Foregrounding Artist Opinions: A Survey Study on Transparency, Ownership, and Fairness in AI Generative Art

Juniper Lovato[1,2], Julia Witte Zimmerman[2], Isabelle Smith[3], Peter Dodds[1,2], Jennifer L. Karson[4]

[1]Department of Computer Science, University of Vermont
[2]Vermont Complex Systems Center, University of Vermont
[3]Department of Mathematics & Statistics, University of Vermont
[4]Department of Art and Art History, University of Vermont
jlovato@uvm.edu

Abstract

Generative AI tools are used to create art-like outputs and sometimes aid in the creative process. These tools have potential benefits for artists, but they also have the potential to harm the art workforce and infringe upon artistic and intellectual property rights. Without explicit consent from artists, Generative AI creators scrape artists' digital work to train Generative AI models and produce art-like outputs at scale. These outputs are now being used to compete with human artists in the marketplace as well as being used by some artists in their generative processes to create art. We surveyed 459 artists to investigate the tension between artists' opinions on Generative AI art's potential utility and harm. This study surveys artists' opinions on the utility and threat of Generative AI art models, fair practices in the disclosure of artistic works in AI art training models, ownership and rights of AI art derivatives, and fair compensation. Results show that a majority of artists believe creators should disclose what art is being used in AI training, that AI outputs should not belong to model creators, and express concerns about AI's impact on the art workforce and who profits from their art. We hope the results of this work will further meaningful collaboration and alignment between the art community and Generative AI researchers and developers.

Introduction

Art has a rich history in co-evolution with the development of new technologies (Etro and Galenson 2023). The advent of Artificial Intelligence (AI) tools used to create or assist in the creative process carries great potential utility for artists but also harbors the significant potential for harm to the art workforce and artistic and intellectual property. To examine this tension between utility and threat, we surveyed 459 artists about their views on AI-generated art. The survey results reveal the nuanced perspectives of artists on Generative AI. When the discussion is thorough and directly related to their art practice, it prevents their views from being reduced to a simplistic binary sentiment, which fails to represent the diversity of their field. By capturing artists' views through a survey, this method moves beyond anecdotal evidence and provides a counterbalance to third-party opinions and those of the Generative AI industry.

Our study surveys artists' opinions on fair practices in the disclosure of artistic works in AI training models, ownership and rights of AI art derivatives, and fair compensation models. Future work will involve conducting long-form interviews with them to understand artists lived experiences. Our research questions for this study can be summarized as follows: RQ1: Do artists see AI models as a threat to art workers?; RQ2: Do artists see AI art models as a positive development in the field of art?; RQ3: What types of disclosure for the use of artwork to train an AI art model do artists consider fair?; RQ4: From an artist's perspective, who should own the derivative artwork created by AI art models?

We summarize our research results as follows: (Q1) 61.87% of survey participants agreed that AI models are a threat to art workers. (Q2) 44.88% of survey participants agreed that AI art models are a positive development in the field of art (22.03% of survey participants agreed that both AI was a positive development and a threat). (Q3) Artists overwhelmingly agreed (80.17%) that model creators should be required to disclose in detail what art and images they use to train their AI models. (Q4) Artists were split between thinking that derivatives from AI art models should be owned by the original artists whose work is represented in the AI art (41.39%) and the user who used the model to create the AI art (39.22%). In our study, 26.80% of the participants agreed that the work and its derivatives should be considered the property of the AI model creators. We also looked into what types of compensation artists consider to be fair in exchange for the use of their artwork to train an AI art model. Half of the respondents (50.97%) stated that they did not need compensation, but what mattered to them was who was profiting from the use of their artwork. The most common fair compensation option selected by the survey participants was "I don't need to profit, but I don't want for-profit companies to profit from my art" (22.80%). This suggests that artists are not entirely against their art being used in AI art model training data but oppose for-profit companies benefiting financially.

Our paper draws on our team's multidisciplinary background – in art and art history, social science, complex systems, computer science, mathematics, and statistics – to address these questions and understand the possible impact of Generative AI on artists. We hope this work will provide a modest step towards understanding artists' opinions on Gen-

erative AI technology, its impact on their field, and the role they would like their work to play in the burgeoning field of Generative AI.

Background

There is ongoing disagreement regarding whether the output of Generative AI can be considered art (US Federal Trade Commission 2023; Jiang et al. 2023); we will not examine the topic in depth here. For convenience, we'll refer to AI-generated output as "art" in this paper without asserting or refuting its artistic merit.

Jiang et al. define art as a human endeavor linked to culture and experience (Jiang et al. 2023). They argue that human creativity, distinguished by unique inspiration and corporeal engagement, differs from AI-generated work. Recent court decisions in the US have similarly denied copyright protection to autonomously produced AI creations, underscoring the role of human creativity in copyright law (US District Court for the District of Columbia 2023).

Three compelling arguments against considering AI output as art emerge: art's human origin devoid of "organic pressure," the intimate connection between artist and world, and art's unique role in human experience (Jiang et al. 2023; Huyghe 1970; Pfeifer and Bongard 2006). Regardless of the artistry debate, AI-generated creations can meet commercial standards, like book covers, entering diverse artistic spheres and impacting practitioners.

Artistic use of technology and AI tools (Cetinic and She 2022) is not novel; artists historically adopt tools, including machines and automation, often innovating beyond standard applications (Huyghe 1970). Artists engaging with technology question its role in art and explore its boundaries, raising queries about intentionality, distraction, and artistic agency (Karson 2023).

It is important to note that, although how technology should be incorporated into art may be up for debate, that technology is part of the artistic process is not new or ambiguous (Mazzone and Elgammal 2019). Artists have always used tools (Huyghe 1970), and many art forms involve machines (such as woodwork, metalwork, and digital art) and automation (such as digital printers and cameras). Some artists also choose to use technology in ways other than its intended purpose or develop new art-making techniques by breaking technology. Such artists may be both suspicious and intrigued by the technology they use, questioning the relationship between humans and machines and expanding art practices and processes. When presented with new technology, an artist may ask questions like "How can I keep my work with the machine intentional? How might this technology distract me? Is the machine using me, or am I using it? Is this image-making or art-making? How does this automation serve my artistic intentions and practice?" (Karson 2023).

In the case of machine learning, a crucial question for many artists is "What data was used to train this model?" Such inquiries are part of the artistic tradition of challenging the status quo and are in stark contrast to how the creators of Generative AI models have presented their tools for use to the general public.

Generative AI Models

Generative AI models can generate outputs across various media types based on user inputs, such as text prompts or multimedia (Gozalo-Brizuela and Garrido-Merchán 2023). The field of Generative AI has sparked significant excitement due to its broad range of applications, including both anticipated ones valued in the tens of billions of dollars and substantial academic and industry research efforts (Bubeck et al. 2023; Abdin et al. 2023; Jiang et al. 2023; Stanford Institute for Human-Centered Artificial Intelligence 2023). However, alongside this enthusiasm, there is a pressing need for a more critical examination of the ethical implications associated with these tools and methods (Barnett 2023). The deployment of these tools beyond academia has resulted in adverse consequences such as personal and financial losses, widespread information manipulation, and disruptions to sociotechnical infrastructure (Ferrara 2024; Walker et al. 2023).

