AI & Copyright: A Case Study of the Music Industry

Lila Shroff
Stanford University

I. INTRODUCTION

(Un)Stable Diffusion: The Emergence of Copyright Lawsuits over Generative AI

At the end of 2022, litigators filed a class-action lawsuit against three artificial intelligence (AI) image-generation companies for their use of Stable Diffusion, an AI model that allows for the creation of novel images from text-based prompts (Butterick, 2023; Stable Diffusion Online, 2023). The lawsuit, filed on behalf of three artists, was announced in a blog post explaining that Stable Diffusion is "a 21st-century collage tool that violates the rights of artists" (Butterick, 2023). Not long after, a new website called "Stable Diffusion Frivolous" sprung up (Stable Diffusion Frivolous, 2023). This website, created by "tech enthusiasts uninvolved in the case," offered a line-by-line rebuttal to the initial lawsuit (Stable Diffusion Frivolous, 2023). During this period, in a different case, Getty Images filed a lawsuit against Stability AI, the company behind Stable Diffusion, due to the "brazen infringement of Getty Images' intellectual property on a staggering scale" (Vincent, 2023).

Figure 1. Vincent, 2023. Original image on left, generated image on right. Notice how the copyright information has been modified in the generated image.

While neither of these lawsuits has been resolved—and it might be a long time before they are—they illustrate a rising tension between artists and the growing capabilities of AI in the

realm of creative production. This is not a tension limited only to the visual arts—similar advances using artificial intelligence to generate film, audio, music, and text have led to widespread complaints and concerns from artists of all types (Livni et al., 2023). Some worry about losing their jobs to automation, while others feel that using AI to generate art is cheating the creative process (Du Sautoy, 2019). In part, such anxiety is to be expected: these modern complaints are echoes of comments made by previous artists in eras (Hertzmann, 2018). As Aaron Hertzmann argues, history shows that "the advent of new technologies often causes fear of displacement among traditional artists" (2018, p. 2). For instance, the introduction of photography in the nineteenth century led to significant pushback, especially from portraiture painters who felt their livelihoods were being threatened (Hertzmann, 2018). In an essay that perfectly captures this sentiment, poet Charles Baudelaire writes that photography is "art's most mortal enemy" (1857).

However, despite this historical precedent for many of the concerns currently being voiced by artists, there is a novel question that has emerged rooted in the way that generative AI systems are trained. Namely, the training of generative AI systems requires models to be fed incomprehensibly large amounts of data during a training phase which assist the models in later using pattern recognition to produce novel creations (Brown, 2023). This process results in large amounts of copyrighted materials being used in the training of AI systems, and when these models are operated by commercial agents, we must ask ourselves: what should the rights be of the human actors who have produced creative work which is used—without their permission or credit—to train AI systems deployed for the financial gain of others? Answering these questions will help regulators to create better policies and lead technologists to design more human-centered technology.

Engaging with Artists' Perspectives to Propose HC-AI Principles for Creative Work

Already, many legal scholars, technologists, and corporate lawyers have offered opinions on this topic. However, missing from this conversation are the voices of actual creative workers themselves. Not only do these workers serve as primary data providers for the artificial intelligence models in question, but the ongoing development of these models can be seen as a threat to their financial livelihoods (Eloundou et al., 2023). More simply put, artists' own work is being used to train the very models that may replace them, and to proceed without considering the perspectives of these creative workers is to set a dangerous precedent. The question of what rights human artists deserve in relation to the development of generative AI systems has substantial implications beyond the scope of creative work—in the coming decades, as artificial intelligence becomes increasingly capable of performing feats once considered to be purely human, we must better understand how we can develop the law to protect and prioritize the rights of humans over machines. In the following paper, I develop a set of principles outlining what the rights of artists should be with respect to the use of their work in the training and deployment of generative AI systems. Namely, I examine the music industry to understand artists' perspectives in order to arrive at principles of (a) increased inter-stakeholder communication, (b) dataset transparency requirements, (c) the ability for an artist to opt out of a training dataset, and (d) fair application of "fair use" law.

First, I will introduce background context on how AI models are trained, and specific ways that the design of such technology can be seen as being in tension with the interests of creative workers. Next, I will explore leading legal perspectives on the topic, and evaluate areas where they may fall short in considering artists' perspectives. In this background section, I will

also briefly introduce the concept of Human-Centered AI and its potential for informing a healthier future of AI development ("Human-Centered," n.d.). Then, I will move to a case study examining the use of AI in the music industry and consider artists' perspectives on the promises and perils of generative AI systems. Finally, I'll propose a set of artist-centered principles to inform the development of copyright law that combines the opinions reported by artists with a human-centered approach to AI development.

