Mitigating Generative AI's Impact on Artists: A Research Wishlist

Jingna Zhang*
January 2025

Abstract

While capabilities research in AI advances rapidly, research to protect and support creatives affected by generative AI remains underexplored. This white paper presents a wishlist of tools identified through discussions with working artists, researchers, and machine learning professionals at the NeurIPS 2024 Creativity & Generative AI Workshop, offering a guide for researchers interested in addressing the challenges faced by creative communities impacted by generative AI.

1. Introduction: Creatives and Generative AI

Creative work is an integral part of our daily lives and experiences living in this world. The typography of texts we read, the designs of apps and phones we use, the furniture we utilize at work, the music, films, books, and video games we consume that accompany us through our brightest memories and darkest hours. Every day, we interact with images, designs, and stories crafted by creatives. Their work influence how we see and experience the world, our culture, and how we communicate and connect with one another.

Yet by 2026, it is estimated that approximately 203,800 entertainment industry jobs will be disrupted by generative AI in the US for those working in film, TV, games, and animation[1]. In China, some gaming companies have reportedly reduced illustration staff by up to 70% following AI adoption[2]. The International Confederation of Societies of Authors and Composers (CISAC), which represents more than 5 million creators worldwide, estimated that those working in the music sector will lose almost 25% of their income within the next 4 years due to generative AI[3].

Research for generative AI has taken off at unprecedented speed and scale since the boom of consumer-facing products like Stable Diffusion, Midjourney, and ChatGPT. While the machine learning community continues to push the boundaries of what AI can do, creatives are facing stark and immediate backlash to their livelihoods, professional identities, and creative practices with this technology—all without their consent nor compensation for their work.

This appropriation of intellectual property has occurred at massive scale with no technical safeguards for the creators whose work form an indispensable part of this technology. While legal and legislative processes take their long arcs, I seek to identify potential directions for the research community who wish to work on solutions to support artists and creatives deeply affected by these developments.

As both the founder of Cara—a social and portfolio platform for creatives who want a digital experience free from generative AI—as well as a photographer whose work has appeared in Vogue, Elle, and Harper's Bazaar, I stand at an intersection of technology and art communities that gives me perspectives into both worlds, where I observe both the hype and excitement surrounding AI, as well as the real, tangible harms experienced by working artists whose livelihoods are threatened.

---

*Contact: info@zhangjingna.com

2. Current Impact on Creative Industries

Recent industry reports document significant disruption across multiple creative fields:

- The US entertainment industry is projected to see 203,800 payroll job disruptions by 2026, including 21.4% of film/TV/animation jobs and 13.4% of gaming workforce. Freelancers who are unaccounted for in these estimates could be particularly vulnerable due to a lack of unions and job loss protection. (CVL Economics, 2024)
- The global audiovisual sector is estimated to see incomes shrink by nearly 25% in next 4 years (CISAC, 2024)
- Gaming industry recruiter in China reported a 70% drop in illustrator jobs following the AI boom, with one studio laying off a third of their character design team (Rest of the World, 2023)
- Online labor markets supported an estimated 9.3 million freelancers and 2.2 million worldwide at near full-time capacity in 2020 (Kässi et al.[4]). In 2 separate large-scale analyses of such online marketplaces, the negative impact of before and after generative AI boom was consistent between 2021 to 2024 was consistent: listings for writing jobs declined between 30-33%, translation by 19%, customer service by 16%, and graphic design and 3D modeling roles by 17% after a year with no signs of demand rebounding (Bloomberry, 2024[5]; Demirci et al., 2024[6])

Beyond employment statistics, creatives have reported growing pressure to incorporate AI tools into their workflow, regardless of personal preference. Many face the predicament of using AI or lose work opportunities entirely.

3. Common Misconceptions

In discussions with both research and creative communities at the roundtable and outside of the event, several common misconceptions regarding generative AI technology were discussed:

