International Journal of Pedagogics

VOLUME Vol.06 Issue07 2026
PAGE NO. 75-80
DOI 10.37547/ijp/Volume06Issue07-15

The Mirror Paradox: AI as An Amplifier of Human Cultural Hierarchies in Contemporary Artistic Production

Erkinova Aigerim
Artist, Bishkek, Kyrgyzstan

Received: 10 June 2026; **Accepted:** 08 July 2026; **Published:** 31 July 2026

Abstract: Artificial intelligence is demonized by some, while others are neutral about it. This article offers a third perspective on AI: the "mirror paradox," where AI doesn't copy others' work but replicates the cultural hierarchy already embedded in the training data. The author identifies three types of such “mirroring”: structural (racial and gender biases in the training corpus), aesthetic (emphasizing the “popular” at the expense of originality), and economic (the mutually reinforcing logic of platforms and markets).

Keywords: AI art, mirror paradox, cultural hierarchies, GIGO, algorithmic bias, generative AI, authorship, rigorous creativity, cognitive assemblages, artistic labor.

1. Introduction: When Jason Allen's work "Théâtre D'Opéra Spatial" won first place in the 2022 Colorado State Fair Art Competition, the public response was highly controversial. Artists on social media called it "the biggest heist in art history" (Perhiniak , cited in Collins, 2026). Others, conversely, spoke of the dawn of a new era, where creativity is accessible to everyone using modern resources. The judges later stated that they had not realized that Midjourney had been used to create the work.

Debates about AI in art typically revolve around the same questions whether AI should be understood as a genuine tool or as a threat. They are based on the mistaken idea that AI acts on culture from somewhere outside, that is, either improving it or spoiling it. Generative AI introduces nothing fundamentally alien, drawing on what is already embedded in its training data, amplifying it, and producing results with a confidence that the original sources typically lack.

Collins (2026) notes in relation to the garbage-in, garbage-out (GIGO) principle in art that biases in training datasets reflect existing cultural trends. Since AI learns largely from commercially successful, Eurocentric art, it will reproduce and reinforce these very traits as the norm. And the more realistic the result, the less noticeable the choices embedded in the model's training become. This phenomenon is called the mirror paradox.

Collins's work convincingly describes this problem, but does not propose a clear solution. Manghani and Savage (2025) offer a different perspective on the issue, the concept of “rigorous creativity” as a counterweight to “AI spectacle,” that is, the pursuit of spectacular images at the expense of meaning in large-scale generative projects. Stork (2024), in turn, shows that computer vision is capable of finding patterns in old art that elude human historians, but he is interested in AI as a research tool.

This article brings together these three approaches and proposes a three-part model that explains how AI systems replicate cultural hierarchies at different levels.

2. THEORETICAL FRAMEWORK: FROM GIGO TO THE MIRROR PARADOX

2.1 GIGO as a Cultural Principle

According to GIGO, the quality of the output depends on the quality of the input data. That is, if the input data is distorted, the output will also be correspondingly

distorted. Applied to generative AI in art, this highlights the need to consider the cultural implications of data quality.

According to Collins (2026), learning occurs on vast arrays of images from the internet, archives, social media, and commercial sources. But images are never culturally neutral. They depend on who creates the art, who is allowed onto museum walls, who has the money to digitize the works, and what aesthetic norms are considered “standard.” By learning from this material, the system reproduces these distortions because these patterns are reproduced at scale.

As a result, at the level of individual images, Collins shows that generators often sexualize and racialize bodies, repeating old discriminations. An example is MIT researcher Heikkilä's experiment, where after entering her name and profession, she received hypersexualized fantasy images, while her male colleagues received professional, personalized portraits (Heikkilä, 2022, cited in Collins, 2026). When changing the gender in the request, the picture changed radically.

Collins (2026) traces a chain where racial representations in classical academic drawing methods (eg, the Loomis method) influence the works created by these methods, and the digitization of these works subsequently influences modern training datasets for training AI.

