The Starving Artist

Maddi Smith
msmith@oxy.edu
Occidental College

Abstract

This paper examines the ethical implications of Generative Artificial Intelligence in an artistic context through the lens of power, accessibility, accountability, and representation. While Generative AI is often promoted as a tool for productivity, creativity, and innovation, its rapid adoption has also introduced serious concerns regarding environmental impact, labor displacement, algorithmic bias, lack of transparency, and growing market concentration among major technology firms. This paper explores how these issues affect artists and creative workers in particular, whose labor has historically been undervalued and is now increasingly vulnerable to automation. In addition, this paper considers the unequal accessibility of AI systems, arguing that barriers such as paywalled features, limited technological literacy, and biased outputs prevent all users from benefiting equally. Concepts like fairness, accountability, transparency, Design Justice, and cultural representation are used to evaluate how these systems are developed and who has influence over their future. Ultimately, this paper argues that Generative AI is neither entirely beneficial nor entirely harmful, but that its long-term impact will depend on whether human dignity, consent, sustainability, and inclusive decision-making are prioritized over efficiency and profit alone.

1 Introduction

Throughout the entire process of creating my mini-project, Itera Studio, I have been hyper-aware of the ethical implications. After all, my social media and news feeds are regularly populated with updates of technological advancements, politics, how the two intersect, and their correlation with the "declining" state of the world. Recently, I have become fatigued by these reports, as they are primarily negative and persistent. Isn't technology supposed to make our lives easier? What is the cost of higher productivity in our late-stage capitalistic society? What does it entail for the future of humanity? What are the ramifications for people similar to myself: women, people of color, students, and emerging artists specifically?

I vividly remember the release of generative AI, as tools like Midjourney and DALL-E2 suddenly became ubiquitous almost overnight. I believe the catalyst of AI art was the infamous video of Will Smith eating a plate of spaghetti in 2023. At the time, many netizens and artists, including myself, were not concerned with Artificial Intelligence becoming our competitor: the capabilities of generative AI at the time were limited and easy to distinguish. Mutated frames, extra limbs, malformed hands, and distorted backgrounds were easily discernible, and even somewhat comedic at the time.

In 2026, it is very difficult to differentiate between real art and generative AI. So much so, that animation studios and large corporations like Pixar, Lionsgate, and Netflix are aggressively incorporating Artificial Intelligence into their workflows in order to increase productivity and cut the cost of commissioning traditional animators.

As an undergraduate computer science student, I can see the value in using Artificial Intelligence as a means of productivity and educational scaffolding, but as an artist and humanitarian, I am wary of the lengthy concerns regarding Artificial Intelligence in artistic practice.

In this paper, I aim to clearly establish the larger ethical considerations, power, accessibility, accountability, and cultural implications of Generative Artificial Intelligence in an artistic context.

2 Ways to Think About Ethics Research

2.1 Environmental Impact

Recently, I have noticed that when conversing with peers and reflecting on the drawbacks of Artificial Intelligence, environmental impact is one of the first issues that comes to mind. Before directly addressing these concerns, it is important to understand that there are multiple sub-fields of Artificial Intelligence, each with its own unique architecture. The most notorious, and most commonly used in artistic contexts, is Generative AI.

GenAI is environmentally taxing in a variety of ways: electricity consumption, extraction of rare resources, e-waste production, water consumption, and greenhouse gas emissions. It is important to note that these impacts are not exclusive to AI and extend to other forms of technology,

but the extent to which AI is predicted to worsen them is alarming. Additionally, measuring the cost of GenAI is difficult due to a lack of reliable data, and there are multiple approaches to doing so.

When an individual enters a prompt into a GenAI model, it is broken into smaller units called tokens. The model then goes through data collection, where it is fed data from across the internet. Next, it is trained through neural networks, learning relationships between words, pixels, or sound. For language models, this involves predicting sequences in a sentence, while for image models, it deconstructs existing images and reconstructs them based on the prompt, pixel by pixel. After this, the model outputs an image, audio, or text, which can then be refined through Reinforcement Learning from Human Feedback (RLHF).

Experts at the Association of Computing Machinery emphasize that developing GenAI models is an extremely resource-intensive process, requiring large-scale data collection, careful model selection, and the integration of multiple pre-existing systems [7]. Researchers also use Life Cycle Assessment (LCA) to evaluate environmental impact across a model's full lifecycle, including manufacturing, usage, and end-of-life. However, even measuring this energy use is inconsistent. Estimates based on thermal design power (TDP) often overestimate GPU consumption, while software-based measurements can differ significantly from hardware power meters, making it difficult to accurately determine the true environmental cost of training GenAI models.