We define Generative AI as being a subset of tools or processes that combine massive probability distributions with sense-making algorithms in order to create output that, taken as a whole, was not included in the training data (Jiang et al. 2023). The probability distribution is based on the training data. The sense-making algorithm could be based on several strategies, such as a physical process like diffusion or an architectural approach like transformers/neural nets, although in the end, it should produce something akin to a latent/embedding space; a place to hold meaning (Gozalo-Brizuela and Garrido-Merchán 2023; Zhao et al. 2023; Grand et al. 2022). These models use advanced architectures in conjunction with large amounts of data gathered through brute-force methods across vast swathes of the publicly accessible internet (Xu et al. 2023; Zhao et al. 2023).

In this paper, we are interested in how artists perceive Generative AI, what its impact on art-related careers might be, and what moral, ethical, artistic, and financial implications are coming hand-in-hand with the broader flourishing of Generative AI (Stanford Institute for Human-Centered Artificial Intelligence 2023; Gozalo-Brizuela and Garrido-Merchán 2023; Zhao et al. 2023). Although our paper focuses more on the pressing potential for harm, it is worth noting that there is also growing excitement around Generative AI's capabilities and how it could be incorporated into artistic processes (Karson 2023). For example, artist Cosmo Wenman worked with blind researchers Brandon Biggs, Lindsay Yazzolino, and Joshua Miele to create images using the Midjourney tool. Their team noted that programs like Midjourney can find "weird and thematic intersections" in art that go unnoticed by human eyes. They also discussed the role of chaos in creating "visual intrigue" and how technology can be used to "enhance agency." Wenman called the new technology "essentially moveable type for visual rhetoric" (Schrader 2023).

Related Work

A paper by Jiang et al. (Jiang et al. 2023) analyzes Generative AI art models and their potential impact on the world through the lens of diverse academic sources (both philosophical and scientific). They describe many opportunities

for harm as a result of such products: violation of the consent and privacy of artists (Ho 2024), a lack of dataset transparency (shielding violations and biases) (Ho 2024), a failure to uphold the contextual integrity of content (Nissenbaum 2004), harm to the livelihood of the artist, biases in the training data, automating processes in a way that is detrimental to the creation of art, reduced accessibility, erasure of human works by obscurity, and increasing reluctance by artists to share their works openly.

The harms Jiang et al. (Jiang et al. 2023) foresee for artists fall into two categories, financial and emotional, which manifest as the replacement of a human artist on a project, the impersonation of an artist without consent (Gozalo-Brizuela and Garrido-Merchán 2023), or plagiarism. Impersonation, in this context, could be weaponized for nefarious purposes (Ho 2024). Moreover, these two kinds of harm, enacted against individuals at scale, create societal-level harm. Specific societal harms Jiang et al. discuss in their work include data laundering – through academic or open-source projects where public funds are funneled into research with private profits – the proliferation of stereotypes and "hegemonic views," and erasure by obscurity (i.e., that AI-generated outputs can be created much more quickly than human ones). They also point out that US copyright and intellectual property jurisprudence needs to be better equipped to deal with these new tools and that open sourcing, often called on as a solution to software-related problems, is incomplete in this case (Jiang et al. 2023). As possible solutions, Jiang et al. bring up regulations regarding opt-in consent (for training data), independent funding (much current research is supported by industry), and education about how power interacts with technology. In particular, they note there are many incentives to label work as "apolitical" and that the spin capitalism is putting on Generative AI is the "democratization of art" (and other fields) (Jiang et al. 2023).

Going beyond Jiang et al.'s list of solutions, Shan et al. (Shan et al. 2023) propose tools allowing artists to "poison" models that scrape their work without permission. The tool they created (Nightshade) can "destabilize general features" in the targeted generative model and "bleed-through" to semantically similar prompts without injecting large amounts of tailored data (Shan et al. 2023).

Economic Context: Livelihood of the Artist

AI's integration into the global economy is rapidly expanding, as evidenced by the increasing number of AI-related job postings across various economic sectors (Stanford Institute for Human-Centered Artificial Intelligence 2023; Squicciarini and Nachtigall 2021). Employers worldwide are seeking AI-skilled workers, with many companies adopting AI tools each year. The primary benefit reported by these companies is a reduction in costs, often achieved through layoffs and the substitution of human roles with AI-driven solutions (Stanford Institute for Human-Centered Artificial Intelligence 2023; Gmyrek, Berg, and Bescond 2023; Alekseeva et al. 2021). Some of these impacts are already coming to bear on the economic reality of artists (Chow 2020; Henderson 2010). Entry-level positions in film, TV, and gaming seem particularly at risk, and the conceptual period of a project seems to be a popular phase in which to replace artists with AI. The output of Generative AI can now fill traditional niches like illustrating book covers, and major employers can reduce the number of employees needed (for example, Netflix Japan's layoff of animators) (Jiang et al. 2023).

The economic impact on artists is not solely determined by the output of AI models. Harm, both financial and emotional(Caporusso et al. 2023), can arise from the use of content without explicit consent. This includes situations where individuals may not have anticipated their work being repurposed by machine learning models or the entities controlling and utilizing these technologies (Nissenbaum 2004). For instance, consider an artist with a physical gallery presence whose work is photographed by a visitor and subsequently used in ways they did not foresee, especially in the intricate landscape of social media (Ho 2024). Moreover, the issue of copyright protection becomes paramount when AI models are trained using human artworks without proper consent, attribution, or adequate compensation for the original creators. This concern is shared among artists, creative industries, and government agencies alike (US Federal Trade Commission 2023; Marcus and Southen 2024).

Problems in the Pipeline of Generative AI Tools

In this section, we outline the process of creating a Generative AI model and focus on bias (Zhao et al. 2023) to show how different stages influence the model's impact. The development process involves several key steps. First, training materials are gathered, often from real-world data, which are then reviewed to address issues like poor quality, irrelevance, or over-representation (Luccioni and Viviano 2021; Zhao et al. 2023). Next, the model's architecture is built, followed by the training phase. For complex Generative AI models with unpredictable behavior, an additional alignment step is included to ensure the model's actions align with the creators' intended outcomes (Zhao et al. 2023).

Early in the pipeline, the selection of training material significantly shapes the eventual behavior of the model. For instance, many important datasets used in machine learning benchmarks predominantly feature lighter-skinned subjects (Buolamwini and Gebru 2018), indirectly introducing bias into the model. This type of latent bias cannot be detected by locally examining each piece of training data for harmful content. Furthermore, datasets often contain objectionable and biased content, including racist hate speech and sexually inappropriate images (Luccioni and Viviano 2021), even after undergoing cleaning procedures.

Even within the internal workings of a trained model, fundamental aspects like the meanings of common English words can encode gender biases (Bolukbasi et al. 2016). These biases can become more apparent as we move closer to the model's output. For example, the Stable Diffusion model disproportionately generates images of white male professionals (e.g., executives, doctors) and black or brown individuals in criminal or service roles (Ghosh and Caliskan 2023). Users have also reported that Generative AI tools often produce images where they appear younger, more sexu-

alized, or with lighter skin tones, even when these traits were not part of the input prompt (Ghosh and Caliskan 2023). Such biases are widespread across AI models and products, including Generative AI, from the training data to internal representations and outputs (Luccioni and Viviano 2021; Ghosh and Caliskan 2023; Cheng, Durmus, and Jurafsky 2023; Zhao et al. 2023).

