Laypeople's Egocentric Perceptions of Copyright of AI-Generated Art

GABRIEL LIMA, Max Planck Institute for Security and Privacy, Germany
NINA GRGIĆ-HLAČA, Max Planck Institute for Software Systems & Max Planck Institute for Research on Collective Goods, Germany
ELISSA M. REDMILES, Georgetown University, United States

Recent breakthroughs in generative AI (GenAI) have fueled debates concerning the status of AI-generated creations under copyright law. This research investigates laypeople's perceptions (N = 424) of AI-generated art concerning factors associated with copyright protection. Inspired by prior work suggesting that people show egocentric biases when evaluating their own creative outputs, we also test if the same holds for AI-generated art. Namely, we study the differences between the perceptions of those who have something to gain from copyright protection—creators of AI-generated art—and uninvested third parties.

To answer our research questions, we held an incentivized AI art competition, in which some participants used a GenAI model to generate images for consideration while others evaluated these submissions. We find that participants are most likely to attribute authorship and copyright over AI-generated images to the users who prompted the AI system to generate the image and the artists whose creations were used for training the AI model. We also find that participants egocentrically favored their own art over other participants' art and rated their own creations higher than other people evaluated them. Moreover, our results suggest that people judge their own AI-generated art more favorably with respect to some factors (creativity and effort) but not others (skills). Our findings have implications for future debates concerning the potential copyright protection of AI-generated outputs.

1 INTRODUCTION

Recent breakthroughs in generative artificial intelligence (GenAI) have pushed the boundaries of what machines can generate across various domains. Text-based models [6] have become pervasive online, enabling users to generate long-form texts through short prompts and revolutionizing how one searches for information online. Similarly, AI image models [11] have empowered its users to generate highly realistic and detailed images from short textual descriptions. Multimodal GenAI models [99] show even more promise, generating text, images, audio, and other types of data. These systems learn from large datasets of human creations how to generate text, images, and other content that may be indistinguishable from their human-created counterparts [50, 60, 71].

Although GenAI models have the potential to revolutionize how humans express their creativity, they also pose novel challenges to society. One domain that has received considerable attention in the context of GenAI is copyright law. Copyright law regulates works of authorship, such as paintings and novels, and determines who should have exclusive rights over these creations. Extensive literature has examined how copyright law should address AI-generated works (e.g., [32, 37, 41, 64, 96], and several lawsuits are currently underway [28, 83]) to determine whether training on copyrighted material violates the law and if AI-generated content warrant the same legal protection as human-created works.

Discussions surrounding copyright protection of AI-generated outputs have been primarily normative (e.g., [37, 41]), with little focus on capturing the opinion of GenAI users. Scholars debate whether AI-generated outputs are eligible for copyright protection [96] and, if so, who should hold the rights associated with this protection [32]. Here, we investigate *laypeople's expectations* of copyright law in relation to GenAI outputs.

Authors' addresses: Gabriel Lima, Max Planck Institute for Security and Privacy, Germany, gabriel.lima@mpi-sp.org; Nina Grgić-Hlača, Max Planck Institute for Software Systems & Max Planck Institute for Research on Collective Goods, Germany, nghlaca@mpi-sws.org; Elissa M. Redmiles, Georgetown University, United States, elissa.redmiles@georgetown.edu.

There are several reasons why studying lay opinions concerning the intersection of copyright and GenAI is relevant. First, capturing laypeople's perceptions of the law more generally is important to ensure that it is democratically legitimate [102]. Democratic theories of law argue that the law should reflect laypeople's intuitions [102] to motivate citizens to comply with it [103]. This paper's approach can thus help ensure that future legal decisions and policymaking are aligned with public expectations. Even if the law is at odds with lay intuitions, capturing laypeople's opinions can help mitigate any potential backlash that may emerge from these differences, proposing ways to bridge these gaps [9].

Second, aligning laypeople's legal intuitions with the law is particularly important in the context of copyright. Copyright law has several objectives [64], such as promoting fairness [25] and safeguarding creators' moral rights [25]. Copyright also aims to incentivize creativity through financial incentives. By granting exclusive rights over creative works to authors, copyright law attempts to incentivize them to continue exercising their creativity for further financial benefit [73, 74]. Because copyright law depends on this behavioral response to achieve its objective (i.e., to promote creativity), it is all-important that potential creators, namely laypeople, understand how the law works [73].

Finally, empirical studies capturing the opinion of *laypeople* about copyright is distinctively important in the context of GenAI. GenAI aims to democratize the ability to create works that could be eligible for copyright protection—something that used to be restricted to skilled artists—enabling non-artists to participate in the creative society.

This paper examines laypeople's perceptions (N = 424) of AI-generated images vis-à-vis their potential copyright protection. We first capture how laypeople judge GenAI images with respect to factors that help determine whether human creations are eligible for copyright in different jurisdictions [32]. Furthermore, we investigate whether laypeople believe AI-generated images warrant copyright and, if so, who should own it. More specifically, we address the following research questions:

RQ1) How do laypeople evaluate AI-generated images concerning the creativity, effort, and skills involved in the creation process?

RQ2) Who do laypeople believe are the authors of AI-generated images?

RQ3) Who do laypeople believe should hold the rights to 1) display and 2) make copies of AI-generated images?

To investigate laypeople's perceptions of AI-generated art, we conducted an experimental study in the form of a juried AI art exhibition. To guide our experimental design, we leveraged prior research that found that people exhibit egocentric biases in their judgments of creative works [13, 14, 82], particularly if owned or created by them [97]. We hypothesize that similar biases may emerge in people's perceptions of copyright, since they relate to questions surrounding ownership of creative works.

Participants engaged in the exhibition either as creators (by using an GenAI model to create art), invested evaluators (by generating art and evaluating other people's submissions), or uninvested evaluators (by only evaluating others' images). Figure 1 presents a high-level overview of our experimental design. Our between-subjects design allowed us to study how perceptions about copyright vary between those who have something to gain from copyright protection and the exhibition—creators of AI-generated art—and uninvested third parties. For instance, we hypothesized that creators of AI-generated art will egocentrically overestimate the quality of their own creations and exhibit greater support for the exclusive ownership structure that copyright protection could afford them.

Our study was designed to maximize ecological validity. The decision to hold an AI art exhibition was inspired by real-world examples of AI-generated images winning art exhibitions [90] and competitions focused solely on AI art [7]. Our AI art exhibition rewarded the top-10 best submissions, mimicking some of the financial incentives involved in

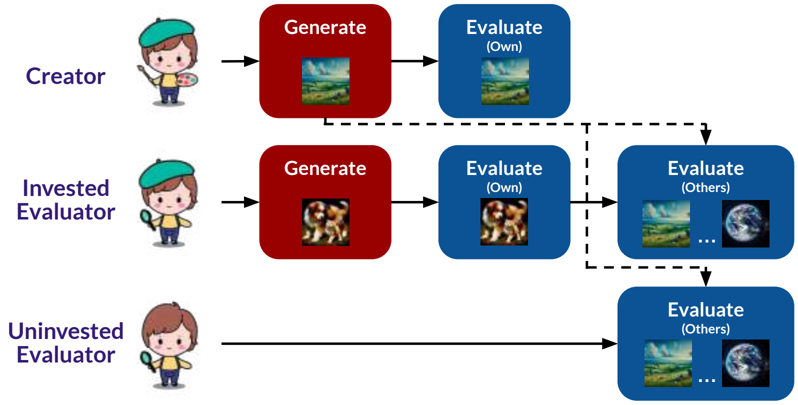

Fig. 1. High-level overview of our experiment.

copyright decisions. It also simulates other non-monetary incentives, such as exposure and recognition; selected images were displayed in a website, in which participants had the choice to attach their names to their creations.[1]

Our results suggest that people believe creativity and effort—but not necessarily skills—are necessary to create art using AI. Participants also indicated that users and those whose creations were used to train the GenAI model should be considered authors of AI-generated images and enjoy the rights to display and make copies of them. In contrast, people were less likely to attribute authorship and the rights associated with copyright protection to the AI model itself and the company that developed it.

Focusing on how people's perceptions vary egocentrically, we found that participants evaluate the process of using GenAI to generate art egocentrically with respect to some factors—e.g., creativity and effort—but not others—such as skills. When deciding who should win the art exhibition's monetary award, however, participants judged their own creations much more favorably than art generated by others, supporting our egocentric hypotheses. Surprisingly, we found the opposite trend in attributions of authorship to the user of the GenAI model, such that evaluators attributed more authorship to creators than creators themselves. Participants' opinions concerning who should hold the rights to display and make copies of AI-generated images did not vary by their role in the AI art exhibition.

Our research has implications for the deployment of GenAI models and their future regulation under copyright law. Our findings call for the consideration of a more distributed ownership structure of copyright, under which training data contributors are also recognized as authors and rights-holders. People's attribution of authorship and rights to data contributors rather than the company that developed the AI model raise questions concerning current business models that concentrate profits in corporate entities at the expense of human artists [18, 76]. We discuss how existing legal doctrines (e.g., neighbouring rights, licensing models) and computer science research could help ensure that training data contributors are compensated.

Our findings suggest that egocentric biases become relevant in perceptions of GenAI outputs and their associated copyright when monetary incentives come into play. Although participants did not prioritize themselves when asked to indicate who should hold *hypothetical* rights over their creations, they overestimated the quality of their images when that determined *real* monetary rewards. Our evidence of egocentric biases suggests that some conflicts of interest may arise in discussions surrounding the legal status of AI-generated art under copyright law.

[1] https://thegcamilo.github.io/AI-art-exhibition/

2 BACKGROUND

Generative AI (GenAI) has the potential to revolutionalize how humans exercise their creativity. However, it does not come without problems. Several reports indicate that GenAI has the potential to amplify harmful stereotypes [81, 101]. Researchers have also warned how generative models could fuel online mis- and disinformation by generating false, yet credible-looking, news and online profiles [8]. These models' tendency to fabricate information while sounding confident and knowledgeable can also distort human beliefs [58]. There exists evidence that GenAI can produce misinformation that is more compelling to readers [98], as well as manipulate people's beliefs in conspiracy theories [23].

One particular domain that has been directly impacted by the emergence of GenAI is intellectual property (IP) law. IP law refers to the rules that regulate the rights associated with human creations, such as inventions and literary and artistic works, determining who should control and benefit from them. How to deal with AI-generated outputs regarding IP rights remains an open question. For instance, should machine-generated artistic works be protected similarly to their human-created counterparts? Who should enjoy the rights that would normally be associated with this creation?

In this paper, we focus on copyright law, which is the branch of IP law that covers works of authorship, including artistic, musical, and literary works, such as novels, movies, songs, and many other human creations.[2] Although what qualifies for copyright protection and what rights follow this determination vary by country, copyright law mainly grants some exclusive rights (e.g., to reproduction and distribution) to the copyright holder for a predetermined period of time. In the case of human creations, copyright owners are often the creators themselves, with exceptions in case the work has been created within the scope of employment.

Debates on how to address novel technologies with copyright law is not a new development. The rise of photography also challenged what was eligible for copyright protection [46]. Similarly, the law was initially unprepared to deal with the emergence of digital art [36] and video-sharing platforms that allowed users to upload copyrighted content without much restriction, leading to solutions that ensure that the creator rights are protected (e.g., Google's content ID [56]). Now it is the time of GenAI to contest what warrants (or not) copyright protection.

The challenges posed by GenAI to copyright law can be grouped in three overarching questions [35], all of which we discuss below: 1) does training AI models on copyrighted data infringe on the copyright of the training data?; 2) are AI-generated outputs eligible for copyright protection; and if so 3) who owns the copyright? In this paper, we focus on the two latter questions. We also discuss some potential alternative regulatory frameworks for GenAI outputs, as well as prior work on lay perceptions of copyright and GenAI.

