 27; “Training and Development” in the UK; database right infringement; trademark infringement (“passing off”).
<b>Status</b>	Dismissed with leave to amend on October 30, 2023.	Proceeding with jurisdictional discovery.	Proceeding to trial.

Copyright typically serves to balance the exclusive rights of the authors with the public’s interest to engage new ideas and creative works. *Fair use* permits the sampling of copyrighted material for academic research and other uses, towards the advancement of public knowledge

<sup>9</sup> “BRIA Partners with Getty Images to Transform Visual Content Through Responsible AI,” October 25, 2022, Getty Images, <https://investors.gettyimages.com/news-releases/news-release-details/bria-partners-getty-images-transform-visual-content-through>.

<sup>10</sup> Plaintiffs: Sarah Andersen, Kelly McKernan, Karla Ortiz. Defendants: Stability AI, Ltd., Stability AI Inc., Midjourney Inc., Deviant Art Inc.

## RESISTANCE TO TEXT-TO-IMAGE GENERATORS IN CREATOR COMMUNITIES

and accessibility of scholarly materials. For this reason, text and data mining activities are often protected under *fair use* (*Authors Guild, Inc. v. Google, Inc.*, 2015; *Authors Guild, Inc. v. HathiTrust*, 2014). However, state proceedings have limited *fair use* as a defense where circumvention mechanisms are used to gain an unfair market advantage or avoid paying for licensing agreements (*Associated Press v. Meltwater U.S. Holdings, Inc.*, 2013). DMCA §512 outlines safe harbors for online service providers given they comply with removal of infringing material (notice and takedown). However, text-to-image generation complicates these protocols in that artists may identify their work and influence embedded throughout the models' internal structure, including its training images, text embeddings, and prompt options. While copyright provides one framework for addressing these emergent issues, other forms of governance are required to supplement this approach.

### COMMERCIAL TRADE REGULATIONS FOR GENERATIVE ARTIFICIAL INTELLIGENCE

On October 4, 2023, the Federal Trade Commission (FTC) hosted several unions and organizations — including the Writers Guild of America West (WGA), Screen Actors Guild-American Federation of Televisions and Radio Artists (SAG-AFTRA), and Concept Art Association — to discuss the impact of Gen AI systems on creative production. Overwhelmingly, creative professionals called for legal protections and technical infrastructure that permit greater control over their artifacts and labor. “Consent, credit, compensation, control,” stated Steven Zapata, a representative of the Concept Art Association. “This is what creators reasonably seek in this new era where our work will be used to add tremendous value to these new technologies.” Umair Kazi, the Policy and Advocacy Director of The Authors Guild, outlined an approach rooted in fair compensation for creators whose works are used to train Gen AI systems, and whose names and work titles are incorporated into prompt engines. This approach also calls for permission pathways between creators and Gen AI systems seeking to train models on existing works; disclosure of data repositories used for training; and disclaimers to be included alongside content generated by artificial intelligence systems. FTC Commissioner Rebecca Slaughter stated:

“Many but not all states have laws that provide rights to publicity, which may provide avenues for legal protection and compensation, and as the chair noted, the FTCs prohibitions against unfair and deceptive practices and unfair methods of competition

## RESISTANCE TO TEXT-TO-IMAGE GENERATORS IN CREATOR COMMUNITIES

apply to applications of AI just as much as they have to every other new technology that's been introduced in the market over the last hundred years.”<sup>11</sup>

The Federal Trade Commission Act, or FTC Act (15 U.S. Code § 45), prevents businesses and services from engaging in unfair competition methods, including “unfair or deceptive acts or practices.” The Federal Reserve defines acts as “unfair” when they inflict injury onto consumers that the consumer cannot “reasonably” avoid, while a “deceptive” act is an omission or misrepresentation of information that might mislead consumers under reasonable interpretations.<sup>12</sup>

The FTC Act has significant implications for services attempting to conduct AI-related business in the U.S. and/or with U.S. customers. On January 25, 2024, the FTC issued compulsory notices to five major companies and cloud services requiring information on their partnerships and investments in generative artificial intelligence.<sup>13</sup> This audit required the companies to share details on the competitive impacts of generative AI, including its training data. Similarly, the European Commission has requested information from leading online service providers for information about products and services using generative artificial intelligence, including potential harms concerning false information and intellectual property violations. The Commission is empowered by the Digital Services Act (DSA), which requires online service providers to comply with requests on threat of fine.<sup>14</sup> However, the European Union Artificial Intelligence Act is the first comprehensive legislation to outline risks and regulations for generative artificial intelligence as it affects EU members. Article 53 1(d) mandates that AI developers provide summaries of training content, including copyrighted material, to increase transparency and accountability for data inputs.

11 “Creative Economy and Generative AI,” October 4, 2023, Federal Trade Commission, <https://www.ftc.gov/news-events/events/2023/10/creative-economy-generative-ai>.

12 “Federal Trade Commission Act Section 5: Unfair or Deceptive Acts or Practices,” (n.d.), Consumer Compliance Handbook, Federal Trade Commission Act, <https://www.federalreserve.gov/boarddocs/supmanual/cch/200806/ftca.pdf>.

13 “FTC Launches Inquiry into Generative AI Investments and Partnerships,” January 25, 2024, Federal Trade Commission, <https://www.ftc.gov/news-events/news/press-releases/2024/01/ftc-launches-inquiry-generative-ai-investments-partnerships>.

14 “Commission sends requests for information on generative AI risks to 6 Very Large Online Platforms and 2 Very Large Online Search Engines under the Digital Services Act,” March 14, 2024, European Commission, <https://digital-strategy.ec.europa.eu/en/news/commission-sends-requests-information-generative-ai-risks-6-very-large-online-platforms-and-2-very>.

RESISTANCE TO TEXT-TO-IMAGE GENERATORS  
IN CREATOR COMMUNITIES

**Table 2: How Commercial Trade Regulations Address Generative Artificial Intelligence**

Statute	Regulatory Agency	Summary
EU Artificial Intelligence Act, Article 53, 1(d)	European Commission	Requires providers of general purpose AI models to publish summaries of training content, including copyrighted material.
Digital Services Act, Article 7(2)	European Commission	Online service providers may be fined if they do not supply correct and complete information about business activity upon request.
15 U.S. Code § 46(b)	Federal Trade Commission	Requires businesses to disclose information regarding conduct, management, commercial activity, etc. (“Report of Persons, Partnerships, and Corporations”).
15 U.S. Code § 45	Federal Trade Commission	Prohibits businesses from engaging in unfair methods of competition, including unfair or deceptive acts.