After consulting multiple sources in order to find concrete data, I found that measuring factors such as water usage is still difficult. According to Molly Taft for Undark, who interviewed individuals working at data centers, the amount used depends heavily on the individual data center: using more water means data centers can avoid running electric cooling systems, while using more electricity lessens the water footprint but increases the power bill and greenhouse gas emissions [10]. However, it is important to note that this is not an insignificant issue, as access to clean drinking water is already limited. According to the Environmental and Energy Study Institute, large data centers can consume up to 5 million gallons per day, equivalent to the water use of a town populated by 10,000 to 50,000 people [6].

Although there may not be one exact number for the total amount of e-waste, CO2 emissions, and water usage, holistically, the environmental impact of GenAI is far from insignificant.

2.2 Power

In our post-cancel-culture, post-Epstein-file-release society, scapegoating, public outrage, and pointing fingers often feel ineffective when it comes to creating real structural change. Arguably, the average person has at least some awareness of socioeconomic power imbalances, especially in big tech. In *Algorithms of Oppression*, Safiya Umoja Noble argues that many users are unaware of how companies like Google collect user data, shape visibility online, and profit through advertising systems. Her larger point is that these corporations should not be viewed as neutral platforms, but as institutions with significant influence over how information is distributed.

At this point, though, many corporations and elites are no longer widely seen as neutral or trustworthy. We live in an age of disinformation and declining trust in major institutions. It does not take much research to see that control over many leading GenAI systems is concentrated among a relatively small number of firms. A 2023 Bruegel report highlighted growing concentration in cloud computing among companies such as Amazon, Microsoft, and Google, whose infrastructure is central to training and deploying many AI systems [4]. The same report explains that GenAI companies rely on cloud providers to train, run, and deploy their models, while cloud providers use GenAI as a driver for expanding market share. While the tech industry is not entirely uniform, many of the highest positions of power remain concentrated among wealthy white men, particularly in Silicon Valley. These imbalances matter because they shape who benefits from AI systems, who is harmed by them, and whose interests are prioritized.

During the FAccT ACM Conference 2025, scholars from University of Michigan presented research suggesting that marginalized groups often report more negative attitudes toward AI than dominant social groups [3]. Their findings showed that gender minorities expressed significantly less positive attitudes toward AI than cisgender participants, and that women reported less positive attitudes than men. These views likely reflect documented harms tied to algorithmic systems, such as biased healthcare tools that underserve disabled people, facial recognition systems that disproportionately misidentify Black individuals and transgender people, and surveillance technologies used without meaningful consent.

Additionally, artists and creatives may be especially vulnerable to labor disruption caused by generative AI. On March 5th, 2026, Anthropic released research estimating occupational exposure to AI systems, with arts and media ranking among the more exposed categories [2]. This is especially concerning because many creative fields already involve unstable income, freelance precarity, and oversaturated labor markets. Since artistic labor has historically been undervalued, generative AI adds another layer of uncertainty for workers who were already economically vulnerable.

As for global impact, AI is increasingly being framed as

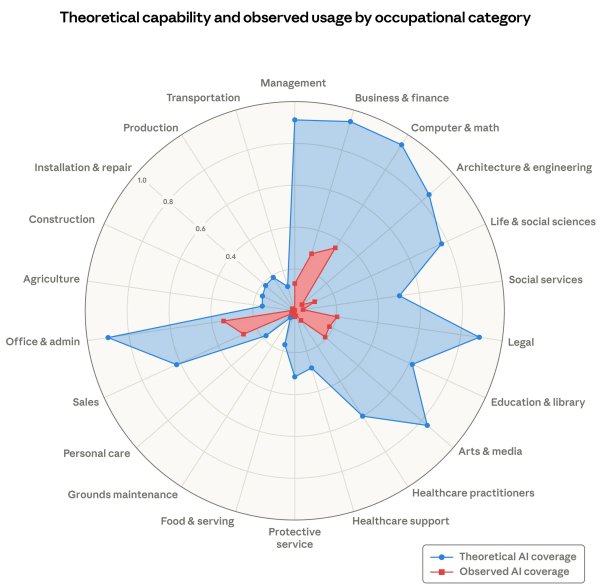


Figure 1: Occupational exposure to AI systems...