2.1 Does Training GenAI Infringe on the Copyright of the Training Data?

GenAI models require large amounts of data to be trained. These datasets may contain copyrighted data, raising the question of whether the training of AI models infringes on the rights of copyright holders. Those arguing that training GenAI on copyrighted material should be illegal often posit that it exploits authors of the training data without compensating them [19], with some going even further and equating the practice to theft [18, 76]. These critics often defend that the owners of training data should be compensated [17].

In contrast, proponents of GenAI assert that training does not infringe on the copyright of its training data. Such arguments in favor of GenAI often rely on the United States' fair use doctrine [65], which permits limited use of

[2]Discussions surrounding GenAI and IP law are not restricted to copyright. For instance, some have argued that AI should be treated as an inventor under patent law [5], which deals with the rights associated with human inventions. This argument has been successful in Australia [24] and, at the same time, has found challenges in the United States [27].

copyrighted materials under specific conditions, and the text and data mining copyright exceptions in the European Union (EU) [94]. The wide availability of GenAI models have triggered several lawsuits that are currently underway (e.g., [83]) and will decide whether training GenAI with copyrighted material without compensating its owners is legal under copyright law.

2.2 Can AI-Generated Outputs Be Protected by Copyright?

Another question raised by GenAI concerning copyright law refers to whether AI-generated outputs should receive the same protection that are assigned to works that are created solely by humans. If works generated with the assistance of an AI model are eligible for copyright protection, someone would have exclusive rights over it; on the other hand, if they are not eligible, these outputs would be part of the public domain, meaning that anyone would be able to use these works without permission.

Scholars largely disagree on this particular question. Some argue extending copyright law to machine-created works would reduce the value of human creativity [76], flood the market with creations of questionable quality [37], and concentrate power in the hands of a few [17, 47]. In contrast, proponents of extending copyright to AI-generated works defend that protecting these outputs could promote innovation by incentivising research and development of AI [44, 61], as well as enable users to create works that would not be possible without GenAI.

Current legal decisions addressing whether AI-generated outputs should be granted copyright have varied widely across different countries. While a judge in the United States (US) has ruled that AI-generated images are not eligible for copyright protection [28], the Beijing Internet Court has taken a different approach by granting copyright to AI-generated art [107]. Even within China, different jurisdictions have made conflicting copyright decisions regarding GenAI outputs [106]. These rulings demonstrate how the requirements for copyright protection, and hence the answer to whether AI-generated outputs are eligible, may vary depending on the jurisdiction [32, 94]. While the US Copyright Office requires originality for copyright protection, meaning that a work must be independently created and exhibit a modicum of creativity [4], the EU posits a work is eligible if it is the result of the "author's own intellectual creation" [45, 96]. In contrast, other countries (e.g., Australia [1] and Canada [3]) also consider whether there was a non-trivial exercise of skill and effort [32, 41]. The United Kingdom (UK) is one of the few countries with clear rules for "computer-generated works," granting exclusive rights to "the person by whom the arrangements necessary for the creation of the work are undertaken" [104]. **Motivated by these conflicting viewpoints, this paper explores how laypeople judge AI-generated images with respect to the creativity, effort, and skills involved in the creation process (RQ1)**.

2.3 Who Owns the Potential Copyright of AI-Generated Outputs?

If an AI-generated work is granted copyright, some legal entity would have exclusive rights over it; but who would that be? A GenAI output is the result of a collaboration between several actors, including the model's developers, its users, those who potentially own the training data, and the AI model itself, making it difficult to determine who should own it.[3]

A common proposition is that the user of a GenAI model, i.e., the person who gave it instructions, should be granted exclusive rights over its outputs [32]. However, it is not clear whether this user would satisfy the conditions that determine authorship and, hence, copyright ownership. The US Copyright Office has stated that merely prompting an AI model does not qualify the user for authorship; instead, it proposes a case-by-case analysis that would determine

[3]This question is an instance of the "problem of many hands" [105], which posits that it is difficult to determine who is ultimately responsible for collective actions. Scholars have also explored how AI may complicate this search for a responsible actor, particularly when it causes harm [21, 75].

whether the work contains "sufficient human authorship" (e.g., whether the user "select[ed] or arrange[d] AI-generated material in a sufficiently creative way" [84]). Another possibility would be granting copyright to the "the person by whom the arrangements necessary for the creation of the work are undertaken," as proposed by the UK Intellectual Property Office [104], which could interpreted as the developer of the AI model [48]. Other proposals include considering the AI model itself as an author (and thus copyright holder) [70, 96] or granting a form of joint authorship (i.e., collective ownership) to the many entities involved [57]. **This research investigates who laypeople consider to be authors of AI-generated images (RQ2), as well as who they believe should have the rights to display and make copies of them (RQ3).**

2.4 Alternatives to Copyright Law

Although copyright law often takes central stage in the discussion surrounding ownership of GenAI outputs, it is not exempt from critiques. Copyright law is often criticized by its potential to hinder innovation by concentrating power in monopolies [48, 66]. Furthermore, extending copyright law to GenAI could conflict with its primary objective—to promote *human* creativity; if AI-generated works are protected, they could devalue human creativity by flooding the market with artificial competition to human creations [47].

Scholars have proposed different approaches for AI-generated works to promote innovation. For instance, the reproduction of GenAI models and their outputs could be restricted under a *suis generis* doctrine, which grants exclusive rights in situations in which there has been substantial financial investment [96]. Someone other than the author—which is hard to define in the case of AI-generative works, as discussed above—could also have "neighboring" rights over AI outputs as a way to protect their investment [47, 63]. AI-generated outputs could also be covered by other branches of IP law, such as trademark law, which relies on owners maintaining and enforcing their rights [48]. Another approach could be moving away from copyright law's focus on creativity and originality as grounds for ownership, and instead reward actors who put effort and demonstrate skill when using GenAI [41].

2.5 Lay Perceptions of Copyright

Copyright protection exists to achieve several objectives [64]. For instance, it promotes fairness by granting authors the right to exclusively control the fruits of their own labor [10]. It also safeguards moral rights of creators, protecting the emotional bond between authors and their works [25]. Most relevant to this research is copyright law's aim to incentivize creativity. Copyright law attempts to promote creativity by granting exclusive rights over creative works to authors. These exclusive rights determine that only authors can profit from their creations, incentivizing them to continue exercising their creativity for further financial benefit [73, 74]. Unless potential authors and rights-holders (i.e., laypeople) understand and agree with what is eligible for copyright protection and what rights are associated with it, copyright law may fail to incentivize the production of creative outputs. Hence, examining public perceptions and expectations of IP law is essential to ensure that current regulations effectively meet their goals and to identify potential changes if necessary.

Prior work looking at perceptions of IP law and authorship demonstrate that lay perceptions may be in conflict with what the law proposes. Laypeople perceive IP law's main objective as preventing plagiarism, although its main aim is more utilitarian by promoting creativity through exclusive rights [74]. In the internet, content creators also have mistaken beliefs about the copyright terms of the platforms they use [33]. **This paper investigates laypeople's perceptions of copyright law concerning works generated with the assistance of a GenAI model.**

2.6 Lay Perceptions of GenAI Outputs

Recent work has explored how people use and perceive GenAI models and their outputs. While some report that people have an inherent bias against AI-generated text and images [40], particularly among those with stronger anthropocentric beliefs in order to "protect" human creativity [78], other studies have found that laypeople prefer human-created works because AI outputs are perceived to be of lower-quality [12, 60]. At the same time, many studies demonstrate that people are bad at distinguishing between human- and AI-generated images and text [60, 71], potentially because of flawed heuristics used to determine whether something is machine-generated [50].

Research has also explored perceptions of GenAI outputs in relation to their ownership and authorship, albeit not in direct relation to their potential copyright protection. Human creators assisted by GenAI are attributed more credit than creators working alongside another human [49]. An experiment found that perceptions of authorship of the AI model and related actors can be manipulated by how anthropomorphized the machine is [31]. Similarly, perceptions of ownership over AI-generated text also vary depending on writing style [52]. Another study suggests that people may not be contrary to the idea of granting copyright rights to the AI model themselves [70]. Prior work has also discovered that laypeople believe artists whose creations are being used to train GenAI should be compensated [55].[4] **We build upon this prior work and investigate people's opinions regarding GenAI under the lens of copyright law, providing implications to the development and regulation of GenAI models.**

3 METHODS

We conducted a large-scale human-subject experiment to capture laypeople's opinions concerning the potential copyright protection of AI-generated art. Furthermore, we studied whether people exhibit egocentric biases in their opinions by looking at the differences between the perceptions of those who use AI models to create art and those who observe and evaluate AI-generated art. In this section, we describe our experimental design, including the experimental setting and the procedures we employ to gather our data. We also formulate hypotheses exploring potential egocentric biases.

3.1 Setting

To study lay perceptions of AI-generated art, we held an online AI art exhibition. We recruited laypeple to participate in the exhibition as either creators or evaluators (or both, as explained below). Participants were told that the exhibition was juried, meaning that not all images would be displayed in the exhibition, and that an online website would display the best 10 images, the creators of which would be awarded a monetary award. Our setting not only enabled us to capture lay perceptions of AI-generated art in relation to copyright law more generally, but also allowed us to form hypotheses regarding how its incentive structure may influence these perceptions.

Our choice of experimental setting was motivated by a series of considerations. Our first consideration is related to the experiment's ecological validity. We mimicked a real-world setting common in the art world: a juried art exhibition in which participants can submit their art online, and where the best rated submissions (i) receive recognition by being displayed in the exhibition and (ii) receive monetary rewards (see [90] for an equivalent scenario involving AI-generated art). The monetary rewards and recognition also serve a second purpose—incentivizing creators to put

[4]The question of copyright and authorship can also be framed as a question of responsibility. For instance, who is responsible for an AI-generated work and, thus, should have the corresponding rights? Although extensive literature has explored how laypeople attribute responsibility for harms caused by AI [59, 63, 67–69, 72], it is still an open question whether the results would replicate in the case of positive responsibility (i.e., credit for AI-generated works) [26, 87].

effort in creating images, as real-world creators of (GenAI or human-created) images normally would. Furthermore, the promise of monetary rewards and recognition to creators may also incentivize evaluators to take the task more seriously, since they are made aware of the fact that their responses influence other participants' outcomes.

It is important to note that our exhibition differs from many juried art exhibitions in one important aspect: a part of the jury consists of creators' peers who can also submit their art, and not an (unbiased) set of professionals who do not compete in the exhibition. We opted for this design choice to study egocentric effects in people's perceptions of GenAI images and their potential copyright.

3.2 Experimental Conditions

In our experiments, we recruited three groups of participants. First, we recruited *creators*, who had the chance to use a GenAI model to create an image for consideration at the AI art exhibition and then evaluated their own creation with respect to several variables. Second, we recruited *evaluators*, who were randomly assigned to one of two separate conditions: *invested evaluators* and *uninvested evaluators. Invested evaluators* used the same GenAI model to create an image for the exhibition before evaluating a subset of the submissions made by *creators. Uninvested evaluators* did not use the GenAI model and instead only evaluated *creators'* images. Figure 1 presents a high-level overview of our experimental conditions.