Guidance from the European Commission and Federal Trade Commission will be critical for governing generative AI systems. First, the provision of information on datasets and data inputs will allow greater transparency between AI system developers, governing bodies, and the public. Moreover, the commissions may provide guidance for other governing bodies that do not possess the right to audit commercial developers, but would benefit from a public summary of datasets and copyrighted material in directing their own activities. With such information, commercial governing bodies may be able to address the ethical considerations concerning consumers’ interactions with AI products and services, including:

- The right for users and consumers to be informed or notified about how their data is being used towards generative artificial intelligence services, preferably *before* use.
- The right for users and consumers to withdraw their data from use in generative artificial intelligence services, including clear opt-in and opt-out protocols and clarification of default services, without penalty. This includes the disgorgement of data from user interfaces *and* internal models.
- Clarification of third-party partnerships and training protocols requiring user and consumer data in privacy policies and terms of use contracts.

A second consideration will be to review standards for “informed consent,” whereby users and consumers can *knowingly* and *voluntarily* provide data for emergent services, if they choose. Most critically, if they do not provide consent, services must be equipped to honor users’ and consumers’ requests.

## CENTERING RESISTANCE IN LEGAL REMEDIES

This paper defines **resistance** as the refusal to engage, participate, or contribute to a technological system; or foiling that system through changes in artistic medium, techniques, or workflows. This report studies creative professionals who resist text-to-image generators in their practice, including their mechanisms and aims for resistance. We define a *text-to-image generator* as a deep learning model that allows a user to specify a text prompt and produce an image according to the prompt’s specifications (e.g., style or subject matter).

It is pertinent to consider **resistance** to emergent technologies throughout history, both as an expression of values and a call for policy to mediate where necessary. In 2023, WGA and SAG-AFTRA went on strike, pausing studio productions to negotiate higher compensation and residual packages, as well as protections against displacement by artificial intelligence systems. This came as studios began to implement contractually required body scans and audio retention for screen actors, voice actors, and talent that could be reproduced for future uses. This indicates the first time that artificial intelligence has been a major decision point in a WGA or SAG-AFTRA strike, though not the first time a new media or technology has required negotiation. In 1980, SAG and AFTRA sought more revenue from media released on at-home videos and television packages. Today, the advent of chatbots and image generation in production, and the growth of streaming services as a major distribution source, have prompted similar questions about labor and compensation in creative industries.

There are several modes of **resistance** to technology. These past examples demonstrate the power struggle between developers and adopters of technological systems, and resisters to those systems.

- By the 1800s, the stocking frame had largely mechanized knitting and textile production. English textile workers dubbed “Luddites” entered factories and destroyed the stocking frames. The “Luddites” viewed the stocking frame as decreasing their wages and producing “inferior materials—especially hosiery articles that were flimsy or lost their shape.”<sup>15</sup> By damaging the stocking frames, the Luddites communicated their discontent with textile manufacturers.
- A new information system, software, or technology is introduced in the workplace. This system is designed to optimize a task and increase efficiency (e.g., a financial accounting software system). But ultimately, this new system redistributes power

<sup>15</sup> Kevin Binfield. (2015). *Writings of the Luddites*. Johns Hopkins University Press. Page 24.

## RESISTANCE TO TEXT-TO-IMAGE GENERATORS IN CREATOR COMMUNITIES

away from particular users (e.g., accountants), and changes users' tasks in undesirable manners. As a result, the user refrains from using the system, or resist the system's integration in the workplace by undermining its functions and using it in unintended manners.<sup>16</sup>

- Users camouflage their data signals to avoid surveillance. This practice is known as obfuscation, or “the production of noise modeled on an existing signal in order to make a collection of data more ambiguous, confusing, harder to exploit, more difficult to act on, and therefore less valuable.”<sup>17</sup>

By studying resistance as a site of dispute and value differentiation, this report locates breakdowns between consent and emerging technology, and identifies policy interventions to support human artists and creative professionals.

<sup>16</sup> M. Lynne Markus, “Power, politics, and MIS implementation.” *Communications of the ACM* 26, Vol. 26, No 6, Pages 430–444, June 1983, <https://doi.org/10.1145/358141.358148>.

<sup>17</sup> Finn Brunton and Helen Nissenbaum. (2016). *Obfuscation: A User's Guide for Privacy and Protest*. The MIT Press. Page 46.



## Methodology

Interview participants were recruited through distribution of recruitment letters in creative professional networks and snowball sampling. Participants completed a screening survey asking them to provide their occupation, workflow and materials, compensation, years of experience, and attitudes towards text-to-image generators. As shown in Table 3, the attitude criteria is a six-item checkbox list composed of sentiments and beliefs including: familiarity with text-to-image generators (*Attitude 1, Attitude 4*), use of text-to-image generators in current workflows (*Attitude 2, Attitude 5*), and use of text-to-image generators in future workflows (*Attitude 3, Attitude 6*).

Some participants expressed familiarity with how text-to-image generators work (*Attitude 1*), but were nonetheless resistant to their implementation (*Attitude 4, Attitude 5, Attitude 6*). We identified *Attitudes 4, 5, and 6* as key indicators of resistance, whereby participants refused to adopt text-to-image generators in the workflows and/or combatted acts of AI training and mimicry, and we pursued these populations for interviews. *Attitude 1* was another strong indicator of resistance when paired with *Attitudes 4, 5, and 6*, and a weaker indicator of resistance when paired with *Attitudes 2 and 3*. We interviewed a diverse range of digital and traditional medium artists, as well as early-, mid-, and late-career professionals (Table 4). Additionally, some participants surfaced artifacts (e.g., artwork, creative software programs) for review.

**Table 3: Initial Attitudes Towards Text-to-Image Generators**

Attitude 1	I am familiar with how text-to-image generators work, including their operations and outputs.
Attitude 2	I have used text-to-image generators in my work.
Attitude 3	I am inclined to use text-to-image generators in my future work.
Attitude 4	I am unfamiliar with how text-to-image generators work, including their operations and outputs.
Attitude 5	I have not used text-to-image generators in my work.
Attitude 6	I am not inclined to use text-to-image generators in my future work.