As a writer and technologist myself, I bring my own opinions from experimenting with these technologies in my creative practice. This personal understanding can act as an advantage, as I have direct insight into the space, however, I am also cognizant of the potential biases that may emerge as a result—namely, my position in favor of increased artist-centric regulation. To mitigate these personal biases, I will make sure to draw from a balanced and widely-representative set of artists' perspectives. I will further work to ensure that the set of principles proposed at the end of the essay is representative of the analyzed artists' opinions, and not of my own.

II. THE COPYRIGHT CONVERSATION

Generative AI Systems are Trained on Datasets Containing Copyrighted Data

Before going deep in AI and copyright policy, it is necessary to first understand how AI systems work. Although there are many nuances to the technology, the most important thing to understand is that these models are trained on massive datasets. For example, one of the most popular Large Language Models, the GPT series which powers ChatGPT, is trained on large swaths of the internet including many Wikipedia and Reddit pages (Romero, 2021). Much of the text that GPT-3 is trained on was produced by writers who never expected their work to be used for training AI systems (Xiang, 2022; Esposti et al., 2020). The dataset powering Stable Diffusion is trained using LAION-5B, a "dataset consisting of 5.85 billion … image-text pairs" (Beaumont, 2022). Included in this dataset is everything from "hacked and stolen nonconsensual porn" to "living artists' artwork (Xiang, 2022).

We have already seen that there is a long history of artist concern over technological displacement, however, the present concerns of artists over generative AI present a novel set of fears (Hertzmann, 2018). Currently, the companies that develop generative AI models are engaged in the practice of scraping the internet for as much data as possible to help them to develop better generative models. However, this practice completely ignores the human labor that went into creating this data. In addition, these models are designed to create novel media—eventually as complex as feature-length films or novels—at minimum cost. This means that not only do these models build on artists' work without compensating them, but they are also designed to automate the work that artists do. To make matters worse, the legal policy surrounding generative AI, especially regarding copyright law, is ambiguous and confusing (Moe, 2023).

The Legal Questions: Debates over Data Provider Compensation, Fair Use & Ownership

In a recent piece in the *New York Times*, correspondent Ephrat Livni, who herself is an artist and lawyer, thoughtfully outlines the three major questions being debated about the future of copyright law: (1) "What is owed to the creators of the original material?" (2) "Does 'fair use' apply?" and (3) "Who owns the output of generative AI?" (Livni et al., 2023).

To the first question—on what is owed to the creators of the original material—some have suggested that artists should be able to opt-in to a training set, instead of having companies automatically assume that their work can be used for training AI models (Butterick 2023). Others have argued in favor of artist compensation for data provision, however, given the scale of the datasets used to train AI models, this is often regarded as infeasible—although certain solutions to the issue of scale have been proposed (Lnu, 2023). The second question, on whether "fair use" applies, examines a legal doctrine in U.S. copyright law that permits "the unlicensed use of copyright-protected works in certain circumstances" including in news reporting and education ("U.S. Copyright," 2023; "Fair use," n.d.). For example, a professor copying an article from a magazine to distribute in class is fair use, while the professor posting that same article publicly on her course website would not be considered air use (Pascual, n.d.). Some believe that fair use applies for machine learning purposes—most vocally, Stanford Law Professor Mark Lemley argues in "Fair Learning" that such systems "should generally be able to use databases for training, whether or not the contents of that database are copyrighted," although, certain uses such as explicit intent to reproduce existing creative work must not be protected (Butler 2023; Tavernise et al., 2023; Lemley & Casey, 2021).

The final question considers who owns the output of generative AI. Although some have made the radical suggestion that AI systems themselves should be given ownership rights, most rebuke this perspective. As Aaron Hertzmann argues, art is "a fundamentally social act of expression and communication" (2018, p. 20). Thus, Herzmann continues, granting AI with authorship "falsely attributes emotions, feelings, and ethical weight to that AI" (2018, p. 21). Computational poet and professor Allison Parrish agrees, writing that poems are created through human intention alone—even if language models can generate infinitely poetic language, the action of creating a poem is impossible for a computer (2021). A proposed balanced approach to ownership of output states "the natural person(s) behind the arrangements necessary for the creation or the invention at stake should be considered as the author/inventor" (Vesala et al., 2019, p. 2). However, given the complexity of AI systems, this third argument allows for " i) the designer(s) of the system, ii) the data provider(s), and iii) the user(s) of the system" to all be considered as deserving authorship status (Vesala et al., 2019, p. 14). As has been already addressed in the previous paragraph, the notion of how one determines how much each party has contributed, in order to determine ownership, is far from clear cut.