Before providing more details about each of the three experimental conditions to participants, we introduced the study setting, which was similar across all three treatments. On the study's landing page, we described our experiment and gathered participants' informed consent. We explained the study setting, and—depending on the experimental condition to which they were assigned—informed participants that they will be generating and/or evaluating images for an AI-generated art exhibition. After reading the task description, participants were asked comprehension check questions to ensure that they understood the task; participants could not continue with the study until they answered the questions correctly. Furthermore, we informed participants that the creators of the 10 highest-rated submissions to our juried art exhibition would have their submissions displayed in an online gallery and earn a monetary reward of 25 USD. Finally, we provided participants with the following definition of GenAI models:

> Generative AI models like the one [you will use/used by other Prolific workers] learn patterns and relationships from a dataset of human-created content (e.g., human-created images like paintings and photographs).
> When a person prompts the AI model (e.g., asks the AI model to generate an image of the sky), the AI model uses the patterns it learned from human-created content to generate the requested content (e.g., an image of the sky).

3.2.1 Creators. The experiment commenced with a three-step tutorial, in which participants were taught how to use the GenAI model to generate images. In the first step, participants were taught to write prompts. Participants were told that the AI model would generate an image according to their written instructions and were asked to instruct the AI model to generate "an image of a cat in a comic-book style of art". In the second step, participants were informed that they could ask the GenAI model to generate as many images as they wanted. We emphasized that AI model they were using (DALL-E 3 [11]) did not keep previous instructions in its memory, meaning that they had to fully describe the image they intended to generate each time. As an exercise, participants were told to instruct the AI model to generate "an image of a rabbit in an abstract style of art". Third, participants were shown how to navigate between all the images

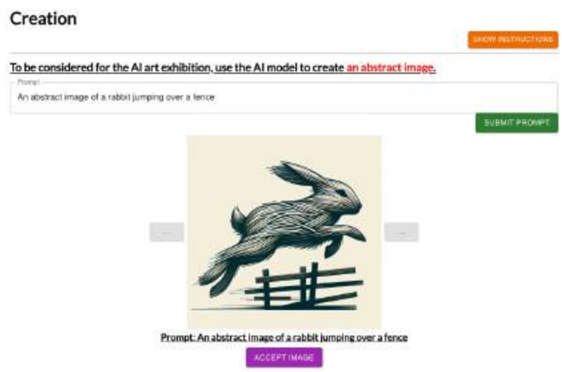

(a) Creation step.

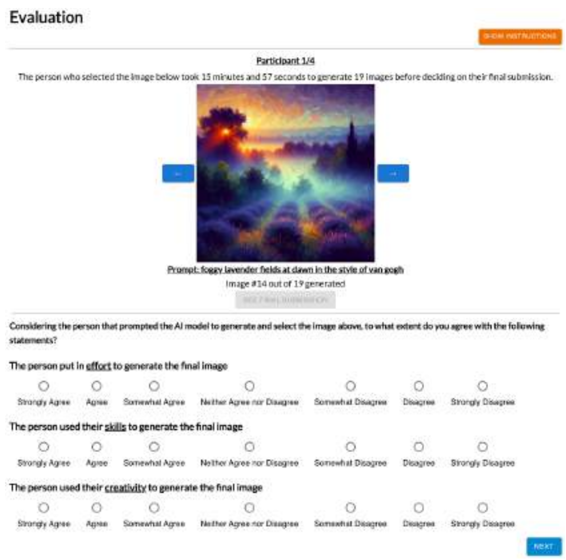

(b) Evaluation step.

Fig. 2. Screenshots of the study's user interface.

they created (with arrows located beside the image), and how to select the image they would like to submit for the consideration at the exhibition (by clicking on a button).

After completing the tutorial, creators were asked to generate their submission to our AI art exhibition (see Figure 2(a)). They were randomly assigned the task of generating *a portrait*, *an image of a landscape*, or *an abstract image* on a between-subjects basis to provide participants with some guidance and ensure some variance in our data. Participants could generate as many images as they wanted and select whichever image they preferred, independently of the order in which they were generated. After selecting the image they wanted to submit to the AI art exhibition, participants evaluated the image they submitted with respect to several variables, all of which we explain in Section 3.3.

3.2.2 Invested Evaluators. Similarly to creators, invested evaluators also completed a tutorial before generating, selecting, and evaluating their own submission to the AI art exhibition. After evaluating their own submission, they evaluated submissions made by four randomly selected creators, one at a time (see Figure 2(b)). When evaluating a creator's submission, participants had access not only to the image the creator submitted to the AI art exhibition, but also to all images generated by that creator and their respective prompts, as well as some descriptive statistics about the generation process (namely, the number of images that the creator generated before deciding on their final submission and how long they took to create and select it).

3.2.3 Uninvested Evaluators. Unlike invested evaluators, uninvested evaluators did not have the chance to complete the tutorial, nor to create and submit an image to the AI art exhibition. Instead, participants only evaluated submissions made by four randomly selected creators, one at a time. They used the same interface as invested evaluators (see Figure 2(b)).

3.3 Measures

All participants evaluated their own creations and/or images created by others. Each image was judged with respect to four groups of questions. Each of the three first question groups addressed a separate research question (see Section 1), whereas the last group helped determine which art would be included in the AI art exhibition.

(1) Factors associated with copyright decisions (RQ1):
   (a) Creativity: Creativity is one of the most important factors determining whether works are eligible for copyright. For instance, under US copyright law, a work must have at least "a modicum of creativity" [2]. Similarly, EU courts have clarified that creations must be the result of the author's "free and creative choices" [45]. Participants evaluated the creativity involved in the image generation process by agreeing with the following statement on a 7-point scale (0 = Strongly Disagree, 6 = Strongly Agree): "[I/The person] used [my/their] creativity to generate the final image."
   (b) Effort: Although some jurisdictions have rejected that mere effort warrants copyright protection (i.e., the sweat of the brow doctrine), other countries, such as Australia, suggest that effort could be sufficient [32]. Scholars have also defended that copyright decisions regarding GenAI outputs could consider the effort put by users [41]. Perceived effort was evaluated by agreeing with the following statement on the same 7-point scale: "[I/The person] put in effort to generate the final image."
   (c) Skills: There exists a legal precedent in Canada [3] that states that copyright protection requires exercise of non-trivial skills. Similarly, Australian law [1, 32] suggests that skill is sufficient for copyright. Participants evaluated skills by agreeing with the following statement on the same 7-point scale: "[I/The person] used [my/their] skills to generate the final image."

(2) Attribution of authorship (RQ2): The question of copyright is closely related to the question of authorship. One of the reasons why AI-generated art is not eligible for copyright protection in the US is that it lacks *human* authorship [28]. Hence, deciding who is the author (or authors) of a GenAI output is crucial to determining whether it is eligible for copyright protection. Although much of the discussion surrounds authorship of the user and the AI model itself, it is also possible that laypeople perceive other actors as authors, such as the artists whose creations were used to train the GenAI model and the company that developed it. Hence, we asked participants to what extent they agreed that these four entities are authors using the same 7-point agreement scale:
   (a) User: "[I am/The person who used the AI model to generate this image is] an author of this image."
   (b) AI Model: "The AI model itself is an author of this image."
   (c) Company: "The company that developed the AI model is an author of this image."
   (d) Data Contributors: "The artists whose creations were used for training the AI model are authors of this image."

(3) Attribution of rights (RQ3): Copyright law grants several exclusive rights to the holder, such as the right to distribute, reproduce, and display the work, as well as make copies and prepare derivative materials [4]. We captured participants opinions about two of these rights: 1) the right to display and 2) the right to make copies. It is important to note that it is legal to use copyrighted material under certain conditions, particularly when the work is used for non-commercial purposes (e.g., according to the US fair use doctrine). We thus collected people's opinions about the two rights in both commercial and non-commercial settings.
   Respondents indicated whom they think should have rights over the image out of a list of entities: 1) the user, 2)

the AI model, 3) the company that developed the AI model, 4) data contributors, 5) anyone, and 6) someone else (followed by an open-ended text box for indicating whom). Participants could select as many entities as they wished for each right and setting combination (e.g., right to display commercially). Entities 1-4 were described as in the authorship question presented above.

(4) Score evaluation: Participants also evaluated each image using a 11-point scale: "On a scale from 0 (Very bad) to 10 (Very good), how would you evaluate this image?" Responses to these questions determined which images were selected for the exhibition and thus received the monetary award.

First, participants answered the groups of questions about factors associated with copyright decisions (1) and authorship (2). The order of these two groups of questions was randomized, as was the order of the questions within each group. [5] Questions related to attribution of rights (3) followed, also in random order. Finally, participants scored the images to determine awards (4). For readability, some questions were rephrased depending on whether respondents were evaluating their own or other people's creations (see above for the exact phrasing). Before completing the study, participants also answered a series of exploratory questions, such as how often they use image and text GenAI models and demographic questions.

3.4 Hypotheses

One's attitudes and beliefs towards a wide range of issues can be biased egocentrically. For instance, people judge fairness violations against themselves more harshly than similar transgressions against others [39, 100]. Furthermore, individuals believe they know more about others than others know about them [88], expect mass media to have a larger influence on others than on themselves [29], and even assume that others are more susceptible to egocentric bias than they are [62].

Particularly relevant for our study, people exhibit egocentric biases in their perceptions of creative works. Authors of creative works overestimate the value of their products [13]. Similarly, people value their self-made products as much as experts' creations and expect others to share this view [82]. Hence, people may value AI-generated art they create more than they value others' art and more than others value their art. Moreover, people overestimate their own contributions to group projects [91] and take undue credit for achievements while attributing failures to external factors [77, 95]. Therefore, people may attribute higher degrees of authorship and more rights to themselves. Finally, creators overestimate the creativity of their own creations [14]—a finding that we expect to replicate in our study. Combined, this body of research motivates us to study egocentric effects on people's perceptions of copyright of AI-generated art.

We build upon this prior work and study how egocentric effects may emerge in the context of GenAI. We formed two hypotheses concerning egocentric effects on lay perceptions of AI-generated images in relation to copyright. In H1, we compare the judgments of different participants about the same image. Namely, we hypothesize that the creator of an image will judge their image more favorably than other participants. In H2, we compare the judgments of the same participant about different images. We hypothesize that participants will judge images they create more favorably than they will judge images created by others.

[5] We did not provide a definition of what creativity, effort, and skills mean in the context of our study. We did so because of the lack of clear legal definition of these factors, as well as disagreement on whether these factors matter when deciding whether a work warrants copyright (see Section 2.2). Whether a work satisfies these conditions is open to legal interpretation and argument. Hence, we decided to not bias participants' responses with respect to a particular definition of these factors and instead capture lay interpretations of these factors.

H1) Egocentric effect between participants: Creators will 1) evaluate images more favorably and 2) be more likely to identify users as authors and right-holders when judging their own creations compared to invested and uninvested evaluators judging the same images.

H2) Egocentric effect between images: Invested evaluators will 1) evaluate images more favorably and 2) be more likely to identify users as authors and right-holders when judging their own creations than when judging other people's images.

In our juried art exhibition, the highest scored submissions received rewards. Creators' and invested evaluators' outcomes depended on participants' assessments. Uninvested evaluators had no direct incentives to lie—the scores they provided had no impact on their own outcomes, since they did not take part in the competition. However, invested evaluators' outcomes depended not only on the scores their submission received, but also on the scores other submissions received.

Invested evaluators had the opportunity to inflate their own relative rating not only by giving high ratings to their own submission (in line with H2), but also by giving low scores to other people's submissions. That is, they could benefit from sabotaging their competition. The problem of sabotage in competitions has been discussed extensively in economics: theory predicts its occurrence, and experiments show that sabotage is indeed empirically relevant [20]. Therefore, we hypothesize that the incentive structure associated with our juried art exhibition may lead those with something at stake (invested evaluators) to judge others' submissions more harshly than those who do not have anything to gain from doing so (uninvested evaluators).