RESISTANCE TO TEXT-TO-IMAGE GENERATORS  
IN CREATOR COMMUNITIES

**Table 4: Participants' Self-Reported Demographics and Attitudes**

Participants	Occupation	Years Active	Attitudes	Preferred Tools
Participant 1 (P1)	Digital Illustrator, Graphic Designer	5+ years	Attitude 6	Clips Studio Paint, Krita
Participant 2 (P2)	Graphic Designer, UX Designer	5+ years	Attitude 1, 3, 5	Figma, Procreate, Adobe Suite
Participant 3 (P3)	Illustrator, Digital Artist	10+ years	Attitude 1, 5	Procreate, Adobe Illustrator, Adobe Photoshop
Participant 4 (P4)	Software Engineer, Media Artist	30+ years	Attitude 1, 2, 3	Stable Diffusion
Participant 5 (P5)	Painter	10+ years	Attitude 4, 5, 6	Oil paint, canvas
Participant 6 (P6)	Special Project Sculpture and Plastic Arts	35+ years	Attitude 5, 6	Plastic, clays, stone, bronze, watercolors, oil paint, acrylic paint, mixed media, Photoshop
Participant 7 (P7)	Photographer	10+ years	Attitude 1, 5, 6	Adobe Lightroom

# Interview Thematic Analysis

## **THEME 1: ARTISTIC MODALITIES AND STYLE PREFERENCES SHAPE ATTITUDES TOWARD TEXT-TO-IMAGE GENERATORS.**

### **Creative Workflows**

Participants' attitudes toward text-to-image generators were shaped by several factors, including individuals' preferred tools, mediums, and workflows; their personal style and style influences; and modalities for creative production. P5 and P6 preferred physical mediums and mixed media, using fewer digital programs to complete their work. P1, P2, P3, and P7 used a combination of computer devices, tablets, and creative software programs. Procreate, Krita, and Adobe Photoshop utilize *raster* graphic editing, or pixel-based image resolution, which allows for natural blending opportunities while painting digitally. Adobe Illustrator utilizes *vector* graphic editing, or geometric forms, to scale 2D/3D objects without losing resolution quality. Some programs, such as Clips Studio Paint, provide both raster and vector graphic editing. Most participants' workflows did *not* include text-to-image generators. Participants continually preferred the agency provided by tactile movement, such as illustrating and shaping materials by hand: this was true for participants who used traditional mediums *and* participants who used creative software programs.

### **Creative Friction**

Some tools or program features did not fit participants' needs and were excluded from workflows. P7 avoided altering aperture values in Adobe Lightroom, and noted, *"I don't find it very helpful. And I also don't like how untrue to the photo it would be."* This participant chose to maintain the integrity of the photo through smaller manipulations. Other programs, such as Adobe Photoshop, felt *"clunky"* and ran poorly on their laptop device (P1). Adobe Photoshop and Adobe Illustrator offer new generative image features, but participants refrained from their use, preferring to translate their own ideas and references onto paper (P3). Some participants disfavored early versions of DALL-E, favoring more open source software with fewer content restrictions (P4). P2 avoided using Figma's generative AI plug-ins because *"relying on AI, or the things that generate the design or generate wireframes or stuff like that, makes me feel a little more disconnected from the final product. And I personally like to have my hands in it, so*

## RESISTANCE TO TEXT-TO-IMAGE GENERATORS IN CREATOR COMMUNITIES

*maybe it's a control issue thing.*" P5 resisted more technology-focused modes of production (e.g., digital painting via iPad). Both P5 and P6 were unlikely to engage AI tools at any point in their workflow.

A recurrent theme among participants was the ability to identify an AI-generated work. Participants identified several rendering errors that signaled that a work was AI-generated, including inconsistent linework, distorted compositions, and discrepancies between digital and traditional mediums. "Rendering" indicates the process of developing an artwork into a "realistic" image with three-dimensional shape and properties. In painting 2D imagery, this may look like adding lines, forms, colors, light, and shadow to increase the overall depth and intricacy. Rendering can tell a story about how a piece was made or executed. P1 observed that, in AI-generated images, *"You'll see strands [of hair] blend into the clothes. Or sometimes you'll just spot certain things where it's rendered really beautifully, but the base thing they're trying to render is wrong."* P3 noted that inconsistencies in faces or hands can signal that a piece is AI-generated; P5 observed that hands possess a range of movements and visual orientations that may elude a computer's comprehension (or the scope of its training). P6 suspected that an oil painting that possessed an "odd" digital contrast was produced by AI. The suspicion may indicate a larger distinction between traditional and digital mediums, whereby digital mediums can achieve similar effects as traditional mediums without using the same tools. However, both traditional and digital artists observed that AI typically possessed an "uncanny" (P2) or "soulless" (P3) quality.

### Unfamiliarity with AI

Participants raised uncertainty about the mastery of generative image algorithms. P2 expressed, *"I haven't quite mastered AI to use it and not get frustrated."* P2 was most frustrated with text-to-image generators' lack of range in producing diverse skin tones, as well as its hazy linework, random blurred effects, and uncanny subject depiction. P2 also questioned AI training processes: *"It could be my assumption that a lot of the data that is already on the internet, they just scraped to use to refine their models or whatever. But I do think there's a level of . . . someone has to categorize these things? Maybe because I don't understand how it works on their side."* P2 surfaces a common uncertainty about labor inputs for AI models.

### The Labor Behind Categorizing Images

The task of “categorizing” images is presently achieved by Contrastive Language-Image Pre-training (CLIP) embeddings. CLIP encodes visual and textual data, and tags images with content-specific captions (e.g., a picture of vehicle may be tagged as a “a photo of a car driving on the highway.”)

Image classification has historically been achieved by thousands of humans manually tagging images through projects like CAPTCHA and ImageNet. But the introduction of self-supervised machine learning algorithms, such as CLIP, have dramatically reduced the number of human laborers required to tag images.<sup>18</sup>

However, generative image algorithms are dependent on human labor in other capacities, including the labor of contributing and moderating data in training sets. As such, it is critical to acknowledge that contributing to these repositories **is an act of labor**, the terms of which may be best agreed upon **prior to** training on an individuals’ work. This requires the **informed consent** of individuals, artists, and copyright holders.

## THEME 2: RESISTANCE ACCESSES DATA NARRATIVES AND PROTECTIONS TO PROTECT WORK FROM AI TRAINING AND MIMICRY, AND TO ALIGN INTERACTIONS WITH THEIR PREFERRED USE CONTEXTS.

### Use Contexts

Several factors determine where, when, and how participants permit their work to be used. Legal contracts are one way to establish permissions. P5 describes the legal frameworks for this process: “As far as with galleries when they have my work . . . everything protects the gallery. They’ll have their commission fees and have their insurance policies and stuff like that.” Contracts are negotiated prior to the completion of a work, or prior to its presentation in a particular venue. This allows the customer or company purchasing an artist’s service to utilize the work for legally agreed uses, and allows the artist to verify these uses through written documentation. P2 and P6 acknowledged that they retain little control over contracted work after the initial purchase, which means that a customer might reprint an image in different dimensions, or widely distribute it without penalty. The granularity of permitted use will vary by contract and signing parties.