3.1 Misconceptions from the research community:

- **Unaware of artist business models:** Many researchers are unaware that creative workers' leverage and business model often rely on copyright and the licensing of their work, and don't realize that generative AI directly displaces their jobs while using their own work in AI training without prior consent or licensing.
- **Assumption that resistance to AI is due to resistance to change and new technologies:** There's an assumption that instead legitimate concerns about consent and compensation, artists are resistant to AI due to being resistant to new technology. In reality, artists have experimented with GANs years before generative AI in non-commercial settings, and new technologies in computer graphics and animation were often developed by entertainment and production studios or workers (for example: Photoshop was created by Thomas and John Knoll while John was at Industrial Light and Magic, a visual effects studio).
- **Assumption that AI will complement rather than replace:** There is strong belief that AI will complement creative work rather than replace workers, "if only artists could see it". But this is contradicted by documented reports of job losses and reduction in wages with the introduction of generative AI tools.
- **Severe underestimation of scale and immediacy of impact:** There is vast underestimation of the scale and immediacy of generative AI's impact on the creative industry, with many researchers not realizing that the technology already negatively impacts the lives of many people today.
- **"Not seeing AI's potential":** Even back in 2015, I would encounter researchers like Yann LeCun at conferences and discuss AI's applications for personal workflows. Similar to the example of GANs—professional artists have long thought about ways that AI *could* be utilized for our personal work. However, right now, I and most artists' main concern is in addressing the ethical issues around consent and compensation for the unauthorized usage of our work, which must be resolved first before any dialogue on potential beneficial applications can take place. The creative communities impacted by generative AI need ways to protect, prevent, and mitigate current harms, and any other discussions would only dilute the conversation around current issues, which is already under-supported in research.

- **Blockchain misconceptions:** Some in the blockchain communities believe that blockchain can help with provenance and proof of authorship, however, the central issue isn't the difficulty of simply proving who posted a JPG on the web—it's in proving authorship for materials that have been ingested for AI training, and letting that information survive through the generative AI process, which blockchain does not solve.
- **That "AI learns just like a human"**: This is a fundamentally flawed perspective for several reasons:
    - **A human is not a machine:** A human artist, unlike a machine, cannot ingest billions of images and then generate billions out based on stored, massive training datasets. The current deployment of generative AI also creates entirely different economic consequences than human learning—a single AI is capable of displacing tens of thousands of workers simultaneously, a human artist is not.
    - **Market participation:** Many human artists financially support creative ecosystems through education, museum admissions, libraries, and purchasing art and reference materials. Yet other than a select number of licensing deals with media companies—current AI training largely circumvents these ecosystems entirely, removing attribution and compensation to individual authors, creators, and artists that put out the very works that teach and inspire.
    - **Respect for artists' copyright**: In traditional creative industries, using others' copyrighted work for commercial purposes typically required licensing or permission from the original artist and/or rights holder. The use of LAION[7] and LibGen[8] shows that this is not the case in the age of generative AI.
    - **Scale of impact:** A human artist produces a limited number of work over decades, AI systems can generate millions of outputs daily that directly compete in commercial spaces.

3.2 Misconceptions from the creative community:

- **An all-or-nothing belief about protective measures being meaningless if it's not perfect:** "If it can't protect against everything, what's the point?" Many are unaware that layered protection is still meaningful. For example, a photographer adding a visible watermark to their photo shows an attempt to protect their copyright prior to generative AI. Even if the watermark can be removed, this protective action could prove helpful in legal action and professional use.
- **Misconception that any one solution might prevent unauthorized scraping**: There is a misconception that surely one technology can emerge that can fully prevent unauthorized scraping. Unfortunately, as of time of writing, there is no known technology that can stop any bad actor from scraping the web unauthorized. This remains the general consensus of all security experts and specialists. The protections available today are steps on a scale that seek to reduce unauthorized use, protect, or help with certain preventions. But no technology can currently stop a bad actor who does not respect directives to not scrap if they intend to ignore such instructions. Furthermore, common suggestions like image splicing and disabling of right clicks on a website can easily be bypassed, leaving only the downside of making an art director's job harder if they wished to save an artist's work for project presentations as viable hiring candidates and give up instead due to the trouble.
- **Underestimating speed of AI improvement:** Those not working in AI can often underestimate how quickly research and technology can improve what humans typically consider difficult (e.g., gen AI hands, fingers, eyes details from 2022 versus now are unimaginably different). Today's goalpost for what looks bad or impossible will likely be possible tomorrow. A false sense of security may cause creatives who have yet to be affected by generative AI to feel secure in their current roles, believing that their jobs would not be impacted.
- **Overestimation of human AI detection ability**: There is an overestimation of human ability when it comes to accurately identifying AI-generated imagery. A 2024 study[9] found that general viewers achieved only 59% accuracy, with 41% chance of both false positives (identifying human work as AI) and false negatives (identifying AI work as human). Professional and expert artists reached 75% and 83% accuracy respectively. In contrast, an AI detector in the study achieved 88% accuracy, and when expert artists used AI detection tools, the accuracy rate improved to 93%.