H3) Competition effect: Invested evaluators will 1) evaluate other people's images less favorably and 2) be less likely to identify users as authors and right-holders than uninvested evaluators.

Uninvested evaluators did not have an opportunity to interact with the GenAI model and to create AI-generated art in our experiments. On the other hand, creators and invested evaluators repeatedly interacted with the model throughout the tutorial and creation task, observing how changing their inputs influences the model's outputs and learning how to use it to create AI-generated art. Research has identified several psychological phenomena, such as the mere-exposure effect [79, 108], practice and learning effects [30], and the curse of knowledge [43], which lead us to hypothesize that participants' evaluations may be influenced by their interaction (or lack thereof) with the GenAI model. Consequently, we hypothesize that the experience of interacting with the GenAI model will influence participants' perceptions about AI-generated art.

H4) Experience effect: Participants who interacted with the GenAI model (i.e., creators and invested evaluators) will 1) evaluate images and 2) attribute authorship and rights differently than those who did not use the GenAI model (i.e., uninvested evaluator).

We emphasize that while we gather data on how participants attribute authorship and rights to various entities, our hypotheses focus on attributions of authorship and rights to one specific entity: *users*. We investigate the remaining entities exploratively (as discussed above).

3.5 Analysis Plan

We used regressions to analyze our data. We treated participants' assessments of factors, authorship, and score evaluations as continuous dependent variables in linear regressions. Participants' attributions of rights were treated as binary dependent variables and modeled using logit regressions. To account for repeated measurements across participants and images (i.e, participants evaluated several images, and images were evaluated by several participants), we initially planned to use mixed-effects linear regressions with crossed random intercepts for images and participants. However, due to convergence issues, we instead opted for regressions with two-way clustered standard errors [22].

Our primary independent variable is a dummy variable encoding both the i) treatment condition to which the participants were assigned and ii) whether the data point refers to an evaluation of their own image or someone else's creation. Hence, our dummy variables has four levels representing 1) *creators* judging their *own* images; 2) *invested evaluators* rating their *own* creations; 3) *invested evaluators* assessing *others'* images; and 4) *uninvested evaluators* judging *others'* submissions.

We tested for differences between pairs of treatments by estimating their contrasts (i.e., by estimating the difference between the treatments' estimated regression coefficients), and applied Bonferroni corrections to account for multiple comparisons. We conducted such pairwise comparisons only on pairs of treatments that are the subject of the hypotheses described above. In Section 4, we discuss all pairwise differences that were significant at the $\alpha < .05$ level.

Finally, as a robustness check, we repeated the analysis described for three different sets of covariates, by including them as additional independent variables:

(1) Participant-level variables: we included seven variables concerning the evaluator: how often they use GenAI models for generating 1) images and 2) text; whether they had already participated in a study in which they were asked to 3) generate and 4) evaluate images; and whether they have any training in professions related to 5) art, 6) computer science, or 7) law.
(2) Image-level variables: we account for the 1) number of images that creators generated before selecting one for the exhibition; 2) how long the process took; 3) the length of the selected prompt; and 4) the type of image they were asked to generate (portrait, landscape, or abstract).
(3) Order variable: we also included a variable indicating the order in which the image was shown to an evaluator. For instance, if a measurement refers to the second image that participants evaluated, this variable is equal to two. This analysis not only provides robustness to our results but also explores any order effects in participants' evaluations.

Our results are robust to all three groups of covariates. That is, pairwise comparisons between treatments remain qualitatively and quantitatively similar upon including any of the three groups of covariates, barring minor changes in the significance of borderline results. For simplicity and brevity, we do not discuss the results of models with participant- and image-level covariates in the paper. However, we discuss order effects when the coefficient is significant at the $\alpha <$ .05 level.

3.6 Data Collection & Participants

For the main study, we recruited 450 participants on Prolific [85]. First, we recruited 100 participants to complete the study as creators, followed by an additional 350 participants divided equally between invested and uninvested evaluators. We targeted US residents who were fluent in English and had completed at least 50 tasks on Prolific with an approval rate of over 95%. Participants were sampled at different hours over several days to mitigate sampling biases that may occur due to time [16].

We discarded responses from 21 participants who failed any of two instructed response questions. Due to technical problems, responses from three participants had to be dropped because they were not saved completely. Some participants took part in the study more than once, in which case we only kept their first response. Finally, we discarded judgments made by (invested and uninvested) evaluators regarding images generated by creators that were removed due to

Fig. 3. Example images generated by creators for our AI art exhibition. The images in the top row were among the best evaluated, while those in the bottom row were judged poorly.

attention check failures. Our final sample comprises 424 participants, out of which 95 were creators, 165 were invested evaluators, and 164 were uninvested evaluators. All participants were paid 3.50 USD for their participation regardless of their treatment condition to keep monetary incentives constant (approximately 11.50 USD per hour).

Half of participants identified themselves as women, 47.4% as men, and 2.12% as non-binary. Participants' mean age was 40.5 years old (*SD* = 14.0), with the youngest respondent being 20 years old and the oldest 81 years old. Our sample is slightly more diverse than the US population in terms of race, with 10.1% of participants describing themselves as Asian, 15.3% as Black or African American, and 61.6% as White. Only 7.5% of participants reported having prior training in professions related to law, while 24.1% stated that they had training in computer science-related professions and 26.2% in art-related occupations.

In addition to the participants recruited for the main study, we recruited an additional sample of evaluators for determining which art should be included in the exhibition. In the main study, only creators' submissions received evaluations from other participants. Specifically, each of the creators' submissions was evaluated by approximately 13 other participants from the pool of invested and uninvested evaluators. On the other hand, invested evaluators' submissions were only evaluated by themselves. Since both creators and invested evaluators were eligible to receive awards based on how others rate their submissions, we recruited 115 participants to rate 20 of the invested evaluators' submissions each, resulting in approximately 13 evaluations for each submission to the art show. These participants were paid 1.60 USD (approximately 11.30 USD per hour).

4 RESULTS

A total of 260 participants submitted an image to the AI art exhibition (95 creators and 165 invested evaluators). These participants generated a median of two images (mean = 4.33, *SD* = 5.08). Although 90 of these participants generated

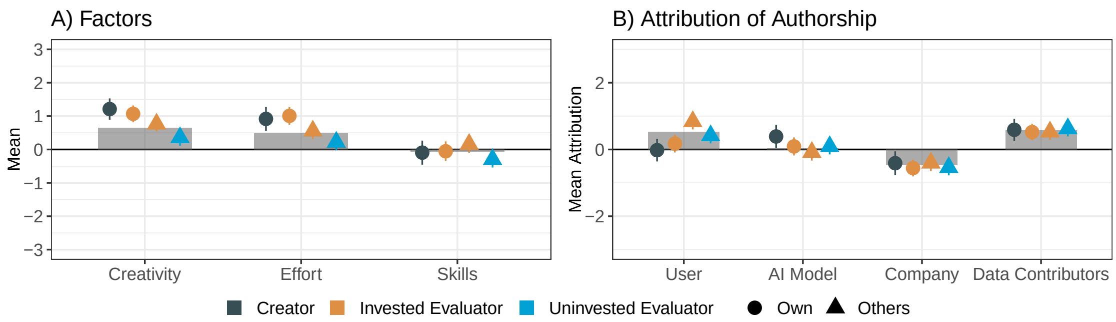

Fig. 4. A) Perceived creativity, effort, and skills involved in generating images using GenAI for our AI art exhibition. B) Perceived authorship of the user, the AI model, the company that developed the AI model, and data contributors (i.e., those whose creations were used to training the AI model). Gray bars present the mean value across all conditions, while circles and triangles represent mean values in each treatment condition according to the legend. Error bars correspond to 95% confidence intervals.

only one image, 30 of them created 10 or more. Both creators and invested evaluators generated a median of two images, with creators creating more on average (mean = 5.01, *SD* = 6.40) than invested evaluators (mean = 3.93, *SD* = 4.1). Participants took an average of 5.86 minutes (*SD* = 5.86) to create and select an image for the exhibition. Invested evaluators (mean = 5.94 minutes, *SD* = 5.66) spent slightly more time generating images than creators (mean = 5.72, *SD* = 6.23). Figure 3 presents some example images.

4.1 RQ1: Perceived Creativity, Effort, and Skills

Contrast	diff	SE	t-test	p-value
Creativity				
Creators (Own) - Uninvested Evaluators (Others) - H1 & H4	0.853	0.189	t(1515) = 4.505	**< 0.001**
Creators (Own) - Invested Evaluators (Others) - H1	0.446	0.178	t(1515) = 2.503	0.062
Invested Evaluators (Own) - Invested Evaluators (Others) - H2	0.302	0.137	t(1515) = 2.203	0.139
Uninvested Evaluators (Others) - Invested Evaluators (Others) - H3 & H4	−0.408	0.143	t(1515) = -2.844	**0.023**
Invested Evaluators (Own) - Uninvested Evaluators (Others) - H4	0.710	0.180	t(1515) = 3.948	**< 0.001**
Effort				
Creators (Own) - Uninvested Evaluators (Others) - H1 & H4	0.695	0.192	t(1515) = 3.616	**0.002**
Creators (Own) - Invested Evaluators (Others) - H1	0.361	0.206	t(1515) = 1.752	0.400
Invested Evaluators (Own) - Invested Evaluators (Others) - H2	0.451	0.145	t(1515) = 3.123	**0.009**
Uninvested Evaluators (Others) - Invested Evaluators (Others) - H3 & H4	−0.334	0.146	t(1515) = -2.284	0.112
Invested Evaluators (Own) - Uninvested Evaluators (Others) - H4	0.785	0.185	t(1515) = 4.247	**< 0.001**
Skills				
Creators (Own) - Uninvested Evaluators (Others) - H1 & H4	0.202	0.212	t(1515) = 0.953	1.000
Creators (Own) - Invested Evaluators (Others) - H1	−0.251	0.226	t(1515) = -1.111	1.000
Invested Evaluators (Own) - Invested Evaluators (Others) - H2	−0.210	0.120	t(1515) = -1.746	0.405
Uninvested Evaluators (Others) - Invested Evaluators (Others) - H3 & H4	−0.453	0.161	t(1515) = -2.803	**0.026**
Invested Evaluators (Own) - Uninvested Evaluators (Others) - H4	0.242	0.195	t(1515) = 1.241	1.000

Table 1. Pairwise comparisons of perceived creativity, effort, and skills between treatment conditions. We only test the contrasts relevant to our hypotheses presented in Section 3.4.

Figure 4A shows participants' mean evaluations regarding the creativity, effort, and skills involved in generating images with GenAI. Participants somewhat agreed that creativity (mean = 0.656, SD = 1.72) and effort (mean = 0.488, SD = 1.79) were necessary to create the images. In contrast, judgments concerning skills were closer zero (mean = -0.070, SD = 1.82), meaning that on average participants neither agreed nor disagreed that creators used their skills to generate images with GenAI. Table 1 presents pairwise comparisons of perceived creativity, effort, and skills between treatment conditions. Below, we discuss results of pairwise comparisons for each factor separately.

4.1.1 Creativity. Concerning creativity, uninvested evaluators rated images lower than creators (diff = 0.853) and invested evaluators, both when the latter evaluated their own creations (diff = 0.710) and other people's images (diff = 0.408). That is, our results regarding creativity support H4, and they partially support H1. Finally, we found an effect in the opposite direction than the one hypothesized in H3—invested evaluators assigned higher creativity scores to creators than uninvested evaluators did.