<sup>18</sup> Alec Radford et al., “CLIP: Connecting text and images,” OpenAI, January 5, 2021, <https://openai.com/index/clip/>.

## RESISTANCE TO TEXT-TO-IMAGE GENERATORS IN CREATOR COMMUNITIES

Participants also established policies for how their work could be shared on social media channels. P1 explained, “*I do communicate that when I do commissions, and it’s also in my social media bios . . . ‘You cannot repost’ and then ‘Please ask for permission first.’*” Participants generally approved educational, classroom uses of their work (P3), and allowed their works to be used as personal profile pictures with credit to the artist (P1). These uses were not categorized as malicious as long as users properly credited the artist. But online distribution can blur the lines between permitted and unpermitted uses. P1 described finding their art reposted on other accounts without permission: “*it’s not an art account . . . they just take images they see on the internet [and] repost them.*” P7 clarified that clients should not add and post filters on top of their original photographic image, nor resell the image. Clients could, however, use the image in modeling portfolios with credit to the photographer. P7 has found their work reposted without permission but said, “*I haven’t had a terrible process trying to get my stuff removed. It’s pretty straightforward. I feel like people are pretty nice about it. There’s little to no malice most of the time, because I don’t think my work invites that.*” Ultimately, each artist sets guidelines for how their work should be used and negotiates these guidelines with audiences, customers, and host sites.

### Artistic Appropriation

Participants identified instances in which generative image algorithms appropriated their work, or their peers’ works, for commercial gain. In this, it is critical to distinguish between *similarity* and *mimicry*. *Similarity* indicates parallels between two artists’ works, where one work takes inspiration from or pays homage to a predecessor’s work, but could also stand on its own merits. P6 described finding a sculpture online that was similar to one they had previously made, with the “*same composition, same body posture, same size and everything.*” In these cases, it can be hard to determine if one work has deliberately taken inspiration from another, unless it is credited as such. *Mimicry* indicates the referencing of specific art styles, aesthetics, and creators within image generator prompts, thereby recreating these styles for AI outputs. Participants noted that more widely known artists are likely to be referenced in AI image generators, including Moebius and the late Kim Jung Gi. Prompts may also mimic aesthetics tied to notable creators, directors, and media brands. *Mimicry* is notable for deliberate aesthetic references that audiences may quickly recognize and identify with an established artist or brand.

The boundaries between *similarity* and *mimicry* can blur. P7 recalled several YouTube videos explaining how to edit photos in the styles of a respected photographer. In some cases, this may simply be *similarity*, but it could also be considered *mimicry*, where videos invoke notable names to attract viewers seeking to emulate these visuals. P3, P5, and P6 distinguished between

taking inspiration from other artists and blending it with their own personal styles and workflows, and the act of generative image algorithms “*mashing them [images] together*” (P3). Participants generally viewed appropriation of work from small creative professionals by larger companies as problematic and unfair. The concerns stem from the uncertainty about whose artwork is influencing an AI-generated image, which makes it difficult to assign credit to the created image. Metadata, including diffusion prompts and LoRAs that specify and fine-tune style, can be a critical source of information for how AI works are related to prior existing works.

## Resistance

Despite varying familiarity with text-to-image generators, participants expressed a shared understanding of the importance of protecting their creative works from exploitation by third-party data services and corporations. In this context, resistance emerged as thoughtful sabotage of text-to-image generators by preventing training and machine readability. P1, P3, and P7 articulated protections such as posting low-resolution images, watermarks, and local file saving features. Posting low-resolution images ensures that others cannot repost or assume authorship of clear, full-resolution versions because they are retained by the original artists. Additionally, if artwork is reposted without permission, watermarks and signatures can help identify the original artist. P3 verified that their work had not been used for AI training through “Have I Been Trained,” a program through which users may search scraped images within the LAION 5-B database and opt-out images from future use. P3 also expressed concern that the Adobe Creative Cloud was utilizing images saved via the Creative Cloud to train AI algorithms. P3 combatted this by saving their files locally and sharing files with collaborators outside the Creative Cloud. Participants who held Attitude 1 (P2, P3, P4, and P7) generally had a capable understanding of text-to-image generators and their operations; when paired with more aversive Attitudes 5 and 6, they could identify how to sabotage those operations. P4 admitted that robots.txt could stall web crawling, but did not prevent users from copying images from websites and submitting as reference material for text-to-image generators.

P1 identified two programs that helped them disrupt AI training and mimicry. Glaze provides a defense against style mimicry by layering a cloak over original artwork, disrupting readability by machine readers.<sup>19</sup> Through the Glaze interface, users may determine the magnitude and render quality of the cloak based on the level of desired protection. This cloak is near imperceptible to the human eye, and does not disrupt the visual appeal of the artwork. Conversely,

19 Glaze: Protecting Artists from Style Mimicry by Text-to-Image Models. Shawn Shan, Jenna Cryan, Emily Wenger, Haitao Zheng, Rana Hanocka, Ben Y. Zhao. Proceedings of 32nd USENIX Security Symposium, August 2023. <https://people.cs.uchicago.edu/~ravenben/publications/abstracts/glaze-usenix23.html>

## RESISTANCE TO TEXT-TO-IMAGE GENERATORS IN CREATOR COMMUNITIES

Nightshade is an offensive program that poisons deep learning models' ability to produce relevant images by mispairing text and image data.<sup>20</sup> For example, the program might sabotage an AI model's concept of a "dog" by labeling an image of a "cat" with text that describes canine species; repeating this process at scale can make the text-to-image generators fail to produce relevant images of a "dog" when prompted by a user. These programs were pioneered by TheGlazeProject at the University of Chicago. The project signifies a significant contribution to human artists seeking to obfuscate their work from generative artificial intelligence models, and undermine the efficacy of these models trained on artwork without permission.

Participants who were less familiar with text-to-image generators expressed uncertainty about how to apply technical interventions. P6 explained that one way to sabotage text-to-image generators was to *"just not post anything online."* However, as many creative opportunities exist online, disconnecting portfolios from internet audiences may also mean a loss of income and positive exposure. This is a critical tradeoff in managing online visibility. Creative professionals struggled to identify resistance strategies specific to text-to-image generators, and largely defaulted to general strategies against misuse of their art work. While some participants saw reporting copyright infringers to site administrators as a possible course of action (P6, P7), they were not confident that reporting produced real consequences, and believed that little could be done after their work was used for AI training.