These misconceptions need to be understood by both communities in order to correct current misunderstandings, and is a vital step in bringing the communities together for collaborative efforts and support.

4. Current Tools, Adoption, and Limitations

Several tools and solutions have emerged in the last two years to help artists protect their work and navigate the web since the rise of generative AI. Below is a non-comprehensive list featuring the most commonly utilized tools by the visual arts community currently:

- **Glaze**[10][11]: Disrupts style mimicry in LoRA training - 6.7 million downloads
- **Mist**[12][13]: Disrupts style mimicry in LoRA training - commercial interest from companies like Dropbox for adoption
- **Nightshade**[14]: Image poisoning to protect unauthorized image use - 2.2 million downloads
- **AI detection services**: Various companies are in the business of AI-generated content detection, with Hive notably valued at $2 billion in 2021[15]
- **Cara**[16][17]**:** Social and portfolio-sharing platform for visual artists that filters out generative AI images - over 1.25 million users. An integration with Glaze launched in 2023 was paused due to overwhelming response causing resource constraints.

4.1 Limitations

Despite the widespread adoption of these tools, they face limitations:

- **Compute requirements**: Tools like Glaze, Nightshade, and Mist all require substantial compute resources, limiting accessibility for local use and present substantial costs for web use.
- **Lack of web and easy-to-access options**: The multi-step setup requirement to use tools like Glaze, Nightshade, and Mist before a user can share their work online creates a high friction point. Cara's Glaze integration was the closest ease-of-use case but was halted due to resource constraints. It is however strong evidence that there is demand for such integrations or supportive tools that can ease this process for social media sharing.
- **Maintenance and arms-race**: Tools need continuous updates to keep up with new AI models, and continued maintenance and distribution for effectiveness for users.
- **Limited coverage of formats**: Most tools focus on static images, leaving video, 3D, and other forms of media files unprotected.
- **Scale challenges**: Independent platforms struggle to offer protection features due to computational costs.
- **Text detection reliability**: AI detection for text continues to be less accurate than images and audio.
- **Lack of protection for text:** While images, and even videos and 3D media can embed files that have some level of protection tools applied, there is none that can currently be done for a writer's words to act as indication that they do not want their work scraped or used without consent. (Participants and I are both aware that this may be an impossible ask, but given the recent news of LibGen being scraped by Meta, I wanted to emphasize the lack of protection for writing here as there is a large gap with no answers for the writers community.)
- **Trust factor**: While C2PA was often brought up as a tool for content provenance to help a human creator prove that something is human-made, due to its committee members being made up of companies currently creating generative AI models such as Adobe, Google, Meta, and OpenAI, there is a sense of distrust about adoption of the system and whether it might be used for other purposes without users consenting—in a way similar to how images are currently being used by many of these AI companies without users' consent.

5. List of Tools Artists Wish For

Based on discussions with working creatives across various disciplines in the visual arts, below is a compilation of tools, features, and research that artists and others in the art industry wish would exist. It is our belief that research in these areas would have a strong and direct positive impact on artist communities and people currently impacted by generative AI.

5.1 For Artists and Creatives

For many artists and creatives, they do not want their work to be scraped and/or used in generative AI without their consent. In any shape or form, they would appreciate tools that can:

1. **Help them protect their work from unauthorized AI use**: Disrupt results whether it's foundation models, LoRA, img2img or other technologies—whether through disruption, exclusion, or other adversarial or currently yet unknown techniques.
2. **Indicate that they were the original creator of their work:** Visible watermarks that survive training, metadata/ author information that survives training.
3. **Preserve opt-out information**: The EU AI Act protects opt-out but it is impossible to enforce. Currently, artists can use robot.txt to opt-out their own websites from AI crawlers. However, when someone else reposts their work on other websites or social media, there is no protection and preservation of their opt-out indication. Is there a way to embed a non-removable signature and rights management information in the image file that AI companies would respect in the same way their crawlers respect the robot.txt instructions?
4. **A way to search or know for certain whether an artist's work has been used to train AI**: It is currently possible to search the LibGen database on *The Atlantic* and LAION on *haveibeentrained*. However, these are only possible by first obtaining the known dataset. The question is—whether combined with the watermark/ surviving training meta embed or not—could there be ways for users to search and find out if their work have been used to train an AI?
5. **Less compute-intensive tools**: Current tools like Glaze, Mist, and Nightshade are all compute-heavy. More lightweight tools that are accessible on the web would be invaluable for many creatives.