A key issue in the AI development pipeline is that these tools are not created in isolation; they are embedded in complex human systems that influence their impact. There are disparities in who is creating, profiting from, and using these technologies (Stanford Institute for Human-Centered Artificial Intelligence 2023), which can lead to unforeseen consequences in their design, use, and societal role (Even-Tov et al. 2023; Jiang et al. 2023). Additionally, the companies behind these AI tools often keep details about their training data, architecture, and implementation secret (Ho 2024), forcing users to interact with them as black boxes (Jiang et al. 2023; Marcus and Southen 2024). This lack of transparency benefits the companies by allowing them to sidestep moral and ethical concerns regarding the model's impact and potential harms (Marcus and Southen 2024).

Methods

Survey Methodology

Our anonymous original survey of artists is an observational study (initial N=516 (includes incomplete answers), final N=459 (only complete answers)) of artists over the age of 18 years that was collected using a convenience sample in phase 1 and supplemented with a Qualtrics survey panel in phase 2 (demographic breakdown can be seen in Table 1. Qualtrics is a reputable academic survey panel platform that has been shown to yield reliable results in comparison to other survey platforms (Boas, Christenson, and Glick 2018). A comparison of the distributions of the US art workforce) based on the demographics is presented in Appendix B. Survey participants were recruited based on their self-identification as artists and must be over the age of 18. All survey questions are original and developed by the authors. Incomplete surveys were dropped from the analysis, leaving a final N of 459. Participants provided consent by agreeing to the survey information sheet. They were then prompted to respond to inquiries regarding their artistic identity and financial interactions within the art domain. Participants were presented with a brief introduction to AI-generated art, including an illustrative example, and were asked to indicate their familiarity with AI art tools and models. Subsequently, participants were queried on their perceptions regarding ownership of art produced by AI systems using a 5-point Likert scale. Further questions focused on artists' preferred compensation criteria for their work being used in AI-generated art models. Participants were also asked to evaluate the impact of AI art models on their field, their perception of these models as potential threats, and whether they believed creators of such models should disclose their data sources using a 5-point Likert scale. Finally, demographic information about the participants' backgrounds was collected. An appendix that enumerates all survey questions is available in Appendix A, which can be found in our supplementary materials (Lovato et al. 2024). This survey was approved as "exempt" by the University of Vermont Institutional Review Board (IRB).

Sampling limitations: In future work, we would like to sample a much bigger representative sample that specifically targets artistic practices and demographics. In addition, we are interested in how the dynamics of these opinions change over time with surveys that assess artist opinions at regular intervals.

Phase 1 Sampling: Survey participants for phase 1 (n = 252) of our observational study were recruited broadly on social media, using flyers locally at the University of Vermont, emailing artist groups, and direct contacts who are professional artists. Survey participants who are over the age of 18 and identify as an artist or work in an arts-related field were eligible. Phase 1 of the surveys was collected through a convenience sample from March 25, 2023, through September 8, 2023. Survey responses were collected on Qualtrics, an online survey platform.

Phase 2 Sampling: Phase 2 (n = 264) was collected via a targeted survey panel conducted by Qualtrics from September 12-23, 2023 on the Qualtrics online survey platform. The inclusion criteria for the study in Phase 2 were the same as in Phase 1.

Data

Our original survey of artists is an observational survey that took place over two sample phases (total N=459 completed surveys over both phases). We transformed all survey response variables of interest into numerical form to analyze our survey results (survey questions 8,9,10,13,14,15 as seen in Appendix A). All Likert survey questions were converted from 'Strongly Agree,' 'Agree,' 'Neutral,' 'Disagree,' and 'Strongly Disagree' to an ordinal scale of 1,2,3,4,5. For our ordinal logistic regression analysis, we combine 'Strongly Agree' and 'Agree' to equal 3, 'Neutral' to equal 2, and 'Disagree' and 'Strongly Disagree' to equal 1. In the original logistic regression, we reduced the dimensionality of the Likert scale to investigate the positive, neutral, and negative sentiment of the response questions in relation to our secondary variables.

All nominal and categorical variables (survey questions 1,2,3,4,5,11,12,16,17,18,19, as seen in Appendix A) were transformed into binary variables. Survey questions that included an "other" option or write-in answers were also categorized (e.g., in question 2, we added an additional art practice, which included digital art from the write-in answers).

Demographic questions (survey questions 16, 17, 18, 19 as seen in Appendix A) included asking survey participants their gender identity to identify whether they were white or a person of color, choosing their age range and country of residence. All demographic questions were transformed into binary variables.

Analytical Methods

Odds ratio In general, the odds of an event are the probability of success over the probability of failure. This can be

Type	Count (%)
Male	273 (59.48%)
Female	163 (35.51%)
Non-binary	23 (5.01%)
White	361 (78.65%)
Person of Color	98 (21.35%)
Age 30-49	233 (50.76%)
Age 18-29	104 (22.66%)
Age 50-64	79 (17.21%)
Age 65 or older	43 (9.37%)

Table 1: Descriptive statistics for our survey participants, N=459. A detailed table of descriptive statistics can be seen in Appendix B.

written as the following equation:

$$\frac{p}{1-p} \tag{1}$$

where p is the probability of success. Unlike probability, odds can be from 1 to positive infinity (Bruin 2011a). The odds ratio (OR, or θ) is simply the odds of one event divided by the odds of another. The OR can be from 0 to positive infinity, with an OR equal to 1 indicating equal odds of success between events. The natural logarithm of the OR is often used, as the transformation renders the distribution symmetric about 0 (Agresti 2012).

Ordinal Logistic Regression Ordinal logistic regression (OLR) is a generalization of regression (either multiple linear or binomial logistic) used for an ordered categorical response variable and any number of categorical or numeric secondary variables. OLR finds the logit (also known as log-odds) of the dependent variable (Y) being less than or equal to a certain level (denoted j) (Bruin 2011b). which can be written as:

$$logit(P(Y \leq j)) = log\frac{P(Y \leq j)}{P(Y > j)} \tag{2}$$

It is important to note that OLR assumes proportional odds, meaning the effect of the secondary variable(s) is the same between all levels of the dependent variable.

Results

459 participants were anonymously surveyed, yielding a sample reflective of the demographics (gender, ethnicity, age) of the broader American adult art workforce (See table 1). We conducted this observational study on the opinions of artists to explore the dynamic interplay between their perspectives on the potential benefits and risks of Generative AI art. This investigation delves into artists' viewpoints regarding the utility and risks associated with Generative AI art models, ethical considerations in disclosing artistic works used in AI art training models, the ownership and rights pertaining to AI art derivatives, and equitable compensation. Our results indicate a strong sentiment among artists in favor of disclosing what art is being used in AI training (RQ3), advocating against generative AI companies retaining ownership of AI-generated outputs (RQ4), and expressing concerns regarding AI's impact on the art industry's workforce (RQ1-2) and who profits from their art.