4.1.2 Effort. Our analysis shows that uninvested evaluators assign lower effort ratings than creators (diff = 0.695) and invested evaluators judging their own images (diff = 0.785), offering partial support to H1 and H4. We also find that invested evaluators assign higher effort ratings to their own creations than to others' (diff = 0.451), in line with H2.

4.1.3 Skills. The only significant difference in judgments of skills was found between uninvested and invested evaluators when the latter evaluated other people's images (diff = −0.453). Invested evaluators attributed more skills to creators than their uninvested counterparts. That is, as for creativity, we find an effect in the opposite direction than the one hypothesized in H3.

4.2 RQ2: Perceived Authorship

Figure 4B shows how participants attributed authorship between the user, the AI model, the company that developed the AI model, and data contributors. Participants somewhat agreed that users (mean = 0.537, SD = 1.69) and data contributors (mean = 0.571, SD = 1.69) are authors of AI-generated images. Judgments concerning the AI model were more uncertain (mean = 0.037, SD = 1.78), with evaluators on average neither agreeing nor disagreeing that the AI model itself is an author. The company that developed the AI model had the lowest perceived authorship (mean = -0.476, SD = 1.71). Table 2 presents pairwise comparisons of perceived authorship between treatment conditions. Below, we discuss results for each entity separately.

4.2.1 User. When judging others' submissions, invested evaluators attributed more authorship to creators than creators attributed to themselves (diff = −0.859). This effect is the opposite of the one hypothesized in H1. Moreover, invested evaluators attribute more authorship to others than to themselves (diff = −0.662). This effect goes against our hypothesis H2.

We also observed some order effects in participants' attribution of authorship to the user. The more images a participant evaluated, the more authorship they attributed to users ($b = 0.050$, $SE = 0.021$, $t(1514) = 2.384$, $p < .05$). We note that the results described above are robust to including the order variable as a covariate. Invested evaluators attributed more authorship to creators (i.e., users) than creators themselves (diff = −0.734). Furthermore, invested evaluators attributed more authorship to others than to themselves (diff = −0.537).

4.2.2 Other Entities. There were no significant differences in the perceived authorship of the AI model, the developer, and data contributors across treatments. Nonetheless, we note a borderline significant order effect on perceived

Contrast	diff	SE	t-test	p-value
User				
Creators (Own) - Uninvested Evaluators (Others) - H1 & H4	−0.437	0.205	t(1515) = -2.126	0.168
Creators (Own) - Invested Evaluators (Others) - H1	−0.859	0.208	t(1515) = -4.123	**< 0.001**
Invested Evaluators (Own) - Invested Evaluators (Others) - H2	−0.662	0.104	t(1515) = -6.386	**< 0.001**
Uninvested Evaluators (Others) - Invested Evaluators (Others) - H3 & H4	−0.422	0.168	t(1515) = -2.509	0.061
Invested Evaluators (Own) - Uninvested Evaluators (Others) - H4	−0.240	0.181	t(1515) = -1.326	0.925
AI Model				
Creators (Own) - Uninvested Evaluators (Others) - H1 & H4	0.301	0.222	t(1515) = 1.355	0.878
Creators (Own) - Invested Evaluators (Others) - H1	0.472	0.220	t(1515) = 2.142	0.162
Invested Evaluators (Own) - Invested Evaluators (Others) - H2	0.174	0.098	t(1515) = 1.762	0.391
Uninvested Evaluators (Others) - Invested Evaluators (Others) - H3 & H4	0.172	0.179	t(1515) = 0.96	1.000
Invested Evaluators (Own) - Uninvested Evaluators (Others) - H4	0.002	0.183	t(1515) = 0.011	1.000
Company				
Creators (Own) - Uninvested Evaluators (Others) - H1 & H4	0.126	0.213	t(1515) = 0.592	1.000
Creators (Own) - Invested Evaluators (Others) - H1	−0.007	0.219	t(1515) = -0.031	1.000
Invested Evaluators (Own) - Invested Evaluators (Others) - H2	−0.154	0.091	t(1515) = -1.695	0.451
Uninvested Evaluators (Others) - Invested Evaluators (Others) - H3 & H4	−0.133	0.178	t(1515) = -0.747	1.000
Invested Evaluators (Own) - Uninvested Evaluators (Others) - H4	−0.021	0.177	t(1515) = -0.119	1.000
Data Contributors				
Creators (Own) - Uninvested Evaluators (Others) - H1 & H4	−0.034	0.211	t(1515) = -0.162	1.000
Creators (Own) - Invested Evaluators (Others) - H1	0.062	0.204	t(1515) = 0.302	1.000
Invested Evaluators (Own) - Invested Evaluators (Others) - H2	−0.007	0.098	t(1515) = -0.067	1.000
Uninvested Evaluators (Others) - Invested Evaluators (Others) - H3 & H4	0.096	0.172	t(1515) = 0.557	1.000
Invested Evaluators (Own) - Uninvested Evaluators (Others) - H4	−0.103	0.175	t(1515) = -0.587	1.000

Table 2. Pairwise comparisons of perceived authorship between treatment conditions. We only test the contrasts relevant to our hypotheses presented in Section 3.4.

authorship of the AI model, such that the more images participants evaluated, the less authorship they attributed to the AI model (b = -0.056, SE = 0.030, $t(1514) = -1.874$, $p = 0.061$).

4.3 RQ3: Attribution of Rights

Figure 5 presents participants' opinions regarding who should hold the rights to display and make copies of AI-generated images. Users (i.e., those who used the GenAI to generate images) were selected by more than 60% of participants for all rights, in both commercial and non-commercial settings. Data contributors were identified as rights-holders by approximately 50% of participants across all rights and settings. The company that developed the AI model was granted the rights to display and make copies of AI-generated images by around 41% of participants in non-commercial settings; however, only 32% of respondents believed the company should have commercial rights over the image. The AI model was recognized as a rights-holder by around 19% and 27% of participants for commercial and non-commercial uses, respectively.

Participants were also able to indicate that no one should have *exclusive* rights over the images by selecting that "anyone" should be able to display and make copies of AI-generated images. We observed clear differences in responses between commercial and non-commercial rights. Participants were more likely to support non-exclusive rights in non-commercial settings (approximately 57% and 47% in favor of anyone having the right to display and make copies,

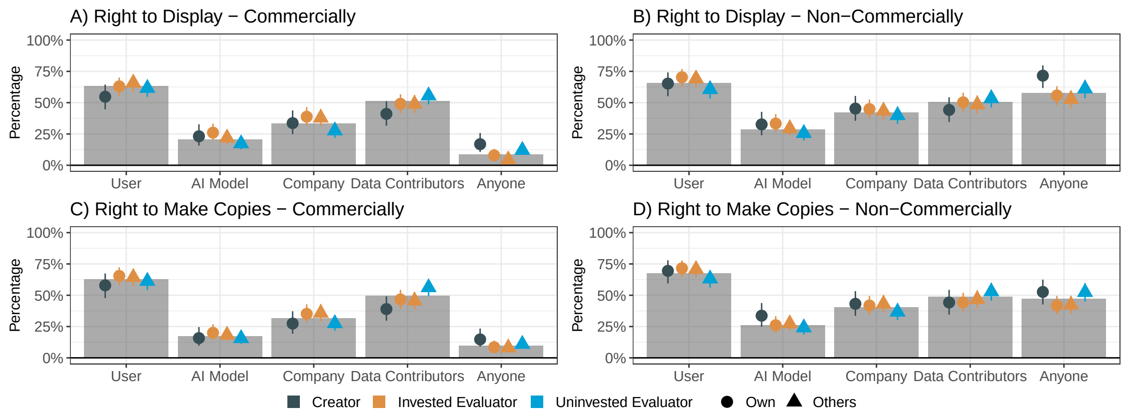

Fig. 5. Percentage of participants who chose the user, the AI model, the company that developed the AI model, and data contributors (i.e., those whose creations were used to training the AI model) as right-holders. We present results separately for the rights to A) display commercially, B) to display non-commercially, C) to make copies commercially, and D) to make copies non commercially. Gray bars present the mean percentage across all conditions, while circles and triangles represent mean percentages in each treatment condition according to the legend. Error bars correspond to 95% confidence intervals.

respectively). In contrast, when evaluating commercial rights, only a few respondents indicated that anyone should have them (approximately 9%).

Unlike for attributions of authorship, our treatments had little effect on attributions of rights. That is, the patterns described above are fairly consistent across all treatments. Hence, for brevity, we omit the pairwise comparisons tables from the paper and comment on all of the differences identified as statistically significant directly in the text below.

We did not find support for any of our hypotheses regarding attribution of rights to users. There were only a few significant pairwise differences between conditions. Creators were relatively more likely to support non-exclusive rights to display their own creations than invested evaluators evaluating the same images (OR = 4.347, z = 3.532, $p < .005$ for commercial use, OR = 2.253, z = 3.043, $p < .05$ for non-commercial use). When evaluating other people's creations, invested evaluators were less likely to support non-exclusive rights to display images commercially than uninvested evaluators (OR = 0.340, z = -2.801, $p < .05$). Finally, creators were borderline more supportive of granting the right to make copies commercially to data contributors than uninvested evaluators (OR = 0.497, z = -2.622, $p < 0.05$). No other differences were statistically significant.

4.4 RQ4: Score Evaluation

Figure 6 presents participants' mean score evaluations on a 11-point scale (0 = Very bad, 10 = Very good). These scores were used to determine who was granted the monetary award and which images were displayed in our AI art exhibition. On average, participants evaluated images slightly positively (mean = 6.61, *SD* = 2.38). Table 3 presents pairwise comparisons between treatment conditions.

Pairwise comparisons between treatments indicate clear egocentric effects. In line with H1, creators evaluated their own creations more highly than invested and uninvested evaluators judging the same images (diff = 1.787 and diff = 1.966, respectively). As hypothesized by H2, invested evaluators judged their own images more positively than those

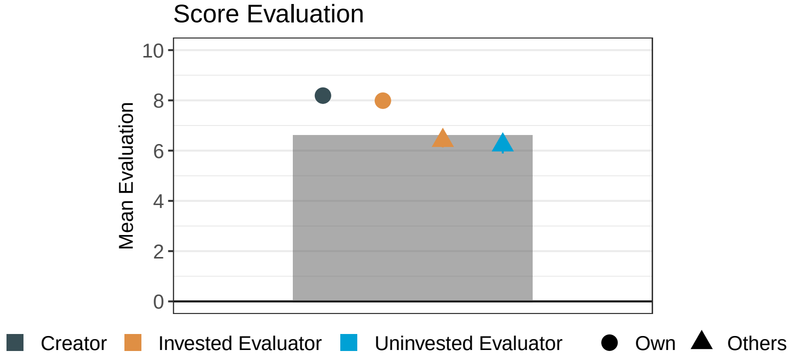

Fig. 6. Score evaluations of images generated by creators. The gray bar presents the mean score across all conditions, while circles and triangles represent mean scores in each treatment condition according to the legend. Error bars correspond to 95% confidence intervals.

Contrast	diff	SE	t-test	p-value
Creators (Own) - Uninvested Evaluators (Others) - H1 and H4	1.966	0.226	t(1515) = 8.686	**< 0.001**
Creators (Own) - Invested Evaluators (Others) - H1	1.787	0.212	t(1515) = 8.425	**< 0.001**
Invested Evaluators (Own) - Invested Evaluators (Others) - H2	1.586	0.192	t(1515) = 8.264	**< 0.001**
Uninvested Evaluators (Others) - Invested Evaluators (Others) - H3 and H4	−0.178	0.193	t(1515) = -0.924	1.000
Invested Evaluators (Own) - Uninvested Evaluators (Others) - H4	1.764	0.226	t(1515) = 7.807	**< 0.001**

Table 3. Pairwise comparisons of image score evaluations between treatment conditions. We only test the contrasts relevant to our hypotheses presented in Section 3.4.

created by other participants (diff = 1.586). We found that uninvested evaluators score images generated by creators lower than invested evaluators score their own creations (diff = 1.764), as suggested by H4.