5.2 For Companies, Brands, and General Uses

Beyond individual artists, there are broader needs both in commercial settings and for general internet users:

Companies are increasingly facing legal and reputational risks from inadvertently using or promoting AI-generated content. Recruiters, art directors and art teams are having to spend more time screening and filtering out job applicants that use AI-generated portfolios. At the same time, public sensitivity towards AI-generated media is growing, and many in the creative sector as well as regular art lovers have expressed frustration with being unable to filter out gen AI media while browsing the web. For these needs:

1. **A way to filter out AI media or results when using the web:** perhaps in the form of a browser plugin, for the entire internet.
2. **Any research that can reduce the environmental impact** or the energy consumption caused by AI.
3. **Tools to protect individuals' likeness** from AI training.
4. **Ways to alert users to AI training risks**: Help non-technical users understand when their data may be used for AI training.
5. **Ways to enable data strike/boycott**[18]: Browser plugin/mobile app that helps users better withhold their data from companies and services they use. Either by indicating settings options that can be changed, and/or policies that may be harmful to them so that they are aware while using such services.

6. Roundtable Findings

This roundtable was attended by 15 participants, with 8 in the arts and 7 in research and technology sectors. Additional feedback were collected from speaking with dozens of other artist and researcher attendees at the Creativity & AI workshop that this roundtable was a part of over the course of two days, as well as surveying users on Cara. Before going further, I would like to note that this paper is meant as a starting point for ideas and brainstorming, and I hope the reader understands that these suggestions may not be comprehensive, as generative AI affects different artistic communities in different ways.

The following are some points of interests from these discussions:

1. **Technical solutions often require playing catchup with new AI models**, as historically, adversarial work is often on the slower side. Participants wondered if there could be solutions that move beyond traditional adversarial work that can be more sustainable and have more lasting outcomes, with less maintenance required of both researchers and end users adopting such technologies.

2. **Incentive focus:** Several participants agreed that rather than relying only on technical solutions, perhaps attention should be turned to seeking ways to disincentivize companies from training on works without permission. For example, if data poisoning becomes standard and widespread, then companies may be incentivized to adopt an opt-in business model where they only obtain data from people who willingly want to contribute their work, leading to unauthorized use becoming less of a standard approach.

3. **Artists being vocal may attract research interest:** Some researchers strongly believe that if artists rallied around ideas for tools they wanted and talked about it loudly and frequently—it would be more likely to draw academic research interest. School labs and PhD students who may not yet be part of industry labs would especially be likely candidates, as they have some degree of freedom in choosing their research topics and direction, and open interest shows that a particular area will have meaningful impact.

4. **Awareness gap in AI:** Despite intense discourse and coverage about generative AI's impact on artists, there seems to still be pockets of academia and tech that are unaware of these discussions. At Ted Chiang's talk preceding the roundtables, multiple audience members indicated that it was their first time hearing about how artists feel. Continued attempts to discuss this topic with the academic community may be helpful.

5. **User experience and technical barriers for adoption:** Adoption and distribution of tools require work beyond technical papers. Many artists and designers would be happy to assist with such projects, so it is the creatives' wish that researchers considering working on such tools would take into consideration the final user experience, and seek help from the artist community if needed.

6. **Starting point for collaborations**: The current gulf between researchers and creative communities can be bridged if the two work together. During NeurIPS, I met many attendees who showed interest in contributing to helping artists but weren't sure about where to start. I recommend establishing a dedicated researcher-led group focused on artist-supporting AI technologies that can provide guidance and outreach advice for academics looking to explore the space.

7. Conclusion and Path to Action

Global spending on generative AI hit $2.3 billion in 2023 and jumped to $13.8 billion in 2024[19]. In the same year, companies in the space raised $56 billion[20]. I cannot predict the future, but based on news and records, it seems likely that generative AI will continue to be adopted commercially by companies across different sectors, though the success is yet to be determined.

What is certain, however, is that without the works of creatives—of every writer, artist, musician, and more who has shared their work on the internet—these models would not exist. And while governments should work quickly to regulate, and class actions will be played out in courts, the research community itself also holds immense power over the outcome and direction of research today.