In the analysis of survey data pertaining to our primary research inquiries (see Table 2), 80.17% of the survey participants agree that model creators should be required to disclose in detail what art and images they use to train their AI models. We also find that 61.87% of survey participants agree that AI models are a threat to art workers. In contrast, 44.88% of survey participants agree that AI art models are a positive development in the field of art, notably with 27.02% of artists selecting neutral on the Likert scale for this question. Interestingly, we see an overlap between artists who agree that AI art is a positive development and those who also agree that it is a threat to art workers; 22.03% of survey participants agreed that both AI was a positive development and a threat, which presents an interesting tension which highlights that the sentiments of artists are nuanced and some hold complex opinions towards generative AI.

We examine a survey question prompting participants to consider a scenario where an AI art model was employed by another individual to create artwork distinctly resembling the participant's style and reflect on the rightful ownership of the resulting work and any derivative creations. Only 26.80% of the participants agreed that the work and its derivatives be considered the property of the AI model creators, while 39.22% of participants agreed that the work and its derivatives be considered the property of the person who used the AI model to generate the artwork. Finally, 41.39% of the participants agreed that the work and its derivatives be considered the property of the artists whose style is being represented by the AI model output. This result presents an interesting tension as well between artists who think the ownership should belong to the original inspiration of the AI art while others may also want to retain the ownership of the resulting work. We see in our results that artists who have used AI art models are more likely to agree that the resulting work should be considered the property of the person who used the model to create the work.

In the following sections, we will present results that identify the variables most strongly associated with survey participants who expressed agreement with the primary research questions.

We also look at what types of compensation artists consider fair in exchange for the use of their artwork to train an AI art model. The results for this question can be seen in Table 3. We find that 50.97% of survey respondents stated that they did not need compensation, but what matters to them is who is profiting from the use of their artwork. The fair compensation option selected most by our survey participants was "I don't need to profit, but I don't want for-profit-companies to profit from my art" (22.80% selected this option), which indicates that artists are not wholesale against art being used in AI art models training data, but they oppose for-profit-companies reaping the financial benefits. Our results show that 36.56% of respondents did state they would find some type of compensation fair, only 1.29% of participants felt that a tax on companies and individuals who profit from the AI art outputs would be a fair option, and 11.18% of participants felt that they would not feel comfortable with any of the compensation options listed.

Descriptive statistics about our survey responses to demo-

Question	Strongly Disagree	Disagree	Neutral	Agree	Strongly Agree
AI threat to art workers	6.32%	14.16%	17.65%	33.77%	28.10%
AI positive development for art	13.07%	15.03%	27.02%	22.88%	22.00%
Required disclosure	2.83%	5.23%	11.98%	32.03%	48.15%
Output owned by model creators	32.46%	24.62%	18.74%	15.25%	11.55%
Output owned by AI user	19.61%	22.22%	21.57%	20.92%	18.30%
Output owned by the artist	9.37%	26.58%	25.05%	22.00%	19.39%

Table 2: Participant answers to our primary research questions by percent (answered on a 5-point Likert scale), N=459.

Type	Counts (%)
I don't need to profit, but I don't want for-profit-companies to profit from my art	106 (22.8%)
I would donate the use of my artwork to train AI art models (no compensation)	82 (17.63%)
A flat fee	55 (11.83%)
A portion of any profit made from the model creators	51 (10.97%)
Not comfortable with any listed options	42 (9.03%)
A portion of any profit from people who profit from derivative works made using the model	32 (6.88%)
A portion of any profit made	32 (6.88%)
I don't need to profit, but I don't want anyone else to profit from my art	27 (5.81%)
I don't need to profit	22 (4.73%)
Other	10 (2.15%)
Tax	6 (1.29%)

Table 3: Counts percent per compensation type.

graphic questions can be seen in Table 1. A full table of descriptive statistics that outlines participants who have used AI art models, their familiarity with Generative AI, professional and artistic status, whether or not they have purchased art and art practices can be seen in Appendix B.

RQ1: Do artists see Generative AI models as a threat to art workers? This research question investigates artists' perceptions of Generative AI models as potential threats to art workers. Our findings reveal that 61.87% of survey participants hold the view that Generative AI models pose a threat to art workers. Thus, while a majority of artists perceive AI as a threat, it is not an overwhelming majority. This finding is related to survey question no. 15 (see Appendix A for survey questions).

An ordinal logistic regression (OLR) analysis was performed to examine the relationship between the Likert scale responses to RQ1 (dependent variable) and the secondary variables (all odds ratios can be seen in Appendix C), focusing on the odds ratios derived from the OLR model (as seen in Figure 1).

The results demonstrate a significant relationship (p-value < 0.05, all detailed odds ratios can be seen in Appendix C) between agreeing (represents a combination of both the responses 'Agree' and 'Strongly Agree' in the Likert scale) that Generative AI models are a threat to art workers and a subset of our secondary variables. Here, we see that artists ages 18-29 are more likely to agree that there is a threat and that artists who have used Generative AI art models to create art are less likely to think that it is a threat to the art workforce.

RQ2: Do artists see Generative AI art models as a positive development in the field of art? This research question investigates artists' perceptions of AI models as positive development to the field of art. We ran the same analysis for our second research question. We found that 44.88% of the survey participants agree that Generative AI art models are a positive development in the field of art. Notably, 27.02% of artists selected 'neutral' on the Likert scale for this question. This analysis pertains to survey question number 14 (see Appendix A for survey questions).

We conducted an ordinal logistic regression (OLR) analysis to investigate the association between Likert scale responses to RQ2 (dependent variable) and the secondary variables with a focus on the odds ratios obtained from the OLR model (as seen in Figure 2). Our analysis indicates a strong relationship (p-value < 0.05, all OLR detailed results can be seen in Appendix C) between agreeing that Generative AI art models are a positive development in the field of art and a subset of our secondary variables. Here, we see that 30-49 years olds, males, persons of color, participants who have used AI art models, and participants who said they would donate their artwork to train models were more likely to agree that generative AI models are a positive development for the field of art. Whereas makers (note that makers were a small n (n=5) with a fairly large range in the confidence interval bars represented in Figure 2), females, and those who identify as artists were less likely to agree that AI models are a positive development for the field of art.

RQ3: What types of disclosure for the use of artwork to train an AI art model do artists consider fair? This research question investigates artists' perceptions of what types of disclosure are fair when using their artwork to train an AI art model. Our results show a consensus among many artists that transparency and disclosure regarding the specific images used in training an AI model are essential prerequisites for them to deem the model fair and acceptable. We see from our survey results that 80.17% of the survey participants agree that creators of AI models should disclose the art and images used to train their models in detail. The analysis for RQ3 is related to survey question 13 (all sur-