Human-Centered AI: Developing Technology to Empower the Human Experience

The concept of "human-centered AI" provides a critical framework for approaching legal questions related to AI (HCAI). IBM defines HCAI as an "emerging discipline intent on creating AI systems that amplify and augment rather than displace human abilities" ("Human-Centered," n.d.). More concretely, as is shown in Figure 2 below, instead of an emphasis on complete and total automation of human activities, proponents of HCAI believe in prioritizing the combined capabilities of humans and machines ("Human-Centered," n.d.).

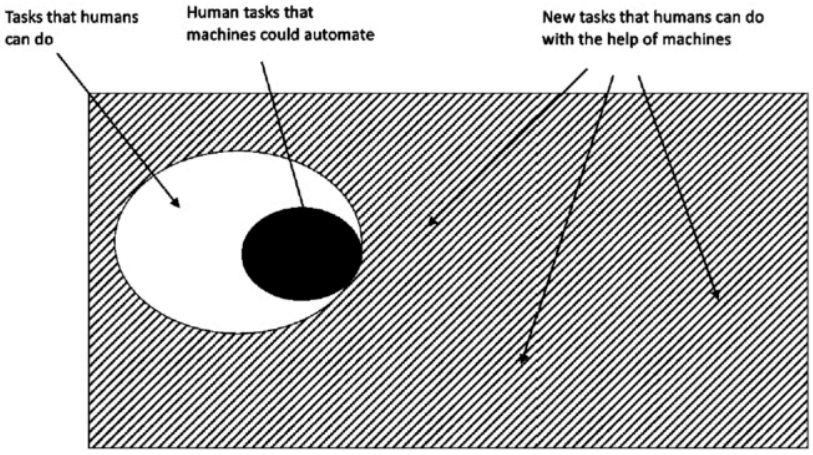

Figure 2. Brynjolfsson, 2022. A visual representation of the HCAI framework.

With regard to the application of HCAI principles to the creative domain, IBM suggests that human users should be responsible for "specification, goal setting, steering, high-level creativity, curation, and governance," while AI systems can "augment human abilities through inspiration, creativity, low-level detail work, and the ability to design at scale" ("Human-Centered," n.d.). In other words, if humans are responsible for the creative vision, machines can help with the operationalizing of such visions.

HCAI can be considered as a subset of a broader philosophy of human-centered design which originated at Stanford University in 1958 ("NeuroDesign, n.d.). Formally, human-centered design is described as "a problem-solving technique that puts real people at the center of the development process" (Landry, 2023). While critics argue that the dominance of human-centered design has led to it being unquestioningly accepted as the sole mode of development in certain technological spaces, it remains a respected framework used widely across fields today (Norman, 2005). The discipline of HCAI is much more recent and its emerging popularity can be at least partially traced to the founding of the Stanford Human-Centered AI Institute (HAI) in 2019. In the post announcing HAI's launch, the institute is described as a space of "multidisciplinary collaboration" which intends "to advance artificial intelligence…to improve the human condition" (Adams, 2019). The US chief data scientist at Deloitte, Jim Guzzcza, aptly summarizes the framework by arguing that "design can help close the gap between AI algorithm *outputs* and improved human *outcomes"* (Guszcza, 2018). Ultimately, this framework provides an important baseline for analyzing artists' perspectives: does the HCAI view of creative work align with artists' own opinions and preferences? If so, how can the HCAI framework be further developed alongside these creative perspectives, with the end goal of helping shape copyright law? If not, what is a better artist-first approach to these legal debates?

III. A LACK OF HARMONY IN THE MUSIC INDUSTRY

A History of Technology and the Music Industry

Over the past couple of decades, the music industry has undergone a series of extraordinary transformations—and generative AI is posed to lead to yet another. An analysis of

the music industry in the digital age suggests that if the transition of music distribution from the physical to the digital has allowed for increased access to music consumption, the coming years will give way to "services and features that allow users to play *with* music rather than merely play music" (Wikström, 2014). This shift is already happening. One of the most popular applications of generative AI to music production is through music synthesis models which allow for the creation of novel compositions of music "in response to text, audio, or image prompts" (Butler, 2023). This year alone has already seen the introduction of over fifteen new such models (Spielburg, 2023). However, the use of these models is not without controversy.