Figure 2: Occupational exposure to AI systems by category, from Anthropic's 2026 labor market impact research [2]. Arts and media rank among the most exposed categories.

a driver of economic growth, productivity, and geopolitical competition. Countries with the capital, infrastructure, and technical talent to implement AI at scale may continue to benefit the most, while countries with fewer resources risk falling further behind. If current trends continue, AI may deepen existing global wealth gaps rather than reduce them.

2.3 Accessibility

ChatGPT, Claude, and various Generative AI models have extensive free tiers. Consequently, one might expect that these tools are easily leveraged and accessible to individuals with varying capabilities. However, not all users receive the same experience or benefit equally from utilizing AI. While it is true that Generative AI has been used to improve closed-captioning, text-to-speech, translation, content simplification, and a variety of other beneficial functions, many of these tools are behind a paywall. Although there is typically a free tier, its capabilities are limited. Additionally, utilizing AI has not always produced better results, depending on the user and how they are generating text, images, or speech. Although physical disability and accessibility issues are a major part of this discussion, media literacy is another significant accessibility barrier.

The popularity of Generative AI has increased significantly, yet not everyone understands whether something is AI-produced or whether the output they are receiving is reliable or high quality. For many users, AI functions as a black box rather than a technology they are encouraged to understand. This lack of technological literacy can disproportionately harm already marginalized groups, who may be more vulnerable to misinformation, bias, or overreliance on flawed outputs.

Recent research presented at the 2026 ACM CHI Conference examined this issue through a study of 344 participants focused on improving media literacy and prompting practices [8]. The researchers argued that coaching users to write more inclusive prompts should not shift responsibility away from developers, but instead serve as a user-centered strategy that helps people recognize and navigate algorithmic bias while using these systems. The study also notes that traditional media literacy education often happens before exposure to a technology, which may be unrealistic in fast-moving digital environments. Many users prioritize convenience and efficiency, and most cannot reasonably be expected to complete training before using a tool. Instead, the researchers propose more seamless interventions that occur during real-time use, allowing users to build awareness while interacting with AI systems. This suggests that accessibility is not only about physical usability or cost, but also whether people are given the knowledge and support necessary to engage with these technologies critically and effectively.

2.4 Fairness, Accountability, and Transparency

According to the ACM Code of Ethics [1], computing professionals are expected to prioritize the public good, avoid harm, and be honest, fair, and transparent in their work. I think these are important values, but ethical codes alone cannot solve ethical problems, especially when there is little accountability or enforcement. In the context of Generative AI, many companies promote innovation while giving limited transparency about training data, labor practices, environmental costs, or how biased outputs are handled. This makes it difficult for users, artists, and researchers to make informed decisions about the tools they are using or supporting. If the field of Computer Science genuinely values fairness and accountability, these principles need to be backed by stronger disclosure requirements, independent audits, clearer consent standards, and meaningful consequences when harm is caused.

2.5 Design Justice and Imagination

Terms like Design Justice appeal to me because it focuses on changing how technology is created, who is included in the process, and who ultimately benefits from it [5]. Rather than waiting until harm is already done, Design Justice asks

whether the people most affected by a system had any voice in shaping it. In the context of Generative AI, that question feels especially relevant. Many creative tools are marketed as revolutionary, yet artists are often treated as sources of training data rather than active collaborators, decision-makers, or people who deserve compensation.

I recently visited Asteria, a company that works with both artists and AI in film production. Seeing this firsthand was both reassuring and unsettling. On one hand, they emphasized how much human labor still goes into the final product. Artists were still refining outputs, directing visuals, revising concepts, and doing the creative work necessary to make anything usable. That challenged the common assumption that AI simply replaces artists with the push of a button. On the other hand, it also made clear how complicated "ethical AI" can be in practice. They discussed cleaner models trained purely on licensed or paid artist work, which sounds ideal in theory, but those models can be more limited in scope, style range, and flexibility. Because of those limitations, broader industry workflows still rely on less transparent models trained on scraped shapes, textures, and images.

That experience made me realize that the issue is rarely as simple as AI being entirely good or entirely bad. Design Justice would suggest that the future of these tools should not be decided only by investors, executives, or engineers, but also by artists, workers, and the communities most affected by them. If AI is going to shape the future of creative industries, then transparency, consent, compensation, and inclusion should be treated as necessities rather than optional features.