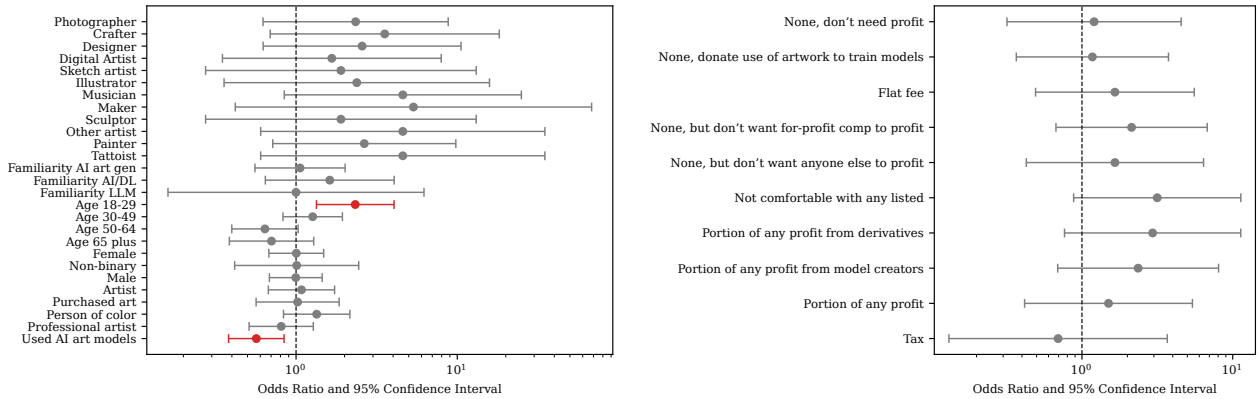

Figure 1: The ordinal logistic regression results show the association between secondary variables and answering 'Agree' or 'Strongly Agree' to RQ1 (Generative AI art models as a threat to art workers) and other survey variables. The x-axis represents odds ratios on a log scale, with 95% confidence intervals; red bars indicate significant results ($p < 0.05$) where the confidence interval does not cross the dotted line.

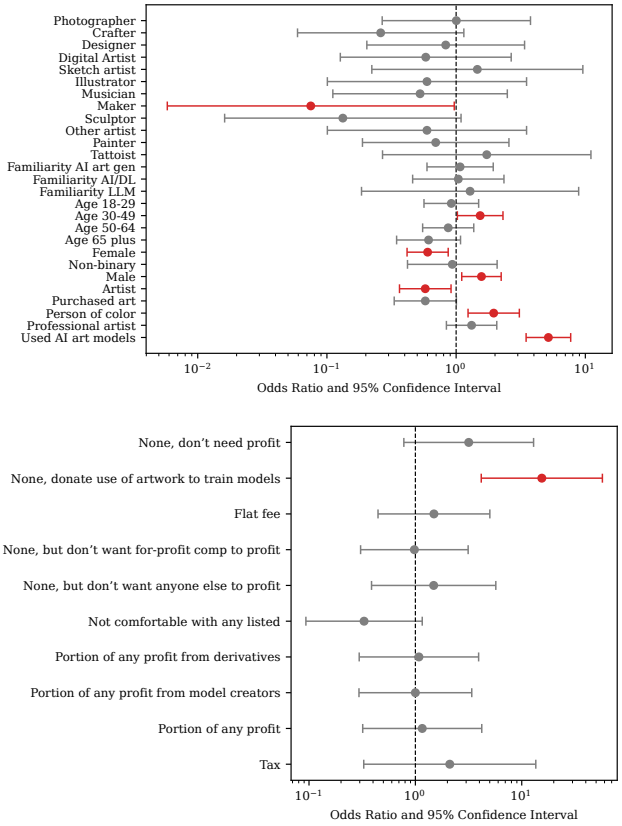

Figure 2: The ordinal logistic regression results show the association between secondary variables and answering 'Agree' or 'Strongly Agree' to RQ2 (Generative AI art models as a positive development). The x-axis represents odds ratios on a log scale, with 95% confidence intervals; red bars indicate significant results ($p < 0.05$) where the confidence interval does not cross the dotted line.

vey questions can be seen in Appendix A), which we used to run an ordinal logistic regression (OLR) to explore the relationship between the Likert scale answers to RQ3 and our secondary variables. The results of the OLR, displayed in Figure 3 as the odds ratio, show a strong relationship (p-value < 0.05, all odds ratios can be seen in Appendix C) between agreeing with the disclosure requirement and a subset of our secondary variables.

Here, we see that participants with a general familiarity with AI art or were between the ages of 18-49 were more likely to agree that AI model creators should be required to disclose their training data in detail. Participants who expressed willingness to accept compensation in the form of a share of the derivatives or profits generated from AI artwork inspired by their work, along with a portion of the model creators' profits, also agreed that disclosure should be required. Participants who identified as illustrators (note that illustrators were a small n (n=9) with a fairly large range in the confidence interval bars represented in Figure 3) or were between the ages of 50-64 years old were less likely to agree that disclosure should be required.

RQ4: From an artist's perspective, who should own the derivative artwork created by AI art models? Finally, we look at the question that asks participants to think about a scenario where an AI art model was used by someone else to produce artwork recognizably in their (the survey participant's) style and think about who should own the work and its derivatives. The analysis for RQ4 is related to survey questions No. 8, 9, and 10 (survey questions can be seen in Appendix A).

RQ4A: If an AI art model was used by someone else to produce artwork recognizably in your style (e.g., in the style of Van Gogh), Should that work and its derivatives be considered the property of the person who used the AI model to generate the artwork? Here, we explore the scenario where the owner of the Generative AI artwork should be the person who used the AI model to generate the

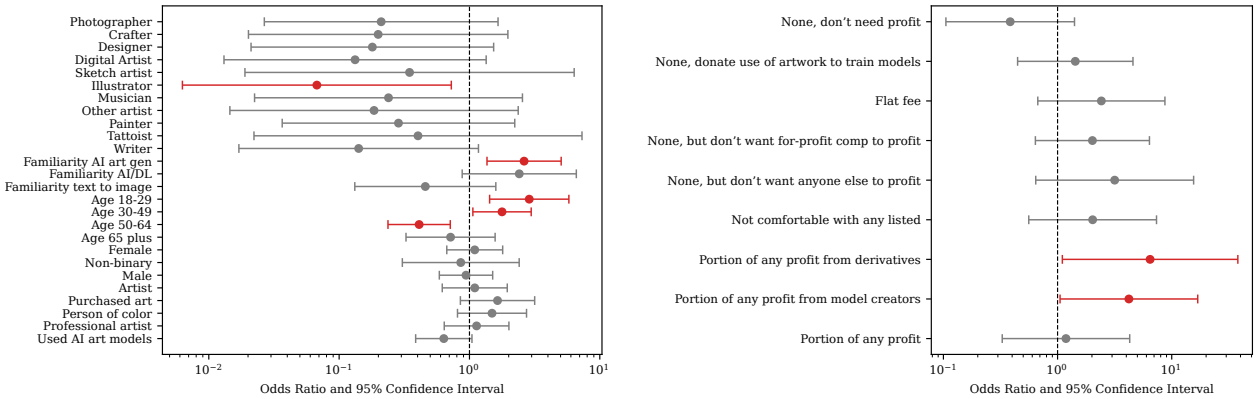

Figure 3: The ordinal logistic regression results show the association between secondary variables and answering 'Agree' or 'Strongly Agree' to RQ3 (Should model creators be required to disclose training data). The x-axis represents odds ratios on a log scale, with 95% confidence intervals; red bars indicate significant results ($p < 0.05$) where the confidence interval does not cross the dotted line.

artwork. We see that 39.22% of participants agreed that the work and its derivatives should be considered the property of the person who used the AI model to generate the artwork. We conducted an ordinal logistic regression (OLR) to investigate the relationship between our dependent variable (Likert scale responses to RQ4A) and a subset of our secondary variables, represented by the odds ratio depicted in Figure 4.