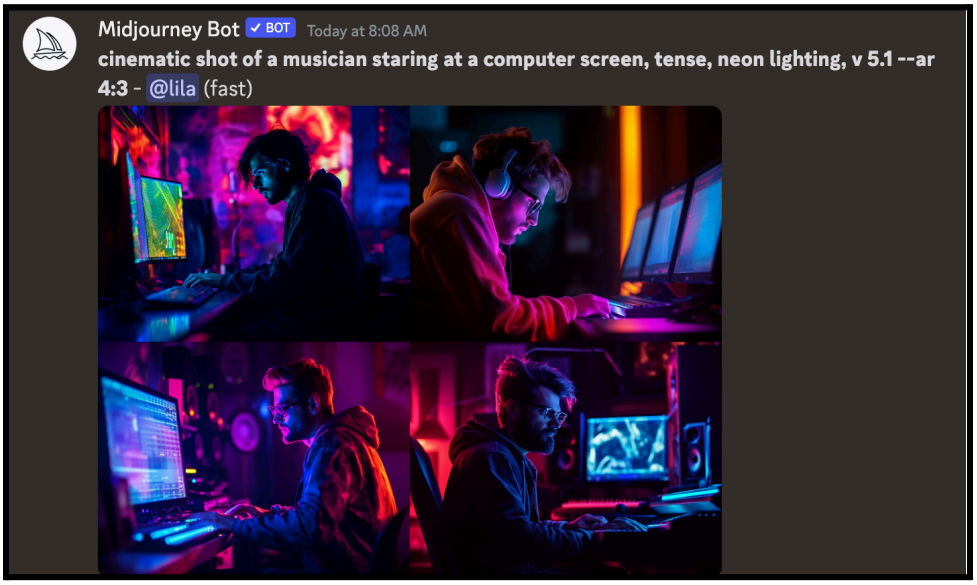

Figure 3. "Midjourney," 2023. Screenshot of a prompt for an image of a "musician staring at a computer screen." [1]

Drake, the Deep Fake

In April 2023, a song called "Heart on My Sleeve" went viral (Tavernise et al., 2023). The song sounded as if it had been produced by Drake and the Weeknd, two hugely renowned musical artists—except, it wasn't (Tavernise et al., 2023). Instead, the song was an example of a "musical deep fake:" it was an AI-generated song that sounded exactly like Drake and the Weekend, produced by an anonymous independent creator (Tavernise et al., 2023). Immediately following the song's release, Drake's label, Universal Musical Group issued a statement criticizing "Heart on My Sleeve" (Tavernise et al., 2023). The group claimed that "training of generative AI using our artists' music" is a breach of copyright law (Brydon, 2023). They then made an even stronger statement positioning music industry stakeholders who reject AI on "the side of artists, fans, and human creative expression" while describing those who embrace AI as "on the side of deep fakes, fraud and denying artists their due compensation" (Brydon, 2023). This led to the song being removed from Spotify, YouTube, and other relevant streaming services (Tavernise et al., 2023). But this example illustrates what some have noted as a "major

[1] Notice how all the musicians generated appear to be white and male-presenting. This leads to the important discussion of AI systems perpetuating historical biases which is outside the scope of this essay but equally important.

contradiction" in the way major music industry players are considering AI (Spielburg, 2023). As one music research agency reports that "nearly every major label has an equity stake in a music AI startup" (Spielburg, 2023). This leaves us to ask: what do the artists think?

Artists' Perspectives in Favor of AI

One of the easiest places to find perspectives of musical artists on AI music is in the spaces where such music is being created. One such space is "AI Hub," a Discord server with tens of thousands of users (at the time of writing this) where members produce, share, and teach others how to use AI to create music (Xiang, 2023). Members of this community argue that AI models increase accessibility and opportunity within the closed-off music industry. As one Ukrainian producer explains, AI offers an opportunity for producers and songwriters to work with singers they'd never otherwise have access to—which can be particularly important for talented creators who don't otherwise have a network (Xiang, 2023). Other artists feel that AI can be a tool for financial benefit. The Canadian artist Grimes shared this opinion from an artist's perspective in a tweet reading:

"I'll split 50% royalties on any successful AI generated song that uses my voice. Same deal as I would with any artist i collab with. Feel free to use my voice without penalty. I have no label and no legal bindings" (2023).