2.6 Cultural Competence and Representation

Cultural competence becomes especially important in Generative AI because the people designing, funding, and deploying these systems are not always representative of the people most affected by them [11]. Throughout this paper, I have discussed how leadership and ownership in the AI space are often concentrated among a relatively narrow group, while many of the harms—such as labor displacement, biased outputs, and lack of transparency—are felt more broadly. When workplaces are not inclusive or do not meaningfully value different backgrounds and experiences, those blind spots can easily be reflected in the technologies they produce.

Representation also matters in creative industries, where Generative AI is increasingly being integrated into workflows. If decisions about the future of art and media are primarily being made by executives, investors, or engineers rather than artists themselves, then creative workers lose influence over tools that directly affect their livelihoods. I have already discussed how many artists feel anxiety over job security, authorship, and compensation. Seeing companies like Asteria involve artists in production showed me that collaboration is possible, but it also highlighted how uncommon that model can feel. Representation is not only about who appears in advertisements or leadership photos, but who has real decision-making power over how technology is built, used, and who benefits from it [9].

3 Conclusion

Throughout writing this paper, I kept returning to the same epiphany: Generative AI is not a simple issue, and it cannot be framed as entirely good or entirely bad. As a computer science student, I can understand the appeal of efficiency, automation, and productivity, especially when these tools can assist with learning, accessibility, and creative workflows. At the same time, as an artist, I cannot ignore the environmental costs, labor displacement, bias, lack of transparency, and power imbalances that continue to shape how these systems are built and used.

What concerns me most is not the existence of the technology itself, but who controls it, who benefits from it, and who is expected to absorb its harms. If the future of Generative AI is decided only by executives, investors, and corporations, then many of the same inequalities that already exist in society will likely be reproduced through new tools. However, if artists, workers, educators, and marginalized communities are meaningfully included in creating and refining these systems, then AI may have the potential to support rather than replace human creativity.

Ultimately, I do not believe the conversation should be about whether AI should exist at all, but about what values guide its development. Efficiency alone should not be the goal. If these technologies are going to become embedded in everyday life and creative industries, then fairness, accountability, consent, sustainability, and human dignity should be treated as necessities rather than afterthoughts.

References

[1] ACM. *The Code affirms an obligation of computing professionals to use their skills for the benefit of society.* en. URL: https://www.acm.org/code-of-ethics.

[2] Anthropic. *Labor Market Impacts of AI*. 2026. URL: https://www.anthropic.com/research/labor-market-impacts.

[3] "Attitudes Toward AI Among Marginalized Groups". In: *Proceedings of the ACM FAccT Conference*. 2025. DOI: 10.1145/3715275.3732081. URL: https://dl.acm.org/doi/10.1145/3715275.3732081.

[4] Bruegel. *Cloud Market Concentration and Generative AI*. 2023. URL: https://www.jstor.org/stable/resrep55201.

[5] Costanza-Chock, Sasha. *Design Justice: Community-Led Practices to Build the Worlds We Need*. en. MIT Press, 2020. URL: https://direct.mit.edu/books/oa-monograph/4605/Design-JusticeCommunity-Led-Practices-to-Build-the.

[6] Environmental and Energy Study Institute. *Data Centers and Water Consumption*. 2024. URL: https://www.eesi.org/articles/view/data-centers-and-water-consumption.

[7] "Generative AI and Environmental Impact". In: *ACM Digital Library* (2025). DOI: 10.1145/3725984. URL: https://dl.acm.org/doi/10.1145/3725984.

[8] "Improving Media Literacy and Prompting Practices for Generative AI". In: *Proceedings of the 2026 ACM CHI Conference on Human Factors in Computing Systems*. 2026. DOI: 10.1145/3772318.3791542. URL: https://dl.acm.org/doi/10.1145/3772318.3791542.

[9] Pantic, Katarina et al. "Drawing a computer scientist: stereotypical representations or lack of awareness?" In: *Computer Science Education* 28.3 (2018), pp. 232–254. DOI: 10.1080/08993408.2018.1533780. URL: https://doi.org/10.1080/08993408.2018.1533780.

[10] Taft, Molly. *The Hidden Water Cost of AI*. 2024. URL: https://undark.org/author/molly-taft/.

[11] Washington, Alicia Nicki. "When Twice as Good Isn't Enough: The Case for Cultural Competence in Computing". In: *Proceedings of the 51st ACM Technical Symposium on Computer Science Education*. SIGCSE '20. New York, NY, USA: Association for Computing Machinery, 2020, pp. 213–219. ISBN: 9781450367936. DOI: 10.1145/3328778.3366792. URL: https://doi.org/10.1145/3328778.3366792.