We see evidence of a strong association (p-value < 0.05) between agreeing and our secondary variables. Namely, males, persons of color, and participants who have used Generative AI to make art were more likely to agree that the owner of the artwork in this scenario should be the person who used the AI model to generate the artwork. Whereas females, participants who had purchased a piece of artwork, and participants who said that they were not comfortable with any type of compensation for the use of their artwork to train a generative AI model were less likely to agree that the owner of the artwork in this scenario should be the person who used the AI model to generate the artwork.

RQ4B: If an AI art model was used by someone else to produce artwork recognizably in your style (e.g., in the style of Van Gogh and you are Van Gogh), Should that work and its derivatives be considered yours? Here, we explore the scenario where the owner of the Generative AI artwork should be the artist whose style is being recognizable and mimicked by the Generative AI art model. Our survey shows that 41.39% of the participants agreed that the work and its derivatives should be considered their property (the artist whose style is being recognizable and mimicked by the Generative AI art model).

To investigate the relationship between the Likert scale responses to RQ4B and the secondary variables, we conducted an ordinal logistic regression (OLR), presented in Figure 5 displaying the odds ratios derived from the analysis. Our results show that there is a strong relationship (p-value < 0.05, all detailed odds ratios can be seen in Appendix C) between agreeing that the artist whose style is recognizable in the output should be the owner of the art and a subset of our secondary variables. We see that participants who are persons of color are more likely to agree that the owner of the Generative AI artwork should be the artist whose style is recognizable and mimicked by the Generative AI art model. Whereas participants who have used Generative AI to make art were less likely to agree that the owner of the Generative AI artwork should be the artist whose style is recognizable and mimicked by the Generative AI art model.

RQ4C: If an AI art model was used by someone else to produce artwork recognizably in your style (e.g., in the style of Van Gogh), Should that work and its derivatives be considered the property of the AI model creators? Finally, we look at our final scenario where the owner of the Generative AI artwork should be the AI model creators. According to our survey results, only 26.80% of the participants believed that the work and its derivatives should be considered the property of the creators of AI models. To examine the relationship between the answers to RQ4C on the Likert scale and other secondary variables, we conducted an ordinal logistic regression (OLR), as illustrated in Figure 6.

Our analysis suggests a strong relationship (p-value $<$ 0.05, all odds ratios can be seen in Appendix C) between agreement on AI companies owning the resulting artwork and a subset of our secondary variables. Participants who considered themselves tattooists (note that tattooists were a small n (n=9) with a fairly large range in the confidence interval bars represented in Figure 6), males, persons of color, or those who said they would donate their artwork to train generative AI models were more likely to agree that the owner of the Generative AI artwork should be the AI model creators. Whereas participants with a general familiarity with Generative AI Art, females, and non-binary participants were less likely to agree that the owner of the Generative AI artwork should be the AI model creators.

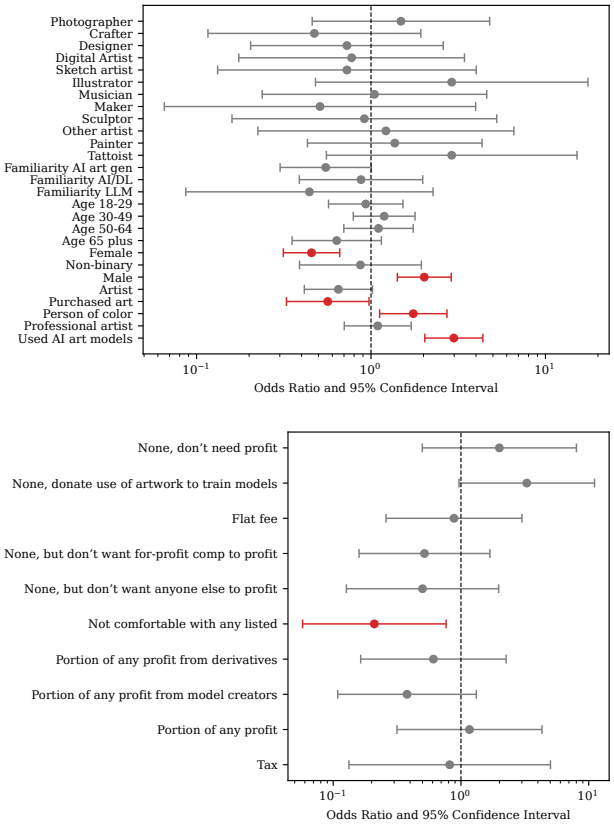

Figure 4: Ordinal logistic regression results showing the odds ratios for answering 'Agree' or 'Strongly Agree' to RQ4A regarding secondary variables and whether AI-generated art in your recognizable style should be considered the property of the AI user. The x-axis represents odds ratios on a log scale, with 95% confidence intervals; red bars indicate significant results ($p < 0.05$) where the confidence interval does not cross the dotted line.

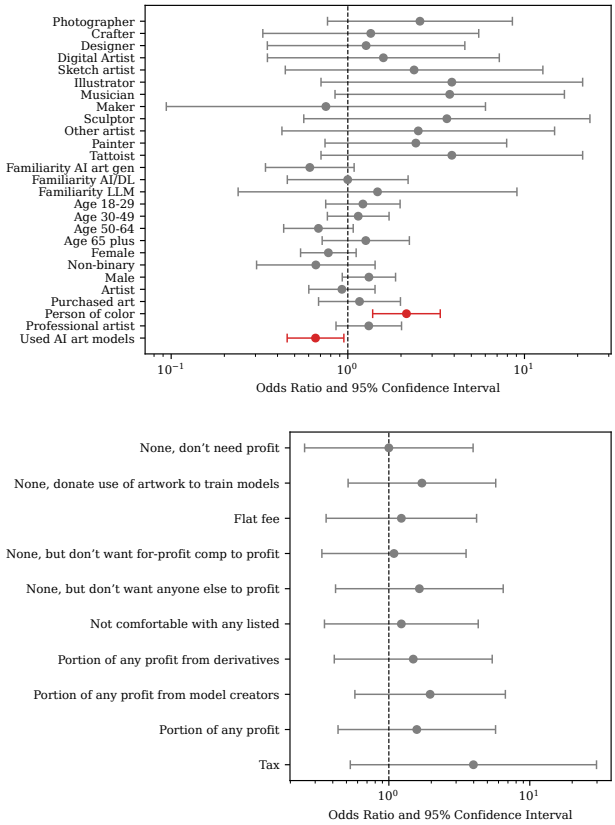

Figure 5: Ordinal logistic regression results showing the odds ratios for answering 'Agree' or 'Strongly Agree' to RQ4B regarding secondary variables and whether AI-generated art in your recognizable style should be considered your property. The x-axis represents odds ratios on a log scale, with 95% confidence intervals; red bars indicate significant results ($p < 0.05$) where the confidence interval does not cross the dotted line.