In response to AI training, participants limited their online presence. P1 and P7 started sharing their work in smaller channels with members opposed to AI-generated art, and on personal accounts, private Instagram profiles, and 'Close Friends' lists. *"How do I control it? Passwords and private accounts,"* said P7, explaining that they only provide the password to their portfolio websites to recruiters and people who asked and were approved. This gave P7 knowledge of who could access their work and when. Another proposed channel was Cara, a platform built exclusively for human artists that included AI detection and filtering tools. These channels provided participants with a community of like-minded artists with similar values and workflows, and a safe place to share one's work. P2 archived and deleted older Instagram posts featuring their art work; P3 mentioned that some other artists removed themselves completely from social media. P1 said they had previously experienced a *"very big need to get followers, a very big need to get likes, comments,"* but after the emergence of AI-generated art, felt that exposure could be quite burdensome. P6 reduced the number of hashtags they used: *"I don't overwork it,*

20 Nightshade: Prompt-Specific Poisoning Attacks on Text-to-Image Generative Models. Shawn Shan, Wenxin Ding, Josephine Passananti, Haitao Zheng, Ben Y. Zhao. To Appear: Proceedings of 45th IEEE Symposium on Security and Privacy, San Francisco CA, May 2024. Preprint: arxiv/2310.13828, February 2024. <https://people.cs.uchicago.edu/~ravenben/publications/pdf/nightshade-oakland24.pdf>



## RESISTANCE TO TEXT-TO-IMAGE GENERATORS IN CREATOR COMMUNITIES

*for example, in the hashtags. I just hashtag one, two, three, four max or something like that, not too many hashtags . . . Not too many . . . The more that you do, the more exposure you have.”*

Aligning with P6’s sentiment that posting artwork online “takes away control,” participants said they were moving to physical media and in-person platforms as a way to continue their work without risk of AI model training. P1, P2, P5, and P6 mentioned that displaying their art at galleries, exhibitions, pop-up shows, and conventions tables earned them exposure and audience connection without requiring an online presence. Instead of selling their art as digital images, P1 and P3 looked to sell their work as physical zines and art books. In part, this may be attributed to a revival in appreciation for art made with traditional media: *“I feel like with any type of technology there is a cycle . . . [Digital was big] a couple years ago and now we’re going back towards physical media [. . .] because there’s a certain look and feel to it”* (P3). The use of physical media represents a unique form of resistance whereby artists claim their work through materiality, and by forgoing algorithmic tools for more personal human connection. While physical media may constrain the reach of one’s work, this tradeoff is beneficial if it reaches audiences who will value it and compensate its makers.

### Data Provenance

Despite holding varying attitudes about generative text-to-image technology, participants shared concerns about data privacy. AI model training accumulates text-image inputs from varying databases (e.g., LAION-5B) in order to refine and further train generative image algorithms to create images with more precision. Many of the preferred tools, such as Adobe Creative Suite and Clip Studio Paint, have begun implementing generative AI features; P1, P2, P3, P6, and P7 said they were hesitant about these technologies and rejected their incorporation in their creative workflows. Similarly, participants pointed out the prevalence of AI-generated images on social media and the role these platforms play in AI model training. These platforms often have terms of service indicating privacy policies and contexts for data use. However, users might not read these policies or know how a particular policy affects their platform usage. P1 explained that the Adobe Creative Suite had performed a *default opt-in* to scrape users’ images for AI/ML training. When made aware, users immediately notified each other to turn off this feature in their privacy settings. These concerns were amplified with the release of Adobe Firefly, which used Adobe Stock contributions to train its algorithm.

Participants shared a ubiquitous sentiment that their data was being scraped, aggregated, and used for AI/ML operations. This lack of control diminished participants’ trust in platforms and services they commonly used. P2 wanted more clear rules around participation: *“I feel like*

## RESISTANCE TO TEXT-TO-IMAGE GENERATORS IN CREATOR COMMUNITIES

there could be more of a like an *opt-in* or a hard stance like, ‘Oh, we only reference like specific artists.’ But then it gets into the territory of the value of art, which is a whole other conversation, where it’s like “Why do these specific artists get chosen to train these models?” The desire for a more consensual approach was echoed by P3 and P7, who felt that *choice opt-in* would allow users who want to participate in AI algorithms to do so, and exclude those who do not. Accordingly, participants condemned *default opt-in* controls that force users to participate in training AI models without opting out. P4 and P7 stated that transparency between users and companies, including information about model inputs, was an important factor in managing these issues. Data provenance, or public documentation of data as it moves through research and development, may facilitate transparency.<sup>21</sup> A lack of transparency and tools for filtering participation were consistent concerns among all participants, regardless of their preferred tools and workflows.

### **THEME 3: RESISTANCE INVOKES DATA NARRATIVES AND PROTECTIONS WHERE COPYRIGHT LAWS FAIL TO BE ENFORCED.**

Creative professionals expressed anxiety about AI causing job displacement and major industry disruptions, and altering job responsibilities. When discussing possibilities for compensation in return for their work to be included in AI training models, participants felt that they deserved monetary compensation, although the type of compensation model that could fulfill their needs and licensing terms, especially retroactively, was unclear. Regarding their work being used for AI training, many participants had a limited understanding of their intellectual property rights or encountered obstacles in enforcing them. Creative professionals felt that the training of AI models was unfair, but that copyright protections were also not positioned to prevent this violation. Participants were also concerned by how the use of their work in AI-generated products newly altered the usage license of their work. Subsequently, additional contracts and union memberships were seen as more active ways to mitigate unfair use and collectively negotiate the changing landscape of AI-generated works.

### **Employment**

Participants saw a need for compensation in a world where AI models would continue to use their assets and work for training. P1 considered the training of models to be analogous to a chargeable service: “*In the end you are using somebody else’s assets for your own thing*

<sup>21</sup> “Data Provenance,” (n.d.), Network of the National Library of Medicine, <https://www.nlm.gov/guides/data-glossary/data-provenance#:~:text=Definition,to%2owhere%2oit%2ois%2opresently.>

## RESISTANCE TO TEXT-TO-IMAGE GENERATORS IN CREATOR COMMUNITIES

*and you do have to pay for that.*” The idea of reparative solutions from work already used for training produced more conflicting responses. Some participants proposed a royalties system, by which artists could be compensated based on the number of times their name or style was prompted (P2) or based on a calculated percentage similar to Spotify’s “pennies per stream” royalties plan for musical artists (P5). On the other hand, some felt that compensation for stolen work was less meaningful or restorative (P7). Others look to artistic unions and associations, such as the Animation Guild, to negotiate protections against AI displacement in contract negotiations (P3). Proposals for compensation in training models were further complicated by the specific terms that creative professionals have control over when negotiating directly with clients. Participants described how work they produce for commercial and non-commercial purposes are compensated differently (P1, P3, P5), but this would be difficult to apply given loose regulations around the permitted use for AI-generated work. P6 stated, *“The only compensation is to relegate that work to me. And that commission, whatever that was, it’s the only way.”*