Midjourney studied the works of over 16,000 artists, but the names of artists, especially non-Western ones or those whose works exist primarily in physical form, were not included (Collins, 2026). The model therefore learned from an archive that was never culturally neutral. Resources for digitalization have historically been concentrated in well-funded Western institutions, meaning that digital accessibility is not a neutral technical condition, but a product of the same institutional hierarchies.

Research by Stork (2024) shows that computational analysis of historical collections can reveal such structural patterns. For example, what poses are most common, how differently men and women are depicted, which cultural preferences persist across centuries.

2.2 Cognitive Assemblages and the Illusion of Neutrality

The GIGO principle clearly demonstrates that AI results carry a cultural imprint. But the question of why it is difficult to notice is partly answered by Katherine Hayles's theory of "cognitive ensembles", which Manghani and Savage (2025) rely on.

Manghani and Savage (2025) use Hayles's idea that "cognitive ensembles" shape behavior, perception, and decision making, often unnoticed by the participants themselves, to show that generative AI is a new part of cognitive systems that have always included non-human elements. Clarke (2003, cited in Manghani & Savage, 2025) called humans “natural-born cyborgs,” organisms that regularly offload some of their mental work to the environment. In later work (Clark, 2025), he applies this specifically to generative AI, calling it an “active partner in idea generation” that co-creates meaning with humans rather than simply following commands.

This creates a false sense of neutrality. When a human artist chooses a subject, style, composition, or color, these choices are generally perceived as conscious acts of artistic decision-making and creative vision. Meanwhile, the result of AI-generated work is often perceived as less authentic. The resulting AI-generated images, as well as the vast amount of training data, only reinforce this perception. It is precisely this illusion of objectivity that is the most persistent feature of the paradox, and the most difficult to destroy.

By computer-analyzing historical collections, Stork (2024) identified certain cultural patterns. For example, how often heads are turned in portraits from different eras, which hand is dominant in self-portraits, and clear gender differences in portraits throughout the centuries. This demonstrates that computational methods can identify patterns that may be difficult for human researchers to detect . This opens the possibility of using AI to identify and analyze patterns in training data from training data.

2.3 Formulation of the mirror paradox

It can be concluded that generative AI systems act as cultural mirrors in a distinctive and potentially troubling way. They reflect current cultural hierarchies at a large scale, rather than simply copying them.

This paradox has two interconnected components, and to understand why this constitutes a paradox rather than simply a form of distortion or amplification, one can take into account these two components.

One of these components is camouflage. The more technically sophisticated a system, the less noticeable its cultural bias. The more compelling the result, the more difficult it is to see which patterns were selected and which were excluded during its development. This is how the mirror paradox differs from “algorithmic bias.”

The second component is authority. Because the actions of generative AI appear autonomous, these results have an unjustified authority. The choice of a human artist is read as a conscious choice of the creator, while the result of AI's work appears as the product of a generalized, statistical and therefore seemingly “objective” process. Accordingly, the result begins to seem natural, rather than constructed and random.

Manghani and Savage (2025) rightly point out that the problem isn't AI itself, but how it's used. The mirror paradox encourages us to understand what these tools do when they are trained on raw data, used mindlessly, and interpreting its raw output.

3. A TAXONOMY OF MIRRORING: THREE TYPES

The mirror paradox manifests through different mechanisms and with different consequences. In practice, the three types (structural, aesthetic, and economic mirroring) are closely intertwined and operate simultaneously.

3.1 Structural Mirroring: Race, Gender, and the Invisible Canon

Structural mirroring touches on the very essence of cultural production: whose bodies, styles, and perspectives are captured in the data and under what circumstances. The analysis draws on empirical data from Collins (2026) .

Collins (2026) documents this on multiple levels. In terms of gender, image generators tend to create images of women oriented towards sexual attractiveness and Eurocentric beauty standards, regardless of the query formulation. Heikkilä's experiment only confirms that the results changed not because the professional context had changed, but because of a change in gender in the query, which triggered a completely different set of statistical associations. This is influenced by the portrayal of women in commercial, digital and artistic spheres.