Others see AI music as a mode of extended fandom: as one artist shared, people are only interested in creating deep fakes that mimic the work of real human musicians because these artists have already done the hard work of creating original music (Xiang, 2023). Through this lens, experimenting with another artist's work is a form of flattery and increased recognition for the original artist, rather than theft. Finally, some believe that most interesting things happening in the musical AI space are among artists working to push the boundaries of music production altogether. Artist Nao Tokui believes that AI allows for a "new form of art expression."

Artists' Fears About AI

Despite such positivity in favor of AI-generated music—there are many fears. Some worry that traditional music labels may begin to demand that artists hand over exclusive rights to their voices. Others are concerned that record labels and streaming platforms will exchange human artists for more profitable synthetic counterparts and that the entire process of music creation will become automated from start to finish (The_Phonecian, 2023). There are also critical concerns about cultural appropriation (Xiang, 2023).

Another difficulty many express concerns over is the distribution of AI-generated music. Some producers have posted AI music with disclaimers invoking the fair use doctrine, only to have their work removed from major platforms like YouTube (Xiang, 2023). Here, we see another contradiction: Google, which owns YouTube, is able to train AI music models on copyrighted material, yet YouTube won't let the output of such models be spread on its platform (Stewart, 2023). Conversely, in a different case, musical technologists CJ Carr and Zack Zukowski were hit with a copyright claim after creating a track where a Frank Sinatra bot sang a Britney Spears song (Stewart, 2023). With help from the Electronic Frontier Foundation, a nonprofit which defends "digital privacy, free speech, and innovation," they were able to argue fair use to convince YouTube to allow the song to remain up ("Electronic", n.d.; Stewart, 2023).

Importantly, it is not as if all musicians fall into the camp of being strictly for or against AI—many creatives hold mixed perspectives. Given the music industry's technological past, Tokui is wary of the present gap "between AI engineers, AI practitioners, and actual artists" (Houser, 2023). Tokui also has concerns over artist dependence on corporate-run technologies, explaining that "it's super important that this process is not governed only by big companies like Adobe or OpenAI" (Houser, 2023). The landscape of AI music is very much still in its infancy and it's critical to have musicians engaging in these conversations while the technology is still in development (Houser, 2023). Pop singer Claire L. Evans agrees that artists need to be in conversation with technologists—although she herself fears the technology will be used to streamline the necessary complexities of the creative process (Stewart, 2023a).

IV. HC-AI PRINCIPLES FOR CREATIVE WORK

I now propose a set of *HCAI Principles for Creative Work.* To do so, I integrate the artists' perspectives outlined in the previous section with the human-first approach of HCAI. In the previous section, we saw that there was no one monolithic perspective that musicians have toward the use of AI technology in music production. Those most directly involved in the production of AI music appreciate the increased accessibility the technology allows, and some are even experimenting with novel methods of music creation. However, the distribution of AI music across contexts has proven challenging and the fair use doctrine is applied unevenly and often in favor of those in power. Artists' fears include displacement, technology dependance, and cultural appropriation. Almost everyone agrees that there needs to be greater dialogue between technologists, artists, and policymakers as these technologies are developed. HCAI aims to augment and not displace human labor. Therefore, an HCAI approach to creative technologies should aim to make more accessible and expand the set of creative possibilities humans have access to—rather than focus entirely on the sole replication of existing human work. If this sounds abstract, let me provide an example: an AI model for producing special effects in films should not aim to replace entire VFX teams, but rather, make such work more broadly available to a wider set of the population, while enhancing the capabilities of VFX experts. Thus, I propose an ethical framework that aims to protect against such displacement while encouraging increased human creativity as aided by advances in machine intelligence.

1. **Interstakeholder communication.** Technology cannot be built in the absence of those whom it most directly affects, just as legislation intended to protect constituents cannot be written without direct attention to the perspectives of the constituents themselves. It is critical that technologists, policymakers, and others working in this space make an explicit effort to engage directly with artists in building and regulating these tools. For concrete tactics that can be used to help foster interdisciplinary conversation, I suggest that large technology companies create or further invest in teams where artists and technologists work side-by-side in the product design and testing processes. These companies can also provide grant funding to artists to help support their work, especially as they test and identify issues with new technologies. Although policymakers can look to technology companies for insight into developing technology regulation, they should make a deliberate effort to balance these biased perspectives with the opinions of artists and everyday citizens.