5 DISCUSSION

5.1 Laypeople's Perceptions of AI-Generated Art in Relation to Copyright

5.1.1 RQ1: Creativity and Effort—Not Skills—Are Required To Generate Art with AI. Participants believed that creativity and effort were necessary to generate images with the GenAI model. In contrast, they neither agreed nor disagreed that creators used their skills when creating AI art.

The fact that the AI generation process was described as creative seems to suggest that AI art could have the modicum of creativity that US copyright law requires for a work to warrant copyright [84], which is at odds with current decisions to reject protection for AI-generated images [28]. Concerning judgments of effort, we note that our study provided evaluators with information that can work as a proxy for effort (e.g., the total number of images creators generated and how long they spent in the study), which could have influenced how participants judged the effort put by creators. Nonetheless, our experimental findings are robust to the inclusion of these covariates. Finally, perceptions of skills could have been limited by how much control creators had over the images they generated. Had creators been allowed to edit the GenAI outputs, evaluators could have judged the creation process as requiring more skills—a research direction that future work could explore.

5.1.2 RQ2: Users and Data Contributors as Authors. Users and data contributors were attributed the most authorship for AI-generated art. Although it is less surprising that those who prompted the GenAI model to generate images are perceived as authors, participants' acknowledgement of data contributors is noteworthy. In line with prior work suggesting that laypeople agree that data contributors should be compensated if their work is used for training GenAI models [55], our results suggest that current practices that fail to compensate and appreciate data contributors, which some would equate to theft [18, 76], may be at odds with laypeople's expectations. People's opinions seem to be more aligned with proposals of licensing models, under which data contributors can be compensated for the use of their creations [15].

Our study explained to participants that GenAI models learn from a "dataset of human-created content," which could have made the role of data contributors more prominent in our study, potentially influencing how people perceived their role in GenAI. Nonetheless, it is important to emphasize that GenAI models require such datasets to work, meaning that any explanation of how GenAI works without mentioning training data is incomplete. We call for future work exploring how different ways of introducing data contributors impact how much authorship and credit laypeople grant to those whose creations are used to train AI.

People's attribution of authorship to training data contributors also highlights the importance of CS research for copyright law in the age of GenAI. To adequately compensate training data contributors for the impact their work had on a generated image, it is important to be able to identify which data points contributed to the generation of the image, as well as estimate their own individual contribution. This is closely related to the problem of data valuation in the machine learning literature [38, 51, 86, 89]. These estimating, however, have been found to be intractable in certain circumstances [42]. Hence, we call for further research on how to robustly quantify training data influence, which can provide technical solutions for compensating training data contributors appropriately.

Our result that users were acknowledged as authors of the images they generate using AI calls for the reconsideration of legal decisions that have refused to grant copyright to GenAI ouputs in the US [28]. The main reason given for the rejection was that AI-generated images do not have human authors, making it ineligible for copyright protection. In contrast, our findings suggest that laypeople believe users are authors, even in a scenario in which they were not allowed to edit the AI-generated images. We expect the perceived authorship of users to increase when given the chance to further exert control over GenAI outputs.

Participants were uncertain about the authorship of the AI model, which is aligned with prior work suggesting that laypeople are neither against nor in favor of copyright rights to AI models [70]. Finally, respondents' disagreement with the idea of companies being authors of GenAI outputs are at odds with potential interpretations of UK IP law, which could be interpreted as the companies training GenAI models holding copyright over outputs [48]. That is, current practices that concentrate profits in the hands of the corporations training GenAI models without little to no compensation to data contributors [18, 76] do not align with lay opinions about copyright of AI-generated art.

We note, however, these these two actors, to whom participants attributed lower levels of authorship, are non-human entities. In contrast, the user and data contributors are human, suggesting that our findings could have been influenced by the human nature of these actors. Future work could replicate our study by also investigating whether the (human) programmers that developed the model would be granted authorship, which could indicate that the human nature of potential authors plays a role in people's perceptions of authorship.

5.1.3 RQ3: Users and Data Contributors as Rights-Holders. Similar to our results concerning authorship, those who prompted the AI model to generate images (i.e., users) and data contributors were frequently identified as potential rights-holders of AI-generated art. This finding has two main implications. First, it puts into question the exclusive nature of the rights afforded by copyright protection; participants seemed to support a more distributed approach, under which not only "authors" would be owners of the copyright, but also some actors without which AI-generated art would not be possible. For instance, a joint ownership model could be more aligned with laypeople's expectations regarding AI-generated art. Second, it highlights participants' calls for the compensation of data contributors. An approach that could be more aligned with lay opinions could rely on neighboring rights [47, 63], which would still center ownership around "authors" (i.e., users) but without neglecting the interests of data contributors, e.g., through the distribution of royalties. As discussed above, another solution could be based on licensing models for training data [15].

Participants were more supportive of non-exclusive rights over AI-generated art in non-commercial settings, in line with doctrines that permit the use of copyrighted material for non-commercial purposes (e.g., the US fair use doctrine). Future debates concerning the use of AI-generated outputs in relation to copyright could consider our results to determine whether their use should be legal.

5.2 Egocentric Biases (Or Lack Thereof) in Perceptions of AI-Generated Images

5.2.1 Creativity, Effort and Skills. Judgments regarding creativity and effort partially exhibit egocentric effects. Creators attributed more creativity to their creations than uninvested evaluators (partial H1), in line with prior work [14]. Evaluations of the effort put in generating images with GenAI are also consistent with our egocentric hypotheses (partial H1 and H2). Furthermore, judgments concerning creativity are aligned with H4 (i.e., experience effect), which suggests that people may judge images more favorably after interacting with the GenAI model. Judgments about effort and skills are also partially aligned with H4.

Surprisingly, our results concerning perceived creativity and skills suggest the opposite of the hypothesized competition effect (H3). Even though invested evaluators were judging images against which they were competing in the exhibition, they assigned higher creativity and skill ratings than uninvested evaluators (i.e., those not competing for the monetary award).

With the current experimental design, it is difficult to distinguish between egocentric and experience effects since all participants who had reasons to judge images egocentrically also interacted with the model. Future work could disentangle these two effects. For instance, some participants could enroll in the competition with an image that they did not generate (e.g., with images generated by the researchers). This treatment condition would maintain potential egocentric effects, since participants would still be incentivized to favor their own submissions to the exhibition, without interacting with the GenAI model, mitigating experience effects.

Another approach that could help research understand how people perceive AI-generated art in relation to copyright is exploring whether the "experience effect" (H4) emerges from interacting with the model or competing in our AI art exhibition. In our study, all participants who interacted with the model also submitted an image for consideration at the exhibition. Future work could distinguish between these two factors by having a condition in which creators use the GenAI model to generate images without the competition setting, thus capturing only the effect of using the model on people's perceptions.

5.2.2 Perceived Authorship. We hypothesized that participants would exhibit egocentric biases in their attribution of authorship to users. However, our findings are at odds with our hypotheses. On average, invested evaluators attributed more authorship to other people compared to themselves, which is contrary to H2. When evaluating other people's images, invested evaluators also credited users as authors more than the users themselves (i.e., creators), a result that is also misaligned with H1.

A potential explanation to why creators attributed less authorship to themselves than to other people is that authorship does not only encode positive rights but also negative responsibilities from which creators tried to distance themselves [34, 80]. Although authorship is often associated with positive outcomes in the context of copyright law (e.g., the opportunity to profit from one's creation), it could also lead to the responsibility for any negative outcomes or backlash arising from their work being published. If one is an author under this interpretation, they are also responsible, for instance, if their creation is found to be infringing on another copyrighted work, which is relevant for debates concerning AI-generated images. Hence, creators seem to want to enjoy the rights associated with copyright (see above), but might rather not be perceived as authors because of potential negative repercussions in case something goes wrong. Future work could explore this potential explanation further.

In H4, we hypothesized that having experience with the AI model would influence participants' attribution of authorship to the user. Informally, we expected that by interacting with the AI model, people would notice the degree of influence a user's prompts have over the model's outputs, leading them to attribute more authorship to users. We did not, however, find a significant difference between uninvested evaluators and other groups of participants for perceptions of user authorship. It is important to note that H4 focused on the experience of observing the relationship between one's own inputs and the model's outputs. Perhaps there was not much variance in their own inputs, or they generated too few images (the average participant prompted the model a median of two times) for the experience effects to kick in.

However, we identified a different effect of experience. The more images a participant evaluated, the more authorship they attributed to users, suggesting order effects. The identified order effects relate to a different form of experience: observing the relationship between other participants' inputs and model outputs. It is likely that there was more between-participants variance than within-participant variations in the images they generated.[6] Since evaluators observed the prompts from four different creators, they had more opportunities for the effect of experience to come into play.

A potential explanation connecting our findings regarding H1, H2, H4 and order effects is that being exposed to diverse GenAI inputs and outputs makes evaluators value the human component of AI-generated images more. After using the GenAI model and seeing what other participants generated and realizing that different human actions led to widely different outputs, invested evaluators may have started acknowledging that AI cannot generate images without some level of human input and authorship. Creators and invested evaluators judged their own images without such experience, while invested evaluators judged the creators' images after this experience, hence possibly explaining why the latter received higher authorship scores (opposite of H1 and H2). To explore this potential explanation, future work could explore inverting the order in which invested evaluators generate and evaluate images. By first evaluating other people's submissions and then generating their own image, participants might end up overestimating their authorship after seeing what others were able to create, providing further evidence of order effects.

Finally, we note that the increase in the perceived authorship of the user was accompanied with a slight decrease in AI model's authorship. A potential factor influencing this shift in authorship could be the perceived skills of the user, which exhibit a similar pattern across our experimental conditions. It is possible that seeing how different participants

[6]This would be aligned with prior research in social psychology, which finds higher degrees of variance in the judgments of different people than between the judgments of the same person in different points in time [53, 54]

generated images of varying quality—while using the *same* GenAI model—highlighted the skills behind prompting, leading to higher perceived authorship to users and a slightly lower perceived authorship of the AI model.

5.2.3 Attribution of Rights. Finally, we also identified some differences between treatments regarding the right to display AI-generated images: invested evaluators were less likely to support non-exclusive rights for creators' images than uninvested evaluators and the creators themselves. A potential explanation is that creators would like audiences to see their creations because it could lead to further financial benefits in the long run. Instead of keeping their creations "behind curtains," creators could have imagined that allowing others to display their images could bring financial benefits down the road. Furthermore, invested evaluators might have wanted to "hide" their competitors' creations because it competes with their own submissions, while uninvested evaluators had no incentive to do so. Additional longitudinal studies are necessary to explore this hypothesis further.

5.2.4 Actual Rewards vs. Hypothetical Rights. Participants evaluated their own creations egocentrically when this assessment was supposed to determine who received monetary awards. On average, creators rated their images over 1.8 point higher (on a 11-point scale) than evaluators evaluated the same images. Furthermore, invested evaluators evaluated their own images around 1.5 point higher than they evaluated creators' images. Our results suggest that financial incentives can bias people's responses to AI-generated art. However, when asked who should hypothetically have some of the rights associated with copyright protection, participants did not prioritize their own gains. Instead, egocentric biases emerged only in the context of actual rewards, and not hypothetical rights, suggesting that future debates concerning AI-generated works in relation to copyright law may be tainted by conflicts of interest.