While many participants suggested models for compensation, there was a recurrent concern that AI would replace certain jobs and disrupt creative industries (P1, P2, P3, P5, P6). P2 observed that their peers were pursuing project management roles over design roles: *“They may be interested in design, but they go into management just because it seems a little more stable . . . If you take it from a management perspective, it makes it a little easier because you’re not doing the labor that would be X’d out.”* However, participants also noted that advancing in the corporate hierarchy might not be viable (or desirable) for everyone, due to limited positions and lack of creative flexibility. P2 projected that *“the design industry collapse could come from the automation of this work.”* P1 lamented that AI might replace entry-level roles and the development of preliminary concepts that are typically drawn by humans. Participants acknowledged that AI could make some processes more efficient, but found that the implementation of AI tools by employers lean more towards displacement and replacement by *“[getting] rid of the human element”* (P3). Participants identified that the actors central to this conflict are not only artists and AI developers, but also upper management, companies, and studios that create demand for AI-generated work (P2, P3, P6). As P6 pointed out, *“robots would be drawing images, but it’s the company that requests the images, right?”*

### Intellectual Property Rights

Participants expressed discontent with how AI models are trained in relation to copyright and intellectual property rights. P4 described the training of AI models as *“mass copyright violation,”* including the inclusion of recent works that are not in the public domain. P1 requested

## RESISTANCE TO TEXT-TO-IMAGE GENERATORS IN CREATOR COMMUNITIES

a redefinition of fair use: *“if corporations are just taking tens of thousands of people’s work under fair use, I don’t think that’s ‘fair.’”* Despite the lack of clarity in copyright law, participants consistently advocated that AI training should be limited to public domain works. Differences in copyright term lengths by country can complicate public domain use, especially if developers are operating services in multiple countries or servicing users in multiple countries. For example, U.S. creators retain copyright for the duration of their life plus 70 years, but several nations have shorter copyright term lengths. Thus, a copyright that has expired in one country may still be active in another.

Regarding their own work, participants cited copyright and trademarks as methods for legal protections, but pursued these steps to varying degrees. P4 was the only participant who had registered their work for copyright, but shared that long wait times, registration fees, and license research can be burdensome: *“Waiting three to six months to see if your registration is approved? That sucks. I have been rejected before by the U.S. Copyright Office.”* Some participants had interacted with copyright and trademark laws for contracted work (P2, P7), but currently do not register personal work (P2, P3, P5, P7). P7 expressed, *“I think it’s general law that if you publish to the internet [...] that you’re generally covered. . . . Not sure that’s really the case anymore.”* P5 did not believe that their work was protected by copyright, joking that they had a written statement affirming their rights to their work, but *“you probably couldn’t take it to court.”* P5 also shared that peers have had their work stolen by major retailers *“and there wasn’t much that they could do because it wasn’t protected in any way.”* P6 noted that seeking recourse for copyright infringement was difficult *“because laws change in every country.”* These examples show the gap between intent and efficacy of copyright law. Evidently, participants felt that copyright law was not a perfect solution for combatting AI use or misappropriation of their creative work by corporate enterprises.

Given the precarity of copyright protections, participants have sought other mechanisms for protecting their work from generative AI image algorithms. P2, P6 and P7 discussed their usage of contracts and licensing agreements. P1, P3 and P7 expressed a need for unions and creator guilds to collectively negotiate pay, standardize protections and privacy, and address interventions to AI-generated work in employment contracts. Participants also emphasized a desire to be credited for their work, regardless of AI use. P5 said, *“I would hope that there would be some sort of way to have my name attached to these things . . . like if the AI stole specifically from my pieces for something like that, I would love to have a little memo that says “inspired by [name].”* Attribution reinforces authorship and ownership, acknowledges the time and energy expended to create a work (P3), and provides exposure for future work opportunities. How-

RESISTANCE TO TEXT-TO-IMAGE GENERATORS  
IN CREATOR COMMUNITIES

ever, participants found it impossible to enforce authorship credit. This desire to be acknowledged as authors might be heightened due to the culture of AI prompts claiming authorship over generated works: *“It’s, analogously, commissioning a person to do it for you, and then calling it your own. Feel like it’s like ghost writing. I also don’t like that either”* (P7). Overall, AI has disrupted standards normalized by copyright, forcing participants to take alternate measures.

## Discussion

A recurrent theme through the interviews was the ethic of “consent” or “informed consent,” defined as participation that is “knowing and voluntary,” or “freely given, specific and informed.”<sup>22</sup> Participants engaged “consent” in several facets, including the inclusion on web pages of robot.txt files to block machine learning and “*stop gratuitous use in building without my consent*” (P4). They also called for contractual agreements of permitted use contexts by clients (P6), and licenses and sub-license agreements (P7). Through these various mechanisms, creative professionals can indicate which projects they want to participate in, and which they do not. Contracts also provide the client with guidelines for how they may use the product or service purchased. For example, P7 allows clients to use purchased photos for modeling portfolios if they include a photographer credit. Creative professionals negotiate relationships between their art and their audience, whereby both the artist and audience are aware and comfortable with the terms of exchange. Additionally, creative professionals negotiate relationships between their art and creative software programs, specifically the use of their art as “data.” Some participants expressed not feeling that their data setting preferences were secure after Adobe saves and sends them to a server (P3); others expressed that DeviantArt and Adobe set users’ preferences as “*default opt-in*,” allowing training on users’ works without manual permission (P1). Distrust over data mistreatment between artists and creative software programs and distributive platforms remains a large point of division in creator communities.