2. **Dataset transparency.** Policymakers can't make regulatory decisions and develop robust AI policies if they don't have basic information about generative AI models. Already, some have begun to develop tools that allow artists to enter text or image queries to see if their work has been used in the training of large AI models ("Have I?", n.d.). Such tooling should be invested in and further expanded across other forms of media. At the same time, policymakers can require technology companies to engage in disclosing greater amounts of information regarding their training data and model architecture. The true challenge here will be balancing the need for greater transparency with corporate interest in protecting trade secrets.

3. **Opting Out.** Even among artists, perspectives toward advances in generative AI are widely varied. For instance, on the topic of ownership of AI model output, some artists were excited about revenue-sharing works created by their fans, while others were very reticent about the idea of any derivative works being made in their style. To the extent that artists are able to opt-out of having their work included datasets, they should have the right to do so. Presently, it is both technically and operationally complex to allow artists and other data providers to remove their data from training datasets. However, instead of buying into technological determinism and allowing the limitations of technology to override human agency, it is critical that technologists actively pursue the development of machine-learning techniques that allow artists to opt-out. If technologists fail to organically pursue the research of such techniques, policymakers can intervene and, through regulation, guide technological development. Inspirational precedent for such policy can be found in existing data privacy legislation: the EU's General Data Protection Regulation (GDPR) requires users to provide opt-in consent, while the California Consumer Privacy Act (CCPA) allows users the option to opt-out ("Opt-In," 2023).

4. **Fair application of the "fair use" law.** Fair learning refers to the invocation of the fair use doctrine as justification for the legality of the use of copyright data in training AI systems. Currently, large commercial players exploit this doctrine to avoid legal responsibilities when training for-profit models on copyrighted material. Yet, the individuals that use these models to produce AI-generated art in the style of existing artists are unable to distribute their derivative works on the very same tech platforms. While the lawfulness of fair learning should be determined by legal experts and policymakers, the present state of how fair learning is applied demonstrates a major power imbalance that favors large technology companies. At the very least, fair learning, or whatever policy is eventually created to replace it, should be designed such that private companies do not receive preferential treatment over individual artists.

These principles do not seek to be all-encompassing, but rather, a foundation upon which future creatives can build. It is the hope that this framework will inform the perspectives of those who are currently making decisions about the design, creation, and regulation of AI technologies that directly affect the livelihoods of artists. I urge technologists and engineers to integrate these principles into their work building new technology. However, if these principles are disregarded by technology builders, the onus will fall onto policymakers to draft legislation such that technologists are required to abide by these principles regardless. With increased discussion about AI policy in Washington, such regulation could be part of an all-encompassing federal

policy—however, there is also room for labor unions and similar organizations to bargain for these rights on behalf of those whom they represent (“Blueprint,” n.d.).

V. CONCLUSION: A PREMONITION OF THE FUTURE

I write this essay on the month anniversary of the Hollywood writers' strike where TV and film screenwriters are calling for—among other things—the regulation of artificial intelligence (Barco, 2023). Many writers fear their replacement by AI and are advocating for bans on the use of AI in the writing and rewriting of treatments and screenplays. They also argue that the writer’s work should not be used as data for training AI systems (Barco, 2023). However, this perspective is not entirely universal—some screenwriters have already begun using this technology in their world and are excited by its potential (Barco, 2023).

The rapid advances in generative AI over the past years and months have brought with it a novel set of concerns regarding where humans fit into a world of increasing AI capability. As Kirby Ferguson writes in *Freethink,* “Artists are feeling this pain most intensely because their skills are among the first to be rivaled” (2023). Not only are artists seeing their creative work be replicated en masse without their consent, but they are also feeling the direct threats of displacement. Some are choosing to embrace the technology while others are stepping away from it. Regardless, the tension between artists and AI models is indicative of the tensions that will be felt in other domains as AI begins to take on more and more types of human work. By thinking through *HCAI Principles for Creative Work,* we also begin to imagine how HCAI principles can be applied to other domains.