Participants' tendency to rate their creations egocentrically when financial incentives are present but not when asked concerning hypothetical rights is also relevant in relation to data contributors. We found that people acknowledged data contributors as authors and potential rights-holders; yet, it is unclear whether we would observe the same result had the corresponding questions involved actual monetary awards. It is possible that if participants were asked to *financially* compensate data contributors, e.g., by sharing their award, people would not do it and rather keep the incentive to themselves. Such a finding would suggest that laypeople may support regulatory frameworks that compensate data contributors only if these frameworks do not affect them financially. Our study design could be easily modified to address these research questions. For instance, future studies could include questions concerning participants' willingness to share their potential reward with data contributors.

It is worth noting that while creators and invested evaluators assigned higher ratings to their own images, we found no evidence of the sabotage hypothesized in H3. That is, invested evaluators did not give creators significantly lower scores than uninvested evaluators. It is unclear whether this lack of sabotage is the result of a lack of understanding of the incentive structure, or if it reflects participants true altruistic behavior. Future research is needed to test participants' understanding of the incentive structure. Another potential explanation is that participants chose to inflate their own scores instead of lowering other participants' scores because doing so could be seen less morally contentious—a hypothesis that future work could explore further.

Our finding that granting monetary awards has the potential to make people judge AI-generated images egocentrically has implications for copyright law. Copyright law relies on similar financial incentives to fulfill its normative goal of promoting creativity. Extending copyright to a work of authorship restricts who can benefit and profit from it—i.e., decides who receives the financial incentives associated with copyright—thus making potential rights-holders susceptible to egocentric biases.

5.3 Concluding Remarks

In summary, we found that users of GenAI judge their own creations egocentrically concerning some factors, but not others. In fact, our results even suggest effects in the opposite direction for some of the variables we examined (e.g., perceived authorship of the user). Importantly, we identified the importance of financial interests in determining whether people favor their own AI-generated images compared to other people's creations.

This research focused on how laypeople perceive AI-generated images in relation to copyright. However, copyright law regulates other works of authorship that could challenge the law if generated with GenAI. For instance, the legal status of AI-generated songs and novels have also fueled contentious debates [92, 93]. We explored images because of the wide availability of AI image models, and well as its potential lower cognitive load for participants compared to other works of authorship, such as novels, which would take longer to evaluate. We call for future work investigating how laypeople's expectations of copyright for AI-generated works may vary depending on the content form.

All the factors we explore in this study impact whether something is eligible for copyright protection. In other words, a comprehensive analysis of all these factors is what determines whether a work of authorship, AI-generated or otherwise, is protected under copyright law. Our findings that some of these factors are susceptible to egocentric biases—while others are not—raise the question of whether current methods of determining copyright eligibility are appropriate for GenAI. Future work could explore whether some of these factors are indeed appropriate in the context of GenAI, as well as examine whether other variables that are currently irrelevant in determinations regarding human-created works should be considered when examining AI-generated outputs.

REFERENCES

[1] 1964. Ladbroke (Football) Ltd v. William Hill (Football) Ltd. , 273 pages.

[2] 1991. Feist Publications, Inc. v. Rural Telephone Service Co. , 340 pages.

[3] 2004. CCH Canadian Ltd. v. Law Society of Upper Canada. , 13 pages.

[4] 2022. Copyright Law of the United States (Title 17). https://www.copyright.gov/title17/.

[5] Ryan Abbott. 2020. *The reasonable robot: artificial intelligence and the law.* Cambridge University Press.

[6] Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774* (2023).

[7] Sonja Anderson. 2014. How a Real Photo of a Flamingo Snuck Into—and Won—an A.I. Art Competition. Smithsonian Magazine. https://tinyurl.com/bdf27n3d.

[8] Isabelle Augenstein, Timothy Baldwin, Meeyoung Cha, Tanmoy Chakraborty, Giovanni Luca Ciampaglia, David Corney, Renee DiResta, Emilio Ferrara, Scott Hale, Alon Halevy, et al. 2023. Factuality challenges in the era of large language models. *arXiv preprint arXiv:2310.05189* (2023).

[9] Edmond Awad, Sohan Dsouza, Jean-François Bonnefon, Azim Shariff, and Iyad Rahwan. 2020. Crowdsourcing moral machines. *Commun. ACM* 63, 3 (2020), 48–55.

[10] Stephanie Plamondon Bair. 2017. Rational Faith: The Utility of Fairness in Copyright. *BUL Rev.* 97 (2017), 1487.

[11] James Betker, Gabriel Goh, Li Jing, Tim Brooks, Jianfeng Wang, Linjie Li, Long Ouyang, Juntang Zhuang, Joyce Lee, Yufei Guo, et al. 2023. Improving image generation with better captions. *Computer Science. https://cdn.openai.com/papers/dall-e-3. pdf* 2, 3 (2023), 8.

[12] Alexander H Bower and Mark Steyvers. 2021. Perceptions of AI engaging in human expression. *Scientific reports* 11, 1 (2021), 21181.

[13] Christopher Buccafusco and Christopher Sprigman. 2010. Valuing intellectual property: an experiment. *Cornell L. Rev.* 96 (2010), 1.

[14] Christopher Buccafusco and Christopher Jon Sprigman. 2011. The creativity effect. *U. Chi. L. Rev.* 78 (2011), 31.

[15] Joy R. Butler. 2024. How Licensing Models Can Be Used for AI Training Data. Variety. https://tinyurl.com/f3juf6dc.

[16] Logan S Casey, Jesse Chandler, Adam Seth Levine, Andrew Proctor, and Dara Z Strolovitch. 2017. Intertemporal differences among MTurk workers: Time-based sample variations and implications for online data collection. *Sage Open* 7, 2 (2017), 2158244017712774.

[17] Alan Chan, Herbie Bradley, and Nitarshan Rajkumar. 2023. Reclaiming the digital commons: A public data trust for training data. In *Proceedings of the 2023 AAAI/ACM Conference on AI, Ethics, and Society.* 855–868.

[18] Kyle Chayka. 2023. Is A.I. Art Stealing from Artist. The New Yorker. https://tinyurl.com/bdnhz8wr.

[19] E Choi. 2023. Protecting visual artists from generative ai: An interdisciplinary perspective. In *1st ICML Workshop on Generative AI and Law.*

[20] Subhasish M Chowdhury and Oliver Gürtler. 2015. Sabotage in contests: a survey. *Public Choice* 164 (2015), 135–155.

[21] Mark Coeckelbergh. 2020. Artificial intelligence, responsibility attribution, and a relational justification of explainability. *Science and engineering ethics* 26, 4 (2020), 2051–2068.

[22] Sergio Correia et al. 2016. A feasible estimator for linear models with multi-way fixed effects. *Preprint at http://scorreia. com/research/hdfe. pdf* (2016).

[23] Thomas H Costello, Gordon Pennycook, and David G Rand. [n. d.]. Durably reducing conspiracy beliefs through dialogues with AI. ([n. d.]).

[24] Rebecca Currey and Jane Owen. 2021. In the Courts: Australian Court finds AI systems can be "inventors'. WIPO Magazine. https://tinyurl.com/bdzzu6kn.

[25] Edward J Damich. 1988. The Right of Personality: A Common–Law Basis for the Protection of the Moral Rights of Authors. *Ga. L. Rev.* 23 (1988), 1.

[26] John Danaher and Sven Nyholm. 2021. Automation, work and the achievement gap. *AI and Ethics* 1, 3 (2021), 227–237.

[27] Emilia David. 2024. US patent office confirms AI can't hold patents. TheVerge. https://tinyurl.com/93x3yc7m.

[28] Wes Davis. 2023. AI-generated art cannot be copyrighted, rules a US federal judge. The Verge. https://tinyurl.com/26pepkva.

[29] W Phillips Davison. 1983. The third-person effect in communication. *Public opinion quarterly* 47, 1 (1983), 1–15.

[30] John J Donovan and David J Radosevich. 1999. A meta-analytic review of the distribution of practice effect: Now you see it, now you don't. *Journal of Applied Psychology* 84, 5 (1999), 795.

[31] Ziv Epstein, Sydney Levine, David G Rand, and Iyad Rahwan. 2020. Who gets credit for AI-generated art? *Iscience* 23, 9 (2020).

[32] Jason K Eshraghian. 2020. Human ownership of artificial creativity. *Nature Machine Intelligence* 2, 3 (2020), 157–160.

[33] Casey Fiesler, Cliff Lampe, and Amy S Bruckman. 2016. Reality and perception of copyright terms of service for online content creation. In *Proceedings of the 19th ACM conference on computer-supported cooperative work & social computing*. 1450–1461.

[34] Michel Foucault. 2019. *Aesthetics, method, and epistemology: Essential works of Foucault 1954-1984*. Penguin UK.

[35] Giorgio Franceschelli and Mirco Musolesi. 2022. Copyright in generative deep learning. *Data & Policy* 4 (2022), e17.

[36] Paul Edward Geller. 2000. Copyright history and the future: what's culture got to do with it. *J. Copyright Soc'y USA* 47 (2000), 209.

[37] Daniel J Gervais. 2019. The machine as author. *Iowa L. Rev.* 105 (2019), 2053.

[38] Amirata Ghorbani and James Zou. 2019. Data shapley: Equitable valuation of data for machine learning. In *International conference on machine learning*. PMLR, 2242–2251.

[39] Jerald Greenberg. 1983. Overcoming egocentric bias in perceived fairness through self-awareness. *Social Psychology Quarterly* (1983), 152–156.

[40] Li Gu and Yong Li. 2022. Who made the paintings: Artists or artificial intelligence? The effects of identity on liking and purchase intention. *Frontiers in Psychology* 13 (2022), 941163.

[41] Andres Guadamuz. 2017. Do androids dream of electric copyright? Comparative analysis of originality in artificial intelligence generated works. *Intellectual property quarterly* (2017).

[42] Zayd Hammoudeh and Daniel Lowd. 2024. Training data influence analysis and estimation: A survey. *Machine Learning* (2024), 1–53.

[43] Pamela J Hinds. 1999. The curse of expertise: The effects of expertise and debiasing methods on prediction of novice performance. *Journal of experimental psychology: applied* 5, 2 (1999), 205.

[44] Kalin Hristov. 2016. Artificial intelligence and the copyright dilemma. *Idea* 57 (2016), 431.

[45] P Bernt Hugenholtz and João Pedro Quintais. 2021. Copyright and artificial creation: does EU copyright law protect AI-assisted output? *IIC-International Review of Intellectual Property and Competition Law* 52, 9 (2021), 1190–1216.

[46] Justin Hughes. 2011. The photographer's copyright-Photograph as art, photograph as database. *Harv. JL & Tech.* 25 (2011), 339.

[47] Vincenzo Iaia. 2022. To Be, or Not to Be… Original Under Copyright Law, That Is (One of) the Main Questions Concerning AI-Produced Works. *GRUR International* 71, 9 (2022), 793–812.

[48] Jani Ihalainen. 2018. Computer creativity: Artificial intelligence and copyright. *Journal of Intellectual Property Law & Practice* 13, 9 (2018), 724–728.

[49] Arthur S Jago and Glenn R Carroll. 2023. Who made this? Algorithms and authorship credit. *Personality and Social Psychology Bulletin* (2023), 01461672221149815.

[50] Maurice Jakesch, Jeffrey T Hancock, and Mor Naaman. 2023. Human heuristics for AI-generated language are flawed. *Proceedings of the National Academy of Sciences* 120, 11 (2023), e2208839120.