The racial dimension extends beyond individual images. Here, Collins (2026) traces how racial representations in classical academic drawing influence works created using these methods, and their digital versions, accordingly, influence contemporary teaching materials. Eurocentric body proportions have become the default statistical norm.

This is also reflected in the formation of the canon. The Midjourney study included works by over 16,000 artists, but excluded non-Western artists and those whose works were primarily preserved physically (Collins, 2026). The AI was trained on what was available in digital format. And this is what distinguishes the works that were given preference. Resources for digitalization have historically been concentrated in well-funded Western institutions, meaning that digital accessibility is not a neutral technical fact, but a consequence of existing institutional hierarchies.

Research by Stork (2024) adds that computational analysis of historical collections can reveal structural patterns such as which poses predominate, how men and women are depicted differently, and which cultural preferences persist over centuries.

3.2 Aesthetic Mirroring: The Tyranny of the Popular Style

Aesthetic mirroring focuses on the specific images being created. A key mechanism in optimizing generative systems for results that, in the words of developer Dr. Mike Pound (quoted in Collins, 2026), simply "look good." As a result, the system systematically favors popular, polished images at the expense of those that are formally unusual, conceptually complex, or marginal in content.

Collins (2026) tested this using Stable Diffusion, asking the system to generate images based on the work of well-known artists (Damien Hirst, Lisa Yuskavage, Anna Weyant) and one "control" artist (Collins herself, who has virtually no online presence). The system performed well on well-known names with a lot of data online but produced substantially weaker results for a little-known artist. The system's performance was directly dependent on the amount of available data, and therefore directly dependent on the artist's existing fame, the baseline that the AI had learned from. Thus, the algorithm reinforces the same recognition hierarchies because it learns only from what already exists. This creates a vicious cycle where curators,

educators, and platforms using AI content will disproportionately encounter and promote images from established artists and distribute them.

In terms of aesthetic implications, according to Collins (2026), AI images gravitate toward the style favored by the commercial world and social media. A system trained on visuals and Instagram will create images that look good on social media. It is less likely to produce a formally strange or conceptually complex image because such forms are statistically underrepresented.

Manghani and Savage (2025) develop this idea at the institutional level, criticizing the “AI spectacle.” They dismantle Refik Anadol's installation Unsupervised (2022) at MoMA, which uses the museum's collection as data to create vibrant, constantly changing projections. Art critic Ben Davis (cited in the authors) notes the problem with Anadol's treatment of art history as a collection of visual fragments to be shuffled rather than as an archive of meaningful practices imbued with social content.

Zylinska , J. is even more blunt, comparing many AI art projects to “a souped-up version of Candy Crush” (2020, cited in Manghani & Savage, 2025). Thus, by optimizing for emotional effect rather than meaning, such systems encourage viewers to perceive an image as separate from the conditions of its creation.

3.3 Economic Mirroring: Algorithmic Amplification of Market Logic

The third type of mirroring, economic mirroring, describes how generative AI systems, when operating within commercial and social media platforms, emphasizes speed and novelty at the expense of deep practice and critical thinking.

Collins (2026) traces this dynamic through the analysis that AI images often receive more positive feedback on social media than human-made products. The reason lies in the speed of production, visual consistency and compliance with what the platform algorithms promote.

Thus, AI content is displacing human art in the competition for visibility, algorithmic promotion, and economic rewards. The economic mirror redirects attention and economic rewards toward content optimized for how platforms operate, rather than toward works that push aesthetic boundaries.

Collins (2026) introduces the concept of the 'fast art fetish' to describe the cultural attraction to quickly produced, easily consumed works. Although this tendency predates generative AI, it was dramatically accelerated by AI. Collins' example of Basquiat and Warhol's collaboration is that market demands forced Basquiat to produce a simplified and racialized image of himself simply to remain visible.

Manghani and Savage (2025) add that large-scale AI art projects are often funded by technology companies and covered by corporate media. This also affects what exactly is promoted.