References

Adams, A. (2019, March 18). Stanford University launches the Institute for Human-Centered Artificial Intelligence. *Stanford News*. https://news.stanford.edu/2019/03/18/stanford_university_launches_human-centered_ai/

Barco, M. del. (2023, June 2). Hollywood writers still going strong, a month after strike began. *NPR*. https://www.npr.org/2023/06/02/1179883812/wga-strike-one-month-mark

Baudelaire, C. (1859). *On Photography*. https://www.csus.edu/indiv/o/obriene/art109/readings/11%20baudelaire%20photography.htm

Beaumont, R. (2022, March 31). *LAION-5B: A NEW ERA OF OPEN LARGE-SCALE MULTI-MODAL DATASETS | LAION*. https://laion.ai/blog/laion-5b

Blueprint for an AI Bill of Rights | OSTP | The White House. (n.d.). Retrieved June 3, 2023, from https://www.whitehouse.gov/ostp/ai-bill-of-rights/

Brown, S. (2023, April 19). *Machine learning, explained*. MIT Sloan. https://mitsloan.mit.edu/ideas-made-to-matter/machine-learning-explained

Brydon, A. (2023, April). *Heart on My Sleeve | AI-generated Drake and Weeknd track taken down after copyright claim from Universal Music Group*. WhyNow. https://whynow.co.uk/read/heart-on-my-sleeve-ai-generated-drake-and-weeknd-track-taken-down-after-copyright-claim-from-universal-music-group

Brynjolfsson, E. (2022). The Turing Trap: The Promise & Peril of Human-Like Artificial Intelligence. *Daedalus*, *151*(2), 272–287. https://doi.org/10.1162/daed_a_01915

Butler, B. (2023, May 8). *What AI can teach us about copyright and fair use*. Freethink. https://www.freethink.com/robots-ai/what-ai-can-teach-us-about-copyright-and-fair-use

Butterick, M. (2023, January 13). *Stable Diffusion litigation · Joseph Saveri Law Firm & Matthew Butterick*. https://stablediffusionlitigation.com/

Carle, E. (2023, April 11). *Ask a Techspert: What is generative AI?* Google. https://blog.google/inside-google/googlers/ask-a-techspert/what-is-generative-ai/

Du Sautoy, Marcus. (2019). The Creativity Code : Art and Innovation in the Age of AI. Belknap Press.

Electronic Frontier Foundation. (n.d.). Electronic Frontier Foundation. Retrieved May 10, 2023, from https://www.eff.org/

Eloundou, T., Manning, S., Mishkin, P., & Rock, D. (2023). *GPTs are GPTs: An Early Look at the Labor Market Impact Potential of Large Language Models* (arXiv:2303.10130). arXiv. http://arxiv.org/abs/2303.10130

Esposti, M. D., Lagioia, F., & Sartor, G. (2020). The Use of Copyrighted Works by AI Systems: Art Works in the Data Mill. *European Journal of Risk Regulation*, *11*(1), 51–69. https://doi.org/10.1017/err.2019.56

Ferguson, K. (2023, May 8). *Why artists shouldn't fear AI*. Freethink. https://www.freethink.com/robots-ai/ai-and-future-of-creativity

Grimes, [@Grimezsz]. (2023, April 23). *I'll split 50% royalties on any successful AI generated song that uses my voice. Same deal as I would with any artist* [Tweet]. Twitter. https://twitter.com/Grimezsz/status/1650304051718791170?s=20

Guszcza, J. (2018, May 1). AI Needs Human-Centered Design. *Wired*. https://www.wired.com/brandlab/2018/05/ai-needs-human-centered-design/

Have I Been Trained? (n.d.). Retrieved May 11, 2023, from https://haveibeentrained.com/

Hertzmann, A. (2018). *Can Computers Create Art?* (arXiv:1801.04486). arXiv. http://arxiv.org/abs/1801.04486

Houser, K. (2023, April 17). *Google's AI music generator is like ChatGPT for audio*. Freethink. https://www.freethink.com/robots-ai/ai-music-generator

Landry, L. (2020, December 15). *What Is Human-Centered Design? | HBS Online*. https://online.hbs.edu/blog/post/what-is-human-centered-design

Lemley, M., & Casey, B. (2021). Fair Learning. *Texas Law Review*, *99*(4). https://texaslawreview.org/fair-learning/

Livni, E., Hirsch, L., & Kessler, S. (2023, April 17). *Who Owns a Song Created by A.I.? - The New York Times*. https://www.nytimes.com/2023/04/15/business/dealbook/artificial-intelligence-copyright.html

Lnu, D. (2023, February 11). Artists enable AI art—Shouldn't they be compensated? *The Gradient*. https://thegradient.pub/should-stability-ai-pay-artists/

Midjourney (2023). *Midjourney* (Version 5.1) https://www.midjourney.com/app/

Moe, A. (2023, April 19). *Copyright law complicated by AI technology | WisBusiness*. WisBusiness. https://www.wisbusiness.com/2023/copyright-law-complicated-by-ai-technology/

Norman, D. A. (2005). Human-centered design considered harmful. *Interactions*, *12*(4), 14–19. https://doi.org/10.1145/1070960.1070976