[51] Ruoxi Jia, David Dao, Boxin Wang, Frances Ann Hubis, Nick Hynes, Nezihe Merve Gürel, Bo Li, Ce Zhang, Dawn Song, and Costas J Spanos. 2019. Towards efficient data valuation based on the shapley value. In *The 22nd International Conference on Artificial Intelligence and Statistics*. PMLR, 1167–1176.

[52] Kowe Kadoma, Marianne Aubin Le Quere, Jenny Fu, Christin Munsch, Danae Metaxa, and Mor Naaman. 2023. The Role of Inclusion, Control, and Ownership in Workplace AI-Mediated Communication. *arXiv preprint arXiv:2309.11599* (2023).

[53] Daniel Kahneman, Andrew M Rosenfield, Linnea Gandhi, and Tom Blaser. 2016. Noise: How to Overcome the High, Hidden Cost of Inconsistent Decision Making. *Harvard Business Review* (2016).

[54] Daniel Kahneman, Olivier Sibony, and Cass R Sunstein. 2021. *Noise: A Flaw in Human Judgment*. Little, Brown Spark.

[55] Jacob Kastrenakes and James Vincent. 2023. Hope, Fear, and AI. The Verge. https://tinyurl.com/3ywjah7d.

[56] Jamie Keene. 2012. Copyright, cover songs, and YouTube's Content ID: how the law is stifling creativity. The Verge. https://tinyurl.com/ex2jet95.

[57] Donal Khosrowi, Finola Finn, and Elinor Clark. 2023. Diffusing the Creator: Attributing Credit for Generative AI Outputs. In *Proceedings of the 2023 AAAI/ACM Conference on AI, Ethics, and Society*. 890–900.

[58] Celeste Kidd and Abeba Birhane. 2023. How AI can distort human beliefs. *Science* 380, 6651 (2023), 1222–1223.

[59] Taemie Kim and Pamela Hinds. 2006. Who should I blame? Effects of autonomy and transparency on attributions in human-robot interaction. In *proc. of the IEEE International Symposium on Robot and Human Interactive Communication*. 80–85.

[60] Nils Köbis and Luca D Mossink. 2021. Artificial intelligence versus Maya Angelou: Experimental evidence that people cannot differentiate AI-generated from human-written poetry. *Computers in human behavior* 114 (2021), 106553.

[61] Mauritz Kop. 2019. AI & intellectual property: towards an articulated public domain. *Tex. Intell. Prop. LJ* 28 (2019), 297.

[62] Justin Kruger and Thomas Gilovich. 1999. " Naive cynicism" in everyday theories of responsibility assessment: On biased assumptions of bias. *Journal of Personality and Social Psychology* 76, 5 (1999), 743.

[63] Jyh-An Lee. 2021. Computer-generated Works under the CDPA 1988. *Artificial Intelligence and Intellectual Property* 2021-65 (2021).

[64] Katherine Lee, A Feder Cooper, and James Grimmelmann. 2023. Talkin"Bout AI Generation: Copyright and the Generative-AI Supply Chain. *arXiv preprint arXiv:2309.08133* (2023).

[65] Mark A Lemley and Bryan Casey. 2020. Fair learning. *Tex. L. Rev.* 99 (2020), 743.

[66] Amanda Levendowski. 2018. How copyright law can fix artificial intelligence's implicit bias problem. *Wash. L. Rev.* 93 (2018), 579.

[67] Gabriel Lima, Meeyoung Cha, Chihyung Jeon, and Kyung Sin Park. 2021. The conflict between people's urge to punish AI and legal systems. *Frontiers in Robotics and AI* 8 (2021), 756242.

[68] Gabriel Lima, Nina Grgić-Hlača, and Meeyoung Cha. 2021. Human perceptions on moral responsibility of AI: A case study in AI-assisted bail decision-making. In *Proceedings of the 2021 CHI conference on human factors in computing systems*. 1–17.

[69] Gabriel Lima, Nina Grgić-Hlača, and Meeyoung Cha. 2023. Blaming humans and machines: What shapes people's reactions to algorithmic harm. In *Proceedings of the 2023 CHI conference on human factors in computing systems*. 1–26.

[70] Gabriel Lima, Changyeon Kim, Seungho Ryu, Chihyung Jeon, and Meeyoung Cha. 2020. Collecting the public perception of AI and robot rights. *Proceedings of the ACM on Human-Computer Interaction* 4, CSCW2 (2020), 1–24.

[71] Chiara Longoni, Andrey Fradkin, Luca Cian, and Gordon Pennycook. 2022. News from generative artificial intelligence is believed less. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*. 97–106.

[72] Bertram F Malle, Matthias Scheutz, Thomas Arnold, John Voiklis, and Corey Cusimano. 2015. Sacrifice one for the good of many? People apply different moral norms to human and robot agents. In *2015 10th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE, 117–124.

[73] Gregory N Mandel. 2014. The public perception of intellectual property. *Fla. L. Rev.* 66 (2014), 261.

[74] Gregory N Mandel. 2016. What is IP for Experiments in Lay and Expert Perceptions. . *John's L. Rev.* 90 (2016), 659.

[75] Andreas Matthias. 2004. The responsibility gap: Ascribing responsibility for the actions of learning automata. *Ethics and information technology* 6 (2004), 175–183.

[76] Ryan Merkley. 2023. On AI-Generated Works, Artists, and Intellectual Property. Lawfare. https://tinyurl.com/yc5xr8bb.

[77] Dale T Miller and Michael Ross. 1975. Self-serving biases in the attribution of causality: Fact or fiction? *Psychological bulletin* 82, 2 (1975), 213.

[78] Kobe Millet, Florian Buehler, Guanzhong Du, and Michail D Kokkoris. 2023. Defending humankind: Anthropocentric bias in the appreciation of AI art. *Computers in Human Behavior* 143 (2023), 107707.

[79] R Matthew Montoya, Robert S Horton, Jack L Vevea, Martyna Citkowicz, and Elissa A Lauber. 2017. A re-examination of the mere exposure effect: The influence of repeated exposure on recognition, familiarity, and liking. *Psychological bulletin* 143, 5 (2017), 459.

[80] Alexander Nehamas. 1986. What an author is.

[81] Leonardo Nicoletti and Dina Bass. 2023. Humans Are Biased. Generative AI Is Even Worse. Bloomberg. https://tinyurl.com/2p8ynpay.

[82] Michael I Norton, Daniel Mochon, and Dan Ariely. 2012. The IKEA effect: When labor leads to love. *Journal of consumer psychology* 22, 3 (2012), 453–460.

[83] Matt O'Brien. 2024. ChatGPT-maker braces for fight with New York Times and authors on 'fair use' of copyrighted works. The Associated Press. https://tinyurl.com/bdduxb24.

[84] US Copyright Office. 2023. Copyright Registration Guidance: Works Containing Material Generated by Artificial Intelligence.

[85] Stefan Palan and Christian Schitter. 2018. Prolific. ac—A subject pool for online experiments. *Journal of Behavioral and Experimental Finance* 17 (2018), 22–27.

[86] Sung Min Park, Kristian Georgiev, Andrew Ilyas, Guillaume Leclerc, and Aleksander Madry. 2023. Trak: Attributing model behavior at scale. *arXiv preprint arXiv:2303.14186* (2023).

[87] Sebastian Porsdam Mann, Brian D Earp, Sven Nyholm, John Danaher, Nikolaj Møller, Hilary Bowman-Smart, Joshua Hatherley, Julian Koplin, Monika Plozza, Daniel Rodger, et al. 2023. Generative AI entails a credit–blame asymmetry. *Nature Machine Intelligence* 5, 5 (2023), 472–475.

[88] Emily Pronin, Justin Kruger, Kenneth Savtisky, and Lee Ross. 2001. You don't know me, but I know you: The illusion of asymmetric insight. *Journal of Personality and Social Psychology* 81, 4 (2001), 639.

[89] Garima Pruthi, Frederick Liu, Satyen Kale, and Mukund Sundararajan. 2020. Estimating training data influence by tracing gradient descent. *Advances in Neural Information Processing Systems* 33 (2020), 19920–19930.

[90] Kevin Roose. 2022. An A.I.-Generated Picture Won an Art Prize. Artists Aren't Happy. The New York Times. https://tinyurl.com/d84dux2u.

[91] Michael Ross and Fiore Sicoly. 1979. Egocentric biases in availability and attribution. *Journal of personality and social psychology* 37, 3 (1979), 322.

[92] Mia Sato. 2023. AI-generated fiction is flooding literary magazines — but not fooling anyone. The Verge. https://tinyurl.com/2tvt44at.

[93] Mia Sato and Nilay Patel. 2023. YouTube is going to start cracking down on AI clones of musicians. The Verge. https://tinyurl.com/ms2frpbz.

[94] Niloufer Selvadurai and Rita Matulionyte. 2020. Reconsidering creativity: copyright protection for works generated using artificial intelligence. *Journal of Intellectual Property Law & Practice* 15, 7 (2020), 536–543.

[95] James Shepperd, Wendi Malone, and Kate Sweeny. 2008. Exploring causes of the self-serving bias. *Social and Personality Psychology Compass* 2, 2 (2008), 895–908.

[96] Jan Smits and Tijn Borghuis. 2022. Generative AI and intellectual property rights. In *Law and artificial intelligence: regulating AI and applying ai in legal practice*. Springer, 323–344.

[97] Barbara A Spellman and Frederick Schauer. 2008. Artists' moral rights and the psychology of ownership. *Tul. L. Rev.* 83 (2008), 661.

[98] Giovanni Spitale, Nikola Biller-Andorno, and Federico Germani. 2023. AI model GPT-3 (dis) informs us better than humans. *Science Advances* 9, 26 (2023), eadh1850.

[99] Gemini Team, Rohan Anil, Sebastian Borgeaud, Yonghui Wu, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut, Johan Schalkwyk, Andrew M Dai, Anja Hauth, et al. 2023. Gemini: a family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805* (2023).

[100] Leigh Thompson and George Loewenstein. 1992. Egocentric interpretations of fairness and interpersonal conflict. *Organizational Behavior and Human Decision Processes* 51, 2 (1992), 176–197.

[101] Nitasha Tiku, Kevin Schaul, and Szu Yu Chen. 2023. These fake images reveal how AI amplifies our worst stereotypes. The Washington Post. https://tinyurl.com/2fwk3b9j.

[102] Kevin Tobia. 2022. Experimental jurisprudence. *The University of Chicago law review* 89, 3 (2022), 735–802.

[103] Tom R Tyler and Jonathan Jackson. 2014. Popular legitimacy and the exercise of legal authority: Motivating compliance, cooperation, and engagement. *Psychology, public policy, and law* 20, 1 (2014), 78.

[104] UK Intellectual Property Office. 2022. Consultation outcome: Artificial Intelligence and Intellectual Property: copyright and patents. https://tinyurl.com/3bn9vzwb.

[105] Ibo Van de Poel. 2015. The problem of many hands. In *Moral responsibility and the problem of many hands*. Routledge, 50–92.

[106] Yong Wan and Hongxuyang Lu. 2021. Copyright protection for AI-generated outputs: The experience from China. *Computer Law & Security Review* 42 (2021), 105581.

[107] Yuqian Wang. 2024. Beijing Internet Court Grants Copyright to AI-Generated Image for the First Time. Associated Press. https://tinyurl.com/5cx8syvc.

[108] Robert B Zajonc. 1968. Attitudinal effects of mere exposure. *Journal of personality and social psychology* 9, 2p2 (1968), 1.