4. CASE STUDIES: THE MIRROR IN ACTION

4.1 Competition Scandals as Diagnostic Events

Controversies surrounding AI-generated works in competitions highlight the tension between the values of the art community and generated content. The case of the Colorado State Fair, where Allen's work "Théâtre D'Opéra Spatial" won first place in 2022, highlighted the issue of disclosure. The Mauritshuis open call in The Hague, which invited participants to submit their own interpretations of Vermeer's Girl with a Pearl Earring, is also significant because the criteria were explicit, and the result demonstrates the simultaneous operation of all three types of mirroring. Here Collins (2026) notes that the winning AI image showed a conventionally attractive white woman, close to the original, while more experimental entries (such as those with a woman of color or an older man) were ignored.

Boris Eldagsen submitted an AI-generated photograph to the Sony World Photography Awards, then withdrew the award in protest. Thus, the judging criteria of art competitions are not adapted to distinguish between the creative process and its result.

Thus, the competition criteria developed to evaluate human creative contribution are not sufficient for evaluating AI-generated works.

5. IMPLICATIONS: BREAKING THE MIRROR

5.1 For Artists: Reflexive Practice as Methodological Commitment

The mirror paradox requires that the artist engage in reflection at all stages, from the choice of tools and the formation of training data to the method of displaying and distributing the work. The artist's willingness to make this process visible, rather than hide it behind a

polished final image, determines whether such reflexivity is effective.

A good example is Anna Ridler's Mosaic Virus (2018), which is described by Manghani and Savage (2025). The work uses a dataset of tulip photographs collected and organized by the artist, making the usually invisible labor of dataset construction part of the artistic process. The usually invisible labor of collecting the dataset is made visible in the work. Ridler's solutions resist aesthetic mirroring by refusing to allow the apparent autonomy of the system to obscure the human decisions that shaped it.

Collins (2026) describes a program called Nightshade that adds subtle distortions to digital images, preventing AI systems from learning from them without the creators' consent. If the data AI systems collect without asking is systematically corrupted, the economic incentive for such data collection is reduced. It is also worth noting the contradiction with the general position of the article, as Nightshade is described not as a tool for better interaction with AI, but as a form of targeted counteraction to data collection in general.

King and McCall (2024), in their study of design thinking and art education, propose cultivating cognitive and emotional qualities that enable artists to engage with these systems critically rather than relying on them blindly. Thus, working with artistic uncertainty, mistakes and repetitions develops the skills needed for critical interaction with creative tools.

5.2 For Institutions: Reflexivity as Infrastructure

The editorial framework proposed by D'Souza and Manghani (2025) clearly outlines the task of creating guidelines, including the validation of AI cataloging, recommender systems, and visitor analytics. The museum, which has historically acquired and exhibited predominantly Western European male artists, has data that reflects precisely these preferences. An AI recommendation system trained on such data will direct visitors to already known works, and the mirror paradox concerns not only the creation of AI art, but also how institutions manage collections using AI .

The 'community archive' model described by Breakell (2025), which seeks to take into account the views of students, staff and alumni, rather than just official institutional records, is highlighted as one way to counter the structural mirroring inherent in conventional archiving. Thus, for those who form AI datasets, the implementation of critical testing is required.

Grier (2022) examines how arts education leaders Alicia Lee and Dale Schmid responded to COVID-19 and systemic racism. Their actions demonstrate how targeted responses can address systemic problems. Drawing a parallel for institutions implementing AI, combating the mirror paradox requires not only technical adjustments to systems, but also ongoing work on institutional culture.

5.3 For Researchers: The Productive Difficulty of AI's Limitations

In general, AI systems often perform poorly when dealing with style, small datasets, imaginary objects, informal rules, abstraction, and semantics. From the point of view of the mirror paradox, these are not just technical shortcomings, but indicators of the limits of what statistical learning can capture in existing cultural production.

In terms of abstraction, the AI produces images that appear abstract but may lack the conceptual intent associated with human-created abstract art. If the viewer sees meaning in what the author puts forward, then the use of AI may produce forms whose conceptual meaning or structure is less clearly defined. Viewer-oriented approaches might challenge this, but in the context of abstraction as a historically rooted practice, the gap between the procedural surface and the conceptual intent is quite real.