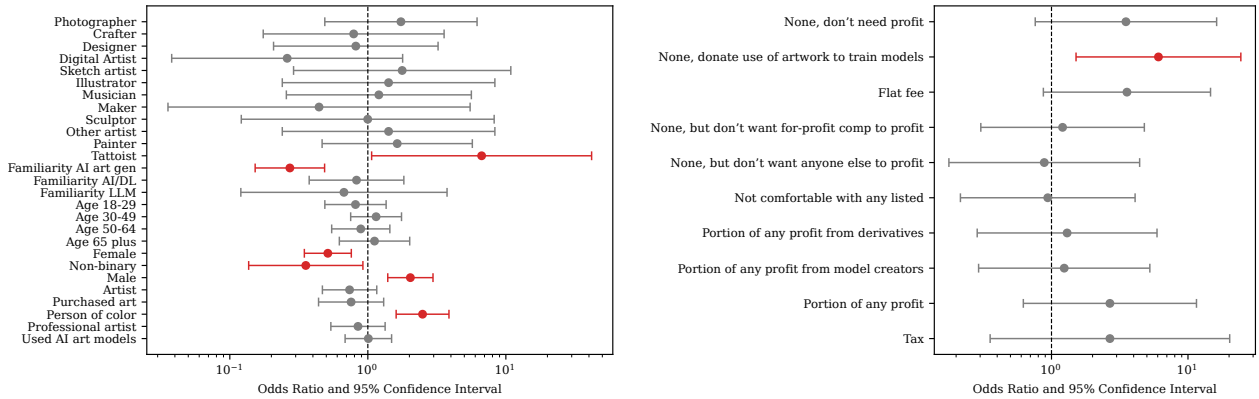

Figure 6: Ordinal logistic regression results showing the odds ratios for answering 'Agree' or 'Strongly Agree' to RQ4C regarding secondary variables and whether AI-generated art in your recognizable style should be considered the property of the AI model creators. The x-axis represents odds ratios on a log scale, with 95% confidence intervals; red bars indicate significant results ($p < 0.05$) where the confidence interval does not cross the dotted line.

Discussion

The impact of Generative AI on artists is a complex global issue with consequences in many arenas, and it must be investigated through further multidisciplinary collaboration. We hope this work contributes to the existing literature on the ethics of AI-generated art and helps to provide a voice for artists on this topic.

In this study, we surveyed how artists feel about Generative Artificial Intelligence (AI) in art. Most participants believe that AI model creators should be required to disclose the art and images used to train them – an aspect of Generative AI tools that is far from transparent. However, only a minority of participants see AI art models as a positive development in art, while the majority see them as a threat to art workers.

When asked who should own the work and its derivatives when an AI art model is used to produce artwork in a particular artist's style, the minority of participants believed that the AI model creators should own the work. In contrast, the majority of survey respondents were split between thinking it should belong to the person who used the AI model or the artist whose style is represented in the output.

Regarding compensation for using their artwork to train AI art models, over half of the participants did not require any compensation but were concerned about who was profiting from their artwork. One concern is that artists will be reluctant to share their work with each other and with the public going forward, and there may be resistance to perceived exploitation of that practice, especially in the service of profit. We think this interpretation is consistent with the result that some participants stated that they do not require profit but do not want for-profit companies to profit from their art. Others considered some compensation fair, while a minority felt taxing companies and individuals who profit from AI art outputs would be a fair option. Finally, some participants stated they would feel uncomfortable with the compensation options listed. Understanding how the field of Generative AI falls short in respecting the values of the people who make the cultural artifacts it relies upon – and how that could be changed – is an important ethical question that we hope to expand on in future work.

In our future work, we plan to expand the sample size to a more representative sample. This will increase our study's statistical power and allow us to look at how opinions change over time. Moreover, we would like to put the current flourishing of Generative AI within the context of art history and taxonomy. We believe that a longer timeline would help consolidate, unify, and clarify the relationship between technology and art. While it is essential to understand what artists think about Generative AI, we need to collaborate with experts from various fields, including art, history, philosophy, psychology, cognitive science, and others, to understand the implications of Generative AI on our understanding of art and its role in our lives.

Appendices and supplementary materials for this paper can be found on the arXiv version of this paper (Lovato et al. 2024). Data and code can be found at the following URL: `https://github.com/juniperlovato/AASP`

Acknowledgements

We are thankful to the anonymous reviewers for their insightful comments. Institutional Review Board Approval: The survey in this project is CHRBSS (Behavioral) STUDY00002379, approved by the University of Vermont I.R.B. on 3/14/2023. The authors would like to thank Maria Sckolnick, John Meluso, Kendall Fortney, Randall Harp, Laurent Hébert-Dufresne, Jean-Gabiel Young, and Winn Gillen, for their comments on early versions of this work. The Alfred P. Sloan Foundation award #2024-22498, The National Science Foundation award #2242829, and the MassMutual Center of Excellence support this work. Any opinions, findings, conclusions, or recommendations expressed in this material are those of the authors and do not reflect the views of the aforementioned financial supporters.

References

Abdin, M. I.; Gunasekar, S.; Chandrasekaran, V.; Li, J.; Yuksekgonul, M.; Peshawaria, R. G.; Naik, R.; and Nushi, B. 2023. KITAB: Evaluating LLMs on Constraint Satisfaction for Information Retrieval. Accepted for a poster presentation at The Twelfth International Conference on Learning Representations (ICLR 2024), arXiv:2310.15511.

Agresti, A. 2012. *Categorical data analysis*, volume 792. New Jersey, USA: John Wiley & Sons.

Alekseeva, L.; Azar, J.; Gine, M.; Samila, S.; and Taska, B. 2021. The demand for AI skills in the labor market. *Labour economics*, 71: 102002.

Barnett, J. 2023. The Ethical Implications of Generative Audio Models: A Systematic Literature Review. In *AAAI/ACM Conference on AI, Ethics, and Society*, AIES '23, 146–161. New York, NY, USA: Association for Computing Machinery. ISBN 9798400702310.

Boas, T. C.; Christenson, D. P.; and Glick, D. M. 2018. Recruiting large online samples in the United States and India: Facebook, Mechanical Turk, and Qualtrics. *Political Sci. Res. Methods*, 8(2): 232–250.

Bolukbasi, T.; Chang, K.-W.; Zou, J. Y.; Saligrama, V.; and Kalai, A. T. 2016. Man is to computer programmer as woman is to homemaker? debiasing word embeddings. *Advances in Neural Information Processing Systems*, 29: 4349–4357.

Bruin, J. 2011a. How do I interpret odds ratios in logistic regression? https://stats.oarc.ucla.edu/other/multpkg/faq/general/faq-how-do-i-interpret-odds-ratios-in-logistic-regression/.

Bruin, J. 2011b. How Do I Interpret the Coefficients in an Ordinal Logistics Regression in R? https://stats.oarc.ucla.edu/r/dae/ordinal-logistic-regression.

Bubeck, S.; Chandrasekaran, V.; Eldan, R.; Gehrke, J.; Horvitz, E.; Kamar, E.; Lee, P.; Lee, Y. T.; Li, Y.; Lundberg, S.; Nori, H.; Palangi, H.; Ribeiro, M. T.; and Zhang, Y. 2023. Sparks of Artificial General Intelligence: Early experiments with GPT-4. arXiv:2303.12712.