NeuroDesign Research. (n.d.). Retrieved June 3, 2023, from https://neurodesign.stanford.edu/about

Parrish, A. (2021, August 13). *Language models can only write poetry*. Allison Posts. https://posts.decontextualize.com/language-models-poetry

Pascual, C. (n.d.). *Research Guides: Copyright and Fair Use: Common Scenarios*. Retrieved June 3, 2023, from https://csulb.libguides.com/copyrightforfaculty/scenarios

Romero, A. (2021, May 24). *A Complete Overview of GPT-3—The Largest Neural Network Ever Created*. Towards Data Science. https://towardsdatascience.com/gpt-3-a-complete-overview-190232eb25fd

Spielburg, Y. (2023, February 1). *A field guide to the latest music AI tidal wave—Water & Music*. Water & Music. https://www.waterandmusic.com/a-field-guide-to-the-latest-music-ai-tidal-wave/

Stable Diffusion Frivolous · Because frivolous lawsuits based on ignorance deserve a response. (2023). Retrieved April 16, 2023, from https://stablediffusionfrivolous.com/

Stable Diffusion Online. (n.d.). Retrieved April 22, 2023, from https://stablediffusionweb.com/

Stewart, B. (2023, May 8). *A view from the creative frontlines*. Freethink. https://www.freethink.com/robots-ai/a-view-from-the-creative-frontlines

Stewart, B. (2023, May 8). *These "AI Artists" don't want to ape existing art—They want to create something new entirely*. Freethink. https://www.freethink.com/robots-ai/nao-tokui-and-dadabots-want-to-create-new-music

Tavernise, S., Reid, W., Johnson, M. S., Feldman, N., Baylen, L. O., Benoist, M., Powell, D., Lozano, M., & Wood, C. (2023, April 28). The Ballad of 'Deepfake Drake.' *The New York Times*. https://www.nytimes.com/2023/04/28/podcasts/the-daily/ai-deepfake-drake.html

The Difference Between Opt-In and Opt-Out Principles In Data Privacy: What You Need To Know. (2023, January 15). Https://Secureprivacy.Ai/. https://secureprivacy.ai/blog/difference-beween-opt-in-and-opt-out

The_Phonecian. (2023, May 8). *Rick Beato Makes Super-Bleak Prediction About AI's Impact on Music Industry: "They're Going to Get Rid of the Artists."* Ultimate Guitar Com. https://www.ultimate-guitar.com/news/general_music_news/rick_beato_makes_super-bleak_prediction_about_ais_impact_on_music_industry_theyre_going_to_get_rid_of_the_artists.html

U.S. Copyright Office Fair Use Index. (2023, February). U.S. Copyright Office. https://www.copyright.gov/fair-use/

Vesala, J., Honkkila, O., & Pihlajarinne, T. (Eds.). (2019). AI-generated content: Authorship and inventorship in the age of artificial intelligence. In *Online Distribution of Content in the EU* (pp. 117–135). Edward Elgar Publishing. https://doi.org/10.4337/9781788119900.00015

Vincent, J. (2023, January 17). *Getty Images is suing the creators of AI art tool Stable Diffusion for scraping its content*. The Verge. https://www.theverge.com/2023/1/17/23558516/ai-art-copyright-stable-diffusion-getty-images-lawsuit

What is "Fair Use"? - Legal Help. (n.d.). Retrieved May 12, 2023, from https://support.google.com/legal/answer/4558992?hl=en

What is human-centered AI? (2022, March 31). IBM Research Blog. https://research.ibm.com/blog/what-is-human-centered-ai

Wikström, P. (2014). The Music Industry in an Age of Digital Distribution. *OpenMind*. https://www.bbvaopenmind.com/en/articles/the-music-industry-in-an-age-of-digital-distribution/

Xiang, C. (2022, September 26). *AI Is Probably Using Your Images and It's Not Easy to Opt Out*. Motherboard: Tech by Vice. https://www.vice.com/en/article/3ad58k/ai-is-probably-using-your-images-and-its-not-easy-to-opt-out

Xiang, C. (2023, April 24). *Inside the Discord Where Thousands of Rogue Producers Are Making AI Music*. Vice. https://www.vice.com/en/article/y3wdj7/inside-the-discord-where-thousands-of-rogue-producers-are-making-ai-music

Xu, W. (2019). *Toward Human-Centered AI: A Perspective from Human-Computer Interaction*. Intel Corporation.