6. CONCLUSION

The mirror paradox offers a new perspective on the role of artificial intelligence in contemporary art. Because AI systems are trained on existing images, art archives, commercial content, and internet materials, what the system creates is based on what works were available for training and what cultural norms were already represented in them. As a result, AI can not only reproduce existing ideas about art, beauty, and authorship, but also enhance them through the scale and speed of generation.

The article identified three forms of such mirroring. Structural mirroring relates to the groups of people, artists, and artistic traditions represented in the training data. If digital archives are dominated by Western and established artists, this situation carries over into the

work of AI. Aesthetic mirroring manifests itself in the tendency of systems to create images that are already considered attractive and popular. This can lead to unusual, complex, or little-known artistic practices becoming less visible. Economic mirroring is associated with platforms and a market where speed of creation, visual appeal, and the ability to attract attention become important criteria for success.

The main problem is that these processes are not always obvious. The more qualitative and convincing the AI output becomes, the easier it is to forget the data and human decisions behind it. The image created may appear neutral and objective, although in fact it is the result of a statistical selection from an existing cultural environment. This is how the mirror paradox differs from algorithmic bias.

However, this does not mean that the use of AI in art should be rejected. On the contrary, the problem requires a more conscious approach to technology. It is important for the artist to understand what data is being used, what decisions the system makes, and what role the artist themselves plays in the process of creating the work.

Therefore, it's important to learn to recognize the cultural and economic mechanisms behind the results of AI-generated content and to monitor the data from which it learns in order to use it to create new content.

REFERENCES

1. Beris, Y., & Gulacti , I. E. (2022). Effects of fine art print artworks on the art viewer in contemporary art presentation. Journal of Graphic Engineering and Design, 13 (2), 21–28. https://doi.org/10.24867/JGED-2022-2-021
2. Breakell, S. (2024). Snail trails and alternative momentum: The Fine Art Critical Practice Archive. Journal of Visual Art Practice, 23 (2), 165–179. https://doi.org/10.1080/14702029.2024.2354086
3. Chen, R., Aghdam, M. R. G., & Khishe , M. (2025). Correction: Utilization of artificial intelligence for the automated recognition of fine arts. PLOS ONE, 20 (1), e0318114. https://doi.org/10.1371/journal.pone.0318114
4. Clark, A. (2025). Extending minds with generative AI. Nature Communications, 16 , 4627. https://doi.org/10.1038/s41467-025-59906-9
5. Collins, A. C. (2026). Garbage in, garbage out? How the monster of AI art reflects human fault , bias, and capitalism in contemporary culture. AI & Society . https://doi.org/10.1007/s00146-025-02799-5
6. D'Souza, R. E., & Manghani , S. (2025). AI & arts: People- centered practice. Journal of Visual Art Practice, 24 (4), 357–363. https://doi.org/10.1080/14702029.2025.2578870
7. Grier, S. (2022). Responding & rebuilding amidst dual pandemics: An interview with state fine arts coordinators Alysia Lee & Dale Schmid. Arts Education Policy Review, 123 (4), 236–244. https://doi.org/10.1080/10632913.2021.2011516
8. Hayles, N. K. (2025). Bacteria to AI: Human futures with our nonhuman symbionts . University of Chicago Press.
9. King, C., & McCall, M. (2024). How the fine arts create the finest students: A design thinking study. Higher Education Quarterly, 78 , 1162–1174. https://doi.org/10.1111/hequ.12521
10. Manghani , S., & Savage, T. (2025). Rigorous creativity: AI, art and electronic life. Journal of Visual Art Practice, 24 (4), 427–449. https://doi.org/10.1080/14702029.2025.2577021
11. Stork, D. G. (2024). Computer vision, ML, and AI in the study of fine art. Communications of the ACM, 67 (3), 68–75.
12. Zylinska , J. (2020). AI art: Machine visions and warped dreams . Open Humanities Press.