The thematic analysis revealed that creators *may not know* whether their works have been included in generative artificial intelligence training. Many expressed that, because their work has been posted online, it has most likely been scraped, crawled, or tagged as training data without their knowledge; even so, they found this hard to confirm. Does the online publication of user content grant permission for training in AI models? Must digital service providers indicate or request permission to train models with user-generated content? Between users and service providers, these questions might be answered in privacy policies and terms of use contracts, which outline permissions provided over user-generated content. For example, a service provider may disclose that it shares data with third-party advertisers and partnerships to enhance its services. By indicating “yes,” the user accepts the conditions of this contract in exchange for using the service or platform. However, legal scholars contend that notice-and-choice models do not capture the scope of privacy risks. Consequently, users might not understand the risks of data sharing, especially where surveillance, aggregation, and dissemination

22 Neil M. Richards and Woodrow, Hartzog, “The Pathologies of Digital Consent,” 96 Washington University Law Review 1461 (2019), April 11, 2019. <https://ssrn.com/abstract=3370433>

## RESISTANCE TO TEXT-TO-IMAGE GENERATORS IN CREATOR COMMUNITIES

are coded into user-friendly services. Thus, “informed consent” often fails to ensure that users have the knowledge necessary to make a suitable decision. In cases where developers crawl multiple websites, a user may not know that training is happening, nor do they voluntarily elect to participate.

Participants proposed three alternative templates for consensual participation:

- (1) *Default opt-out*, whereby users are automatically disenrolled from AI/ML training. *Default opt-out* indicates that users are excluded from AI/ML training unless they take further action to opt in.<sup>23</sup>
- (2) *Choice opt-in*, whereby users are notified and elect to participate in AI/ML training. If a user manually *opts in* to a service, then it may be assumed that they are aware of the conditions and are electing to participate by choice.<sup>24</sup>
- (3) *Training exclusively on public domain works*, whereby model developers select text, images, and other audio-visual data exclusively from the public domain (including works made prior to 1929, and works intentionally labeled or licensed for free distribution) if based in the U.S.

These templates surface interesting tensions between research ethics and research development. P4 explained that a generative artificial intelligence system limited to public domain materials would be a “*cool visceral experience for people to understand the legal regime,*” and could help reveal how historical limits change the products of image generation. Restricting training materials could decrease the quality and range of model outputs: “*I definitely feel like people would be a little bit disappointed with the results. The reason why the text-to-image generators get a lot of crap is that they’re impressive. . . . So I feel an opt-in model would be nice, but it would still limit the abilities of AI*” (P7). P7 shares a criticism that AI-generated outputs are “impressive” due to their aesthetic similitude to past and present artistry, and that a limited training set would make their outputs less impressive. Even so, participants preferred to have a *choice opt-in* approach (P2, P3), to be excluded from training altogether via *default opt-out* (P1), or to have a combination of both (P7). In either template, system developers are obligated to request permission to use artists’ work prior to training AI/ML models and on an individual basis.

23 This is **not** synonymous with *opt-out* whereby a service enrolls users into AI/ML training first, and users must take additional steps to disenroll.

24 This is **not** synonymous with default *opt-in* whereby a service enrolls users into AI/ML training first and users must intentionally disenroll, or model developers assume that a work is eligible for training by nature of accessing it via the Internet (e.g. Common Crawl web crawling).

## RESISTANCE TO TEXT-TO-IMAGE GENERATORS IN CREATOR COMMUNITIES

Among the least-cited mechanisms for resistance was the reporting tool, which would allow creators to file complaints about the use of their works in AI models after the fact. P7 seemed most likely to report plagiarism, mimicry, or misuse of copyrighted works, or to motivate peers on social media to report misuse on their behalf. P6 was likely to gather photo documentation of a violation of their copyright, but was not keen to pursue future action. Reporting individual copyright infringements can be time-consuming, and does not stop the potential for future infringement. Takedown notices have the potential to be abused by a high volume of automated requests, as well as large studios and publishers exercising copyright protections where *fair use* may apply. Instead of reporting, participants tended toward obfuscation, such as adding Glaze layers to disrupt style mimicry and reducing images to low-resolution quality, and “poisoning” AI training sets through programs like Nightshade. Participants also modeled resistance by refusing to use systems that removed their agency. This was evident in the use of local saving features instead of cloud services, which would allow them to take advantage of Adobe’s software without complying with the company’s AI/ML training. These methods are *pre-emptive to harm*, whereas reporting is *retroactive to harm*. By the time a copyright holder reports infringement, they may have experienced a significant loss of money or agency.

Nonetheless, reporting can be an important tool of online governance and mediation, as it creates a documented trail of copyright infringement. Platforms have even designed reporting tools to accommodate individual creators’ needs. For example, YouTube’s Copyright Match Tool scans for re-uploaded content, and helps channels take appropriate actions based on the sampled use. This tool is suited to its more than two million Partnered Program channels, highly visible creators who “experienced a higher amount of reposting of their copyrighted content,” but were not necessarily movie studios or corporate publishers.<sup>25</sup> Reporting may offer an avenue for recompense, as well as documentation of misuse involving text-to-image generators.

When pursuing a takedown request, copyright holders might identify the following:

- (1) *Inclusion of copyrighted works in training data sets for generative artificial intelligence systems.* This indicates images, texts, and/or audio used for training models for which copyright holders have not provided express permission or whose existing permissions do not allow for their inclusion.

25 “YouTube Copyright Transparency Report H1 2022,” YouTube, [https://storage.googleapis.com/transparencyreport/report-downloads/pdf-report-22\\_2022-1-1\\_2022-6-30\\_en\\_v1.pdf](https://storage.googleapis.com/transparencyreport/report-downloads/pdf-report-22_2022-1-1_2022-6-30_en_v1.pdf).



RESISTANCE TO TEXT-TO-IMAGE GENERATORS  
IN CREATOR COMMUNITIES

(2) *Appropriation of copyrighted works for referential study in generative artificial intelligence systems.* This includes the use of copyright holders' names as generator prompts, and user submission of copyrighted works as reference material for generators to modify existing and generate new visual, text, and sound elements. These activities are done without permission by the copyright holder.

The scope of generative artificial intelligence systems limits copyright holders' capacity to identify every misuse of copyrighted work. For example, a copyright holder may locate several users posting generated images or text that closely replicate their material. This is possible by prompting a text or image generator with the copyright holder's names or copies of their works (due to *appropriation of copyrighted works towards referential study in generative artificial intelligence systems*). A copyright holder may also locate their work in a training database for which opting-out is difficult and ineffective, and does not reverse prior model training (due to *inclusion of copyrighted works in training sets for generative artificial intelligence systems*).

The following considerations are posed to provide a record of emergent technology in take-down notice procedures and reinforce copyright holders' consent power:

(1) *Notice-filers may complete a survey indicating whether a generative artificial intelligence system has contributed to an infringing use(s).* A DMCA notice of copyright infringement requires the description and identification of infringed copyrighted works and infringing material. Online service providers may provide an accompanying, optional survey inquiring whether:

(1a) The copyright holder has sufficient belief that a generative artificial intelligence system has contributed to the infringing material;

(1b) The infringing material involves the *inclusion of copyrighted works in training sets for generative artificial intelligence systems*;

(1c) The infringing material involves the *appropriation of copyrighted works towards referential study in generative artificial intelligence systems*; and

(1d) The copyright holder wishes to make this information available to the alleged infringer.