I invite researchers to consider the topics and ideas discussed in this white paper, and I urge you to consider this direction of high-impact work that could meaningfully support some of the people most affected by current generative AI. If it would be possible to find new ways to mitigate some of the harms that this technology has inadvertently caused, your work can change lives.

As the organizer of this roundtable and founder of Cara, I welcome questions from researchers interested in pursuing these directions[21].

Acknowledgements

This white paper emerged from discussions at NeurIPS 2024 and the Creativity & Generative AI Workshop. I'd like to thank the workshop organizers for providing a meaningful opportunity for creatives to share their experiences and viewpoints and connect with the research community. I hope this event will inspire other similar initiatives in the future.

References

[1] CVL Economics, 2024. Future Unscripted: The Impact of Generative Artificial Intelligence on Entertainment Industry Jobs. https://animationguild.org/wp-content/uploads/2024/01/Future-Unscripted-The-Impact-of-Generative-Artificial-Intelligence-on-Entertainment-Industry-Jobs-pages-1.pdf

[2] Rest of World. 2024. AI is already taking video game illustrators' jobs in China. https://restofworld.org/2023/ai-china-video-game-layoffs-illustrators/

[3] Burke, K. 2024. Music sector workers to lose nearly a quarter of income to AI in next four years, global study finds. The Guardian. https://www.theguardian.com/music/2024/dec/04/artificial-intelligence-music-industry-impact-income-loss

[4] Kässi et al. 2021. How many online workers are there in the world? A data-driven assessment. https://pmc.ncbi.nlm.nih.gov/articles/PMC10445891/#sec8

[5] Bloomberry. 2024. The jobs being replaced by AI – an analysis of 5M freelancing jobs. https://bloomberry.com/i-analyzed-5m-freelancing-jobs-to-see-what-jobs-are-being-replaced-by-ai/

[6] Harvard Business Review. 2024. Research: How Gen AI Is Already Impacting the Labor Market. https://hbr.org/2024/11/research-how-gen-ai-is-already-impacting-the-labor-market

[7] Appel et al. 2023. Generative AI Has an Intellectual Property Problem. Harvard Business Review. https://hbr.org/2023/04/generative-ai-has-an-intellectual-property-problem

[8] Reisner, A. 2025. The Unbelievable Scale of AI's Pirated-Books Problem. The Atlantic. https://www.theatlantic.com/technology/archive/2025/03/libgen-meta-openai/682093/

[9] Ha et al. 2024. Organic or Diffused: Can We Distinguish Human Art from AI-generated Images? https://dl.acm.org/doi/10.1145/3658644.3670306

[10] Glaze - Protecting Artists from Generative AI. 2023. https://glaze.cs.uchicago.edu

[11] Shan et al. 2023. Glaze: Protecting Artists from Style Mimicry by Text-to-Image Models. USENIX. https://www.usenix.org/conference/usenixsecurity23/presentation/shan

[12] Mist. 2023. https://github.com/psyker-team/mist

[13] Liang et al. 2023. Adversarial Example Does Good: Preventing Painting Imitation from Diffusion Models via Adversarial Examples. https://arxiv.org/abs/2302.04578

[14] Nightshade - Protecting Copyright. 2024. https://nightshade.cs.uchicago.edu/

[15] Roof, K. 2023. Hive AI to Seek $200 Million in Funding and Valuation Jump. Bloomberg. https://www.bloomberg.com/news/articles/2023-08-11/hive-ai-said-to-seek-200-million-in-funding-and-valuation-jump

[16] Cara: Artist Social & Portfolio Platform. 2025. https://cara.app

[17] Silberling, A. 2024. A social app for creatives, Cara grew from 40k to 650k users in a week because artists are fed up with Meta's AI policies. TechCrunch. https://techcrunch.com/2024/06/06/a-social-app-for-creatives-cara-grew-from-40k-to-650k-users-in-a-week-because-artists-are-fed-up-with-metas-ai-policies/

[18] Vincent et al. 2021. Data Leverage: A Framework for Empowering the Public in its Relationship with Technology Companies. https://arxiv.org/abs/2012.09995

[19] Menlo Ventures. 2024: The State of Generative AI in the Enterprise. https://menlovc.com/2024-the-state-of-generative-ai-in-the-enterprise/

[20] Wiggers, K. 2025. Generative AI funding reached new heights in 2024. TechCrunch. https://techcrunch.com/2025/01/03/generative-ai-funding-reached-new-heights-in-2024/

[21] My contact is: info@zhangjingna.com