Buolamwini, J.; and Gebru, T. 2018. Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification. *Proceedings of Machine Learning Research*, 81: 1–15. Presented at the Conference on Fairness, Accountability, and Transparency.

Caporusso, N.; et al. 2023. Generative Artificial Intelligence and the Emergence of Creative Displacement Anxiety. *Research Directs in Psychology and Behavior*, 3(1).

Cetinic, E.; and She, J. 2022. Understanding and Creating Art with AI: Review and Outlook. *ACM Trans. Multimedia Comput. Commun. Appl.*, 18(2).

Cheng, M.; Durmus, E.; and Jurafsky, D. 2023. Marked Personas: Using Natural Language Prompts to Measure Stereotypes in Language Models. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics*, 1504–1532.

Chow, P.-S. 2020. Ghost in the (Hollywood) machine: Emergent applications of artificial intelligence in the film industry. *NECSUS_European Journal of Media Studies*, 9(1): 193–214.

Etro, F. G.; and Galenson, D. W. 2023. The economics of art history. *Journal of Cultural Economics*, 47(3): 355–358.

Even-Tov, O.; Huang, K.; Trueman, B.; Bogard, J. E.; and Goldstein, N. J. 2023. Sharing names and information: Incidental similarities between CEOs and analysts can lead to favoritism in information disclosure. *PNAS*, 120. E2311250120.

Ferrara, E. 2024. GenAI against humanity: Nefarious applications of generative artificial intelligence and large language models. *Journal of Computational Social Science*, 1–21.

Ghosh, S.; and Caliskan, A. 2023. "Person" == Light-skinned, Western Man, and Sexualization of Women of Color: Stereotypes in Stable Diffusion. In *Conference on Empirical Methods in Natural Language Processing*.

Gmyrek, P.; Berg, J.; and Bescond, D. 2023. Generative AI and jobs: A global analysis of potential effects on job quantity and quality. Technical Report 96, International Labour Organization.

Gozalo-Brizuela, R.; and Garrido-Merchán, E. C. 2023. A survey of Generative AI Applications. arXiv:2306.02781.

Grand, G.; Blank, I.; Pereira, F.; and Fedorenko, E. 2022. Semantic projection recovers rich human knowledge of multiple object features from word embeddings. *Nature Human Behaviour*, 6: 975–987.

Henderson, F. D. 2010. It's your own fault: How post-strike Hollywood continues to punish writers for striking. *Popular Communication*, 8(3): 232–239.

Ho, K. K. 2024. Database of 16,000 Artists Used to Train Midjourney AI, Including 6-Year-Old Child, Garners Criticism. https://www.artnews.com/art-news/news/midjourney-ai-artists-database-1234691955/. Accessed on January 5, 2024.

Huyghe, R., ed. 1970. *Larousse Encyclopedia of Prehistoric and Ancient Art*. Hamlyn House, Feltham, Middlesex, England: The Hamlyn Publishing Group Limited. Originally published in 1957. Text prepared by Emily Evershed, Michael Heron, Corinne Lambert, Hugh Newbury and Wendela Schurmann from the French original.

Jiang, H. H.; Brown, L.; Cheng, J.; Khan, M.; Gupta, A.; Workman, D.; Hanna, A.; Flowers, J.; and Gebru, T. 2023. AI Art and Its Impact on Artists. In *AAAI/ACM Conference on AI, Ethics, and Society*, AIES '23, 363–374. New York, NY, USA: Association for Computing Machinery. ISBN 9798400702310.

Karson, J. 2023. Five Guardrails for Artists Working with Machines and AI. https://jennkarson.studio/2023/07/five-guide-rails-for-artists-working-with-machines-and-a-i/. Retrieved on April 11, 2024.

Lovato, J.; Zimmerman, J.; Smith, I.; Dodds, P.; and Karson, J. 2024. Foregrounding Artist Opinions: A Survey Study on Transparency, Ownership, and Fairness in AI Generative Art. *arXiv preprint arXiv:2401.15497*.

Luccioni, A.; and Viviano, J. D. 2021. What's in the Box? A Preliminary Analysis of Undesirable Content in the Common Crawl Corpus. *ACL Anthology*.

Marcus, G.; and Southen, R. 2024. Generative AI Has a Visual Plagiarism Problem: Experiments with Midjourney and DALL-E 3 show a copyright minefield. https://spectrum.ieee.org/amp/midjourney-copyright-2666872100. Accessed on January 10, 2024.

Mazzone, M.; and Elgammal, A. 2019. Art, creativity, and the potential of artificial intelligence. In *Arts*, volume 8, 26. MDPI.

Nissenbaum, H. 2004. Privacy as contextual integrity. *Wash. L. Rev.*, 79: 119.

Pfeifer, R.; and Bongard, J. 2006. *How the Body Shapes the Way We Think: A New View of Intelligence*. London, England and Cambridge, Massachusetts: A Bradford Book, MIT Press. ISBN 978-0262162395.

Schrader, A. 2023. An Artist Invited Blind People to Use an A.I. Image Generator. The Unsettling Results Could Help Make Art More Accessible. https://news.artnet.com/art-world/ai-art-experiments-blind-users-2368599/amp-page. Retrieved on November 28, 2023.

Shan, S.; Ding, W.; Passananti, J.; Zheng, H.; and Zhao, B. Y. 2023. Prompt-specific poisoning attacks on text-to-image generative models. *arXiv preprint arXiv:2310.13828*.

Squicciarini, M.; and Nachtigall, H. 2021. Demand for AI skills in jobs: Evidence from online job postings.

Stanford Institute for Human-Centered Artificial Intelligence. 2023. Artificial Intelligence Index Report. Technical report, Stanford University, Cordura Hall, 201 Panama Street, Stanford University, Stanford, CA 94305. Retrieved on October 30, 2023.

US District Court for the District of Columbia. 2023. Thaler v. Perlmutter, Civil Action No. 22-1564. (United States District Court for the District of Columbia; August 18, 2023).

US Federal Trade Commission. 2023. Creative Economy and Generative AI Roundtable. https://www.ftc.gov/news-events/events/2023/10/creative-economy-generative-ai.

Walker, J.; Thuermer, G.; Vicens, J.; and Simperl, E. 2023. AI Art and Misinformation: Approaches and Strategies for Media Literacy and Fact Checking. In *AAAI/ACM Conference on AI, Ethics, and Society*, AIES '23, 26–37. New York, NY, USA: Association for Computing Machinery. ISBN 9798400702310.

Xu, Y.; Liu, Z.; Tian, Y.; Tong, S.; Tegmark, M.; and Jaakkola, T. 2023. Pfgm++: Unlocking the potential of physics-inspired generative models. In *International Conference on Machine Learning*, 38566–38591. PMLR.

Zhao, W. X.; Zhou, K.; Li, J.; Tang, T.; Wang, X.; Hou, Y.; Min, Y.; Zhang, B.; Zhang, J.; Dong, Z.; Du, Y.; Yang, C.; Chen, Y.; Chen, Z.; Jiang, J.; Ren, R.; Li, Y.; Tang, X.; Liu, Z.; Liu, P.; Nie, J.-Y.; and Wen, J.-R. 2023. A Survey of Large Language Models. arXiv:2303.18223.