(2) *A counter-filer may respond to claimed use(s) of generative artificial intelligence systems.* If the copyright holder makes the survey results available, and the alleged

## RESISTANCE TO TEXT-TO-IMAGE GENERATORS IN CREATOR COMMUNITIES

infringer does not remove the material, the latter may file a counter-notice, including whether or not the contested material involves the use of a generative artificial intelligence system. Statements should be made in good faith.

(3) *Online service providers may provide definitions for “generative artificial intelligence systems.”* These definitions should be consistent with consensus by governing bodies and stakeholders, as to provide accurate information to notice-filers and counter-filers.

Designated agents may confirm if copyright holders would like to make survey results available to alleged infringers for response, or if they wish to retain answers for online service providers’ records. These mechanisms may hold parties accountable for disclaiming use of generative artificial intelligence systems in creative production and reuse, and can hold online service providers accountable for documenting infringing uses involving generative AI systems. Enhancing reporting and takedown mechanisms may provide copyright holders further means of identification and reclamation.

## Conclusion

This report surfaced key elements of creators' resistance to text-to-image generators. Resistance has emerged as an alternative to copyright law in the management of generative artificial intelligence systems, and a way to pre-empt misuse by other users, companies, and AI developers. To mitigate the threat of generative AI appropriating their creative work, the participants interviewed for this study have reduced their online visibility, obfuscated their work from algorithmic surveillance, and created smaller communities of engagement for their work. They have also adjusted their privacy and data settings to align with their preferences. These strategies revealed that many creators are distrustful of creative platforms due to the potential misuse of data in developing generative artificial intelligence systems.

While these strategies may fill short-term gaps, a range of stakeholders — including policy-makers and industry practitioners — can pursue other actions to support artists' needs. The recommendations below aim to align copyright and data management with artists' desire for consent and identification of AI misuse. We hope these recommendations will serve as a starting point for intentional discussions between artists, program developers, and copyright administrators.

### RECOMMENDATIONS

**We recommend that Congressional legislators adopt H.R.7913 - “Generative AI Copyright Disclosure Act of 2024.”** Sponsored by Rep. Adam Schiff, this proposed bill would obligate AI developers to disclose copyrighted materials used in training sets to the Register of Copyrights. This legislation has the potential to make AI training data transparent to U.S. federal institutions and the general public. Legislators may also consider elements of the Federal Anti-Impersonation Right (FAIR) Act as protection for human artists against impersonation and mimicry for commercial gain.<sup>26</sup> These proposals have the potential to establish liability for AI developers and misusers of AI programs.

**We recommend that the U.S. Copyright Office and online service providers adopt a survey addendum to takedown notice procedures indicating misuse involving text-to-image generators.** In this survey, copyright holders could indicate the type of misuse as *inclusion of*

<sup>26</sup> Dana Rao. “The FAIR Act: A new right to protect artists in the age of AI,” Adobe Blog, September 12, 2023, <https://blog.adobe.com/en/publish/2023/09/12/fair-act-to-protect-artists-in-age-of-ai>.

## RESISTANCE TO TEXT-TO-IMAGE GENERATORS IN CREATOR COMMUNITIES

*copyrighted works in training data sets for generative artificial intelligence systems or appropriation of copyrighted works towards referential study in generative artificial intelligence systems.* This intervention may help stakeholders identify problematic sites and repositories, and document misuse involving these systems (see Discussion).

**We recommend that the Library of Congress support human artists through collaborative learning and education.** The migration of human artists from juggernaut entertainment studios and media platforms to smaller professional communities such as Cara is an important movement. The nation's library may facilitate these emerging communities by establishing working groups, research labs, and educational seminars led by human artists to support and preserve their creative practice. Insights from these groups may inform policies set by the U.S. Copyright Office and Copyright Royalty Board.

**We recommend that developers of creative software programs, social media, and image distribution platforms automatically exclude works created by users from generative artificial intelligence training through a default opt-out.** Users' content should not be included in generative artificial intelligence training without their knowledge or permission. If developers choose to offer generative artificial intelligence services, they should provide *choice opt-in*, whereby users can manually elect to participate (and those who do not wish to do so do not have to take further action, as they are automatically disenrolled). If permission is not granted, then system developers must honor this choice and exclude users from training programs. This structure may result in smaller training sets and more consensual participation.

**We recommend that developers of creative software programs, social media, and image distribution platforms provide interface options that separate generative artificial intelligence features from primary services.** For example, generative artificial intelligence training should be independent of users' ability to save image data to cloud services (a primary service), and not a precondition of using cloud services. Developers should separate these services through written policies, and independent user flows and widgets. Furthermore, developers should create interface options without generative artificial intelligence features, allowing users to select interfaces that align with their preferred use contexts and creative workflows.

## Acknowledgments

This work could be done without the intentional collaboration of the artists and creative professionals who volunteered their expertise to this project. This work was supported by the UC Berkeley Center for Long-Term Security, the AI Policy Hub Fellowship, and the Berkeley Center for New Media Undergraduate Research Support Grant. The authors are deeply grateful for the support of mentors, peers and reviewers throughout the project timeline, including Jessica Newman, Brandie Nonnecke, Deirdre Mulligan, Coye Cheshire, Abigail De Kosnik, and Zoe Kahn.

## About the Authors

**Janiya Peters** is an AI Policy Hub Fellow at the UC Berkeley Center for Long-Term Cybersecurity. She is pursuing her Ph.D. at the UC Berkeley School of Information with a Designated Emphasis in New Media. Her work analyzes asymmetric power structures in user design, and user resistance to algorithmic systems in creative processes. She is a Race and Digital Justice Fellow through UCLA Center for Race and Digital Justice; and a Junior Fellow at the Library of Congress, collaborating with the Connecting Communities Digital Initiative (CCDI) team.

**Bonwoo Kuh** is an undergraduate research assistant on this project through the Berkeley Center for New Media. Currently, they are pursuing a degree in Art Practice and a minor in Data Science. Their main interest lies in the intersection of data, arts, and technology.

**Isabel Li** is a human-computer interaction researcher at the Berkeley Institute of Design. Their work aims to empower and elevate the voices of creative communities, youth, and vulnerable populations in the development of emerging technologies. They have previously conducted research at Adobe Research, the Berkeley Center for New Media, and the UC Berkeley School of Information.



**CLTC**

Center for Long-Term  
Cybersecurity

---

UC Berkeley