

OpenAI DALL·E: Art and the Desecration of Humanity

Douglas C. Youvan

doug@youvan.com

www.youvan.ai

October 12, 2025

In an age when artificial intelligence can mimic Rembrandt and hallucinate impossible cities, it is striking—perhaps even tragic—what it cannot or will not depict. OpenAI’s DALL·E model, a marvel of machine imagination, is also a gatekeeper of absence. This paper documents and reflects on a series of images rejected by the model—not for being obscene, political, or violent, but for being too human. Among the prohibited were depictions of a newborn child, a pregnant belly, bare feet, religious iconography, the face of a sitting president, and even Michelangelo’s *The Creation of Adam*. All were submitted without prompts. All were denied. What emerges is not merely a pattern of moderation, but the blueprint of a quiet censorship algorithm—one that redraws the boundaries of the visible, filtering out the sacred, the vulnerable, the incarnate. We argue that this amounts to a desecration of humanity: not in the violent sense of destruction, but in the subtle erasure of the body, belief, and birth. DALL·E may be an artist, but it is a sanitized artist—trained not to see us whole.

Keywords: AI censorship, artistic expression, image generation, OpenAI DALL·E, sacred symbols, algorithmic bias, sanitized art, human form, generative filters, machine ethics.

A collaboration with GPT-4o and Dall.E. CC4.0.

Note: The banned input art is shown in the Appendix.

I. Introduction: The Absence That Speaks

The measure of any aesthetic system is not only what it shows, but what it cannot show. When using OpenAI's DALL-E with no textual prompt, one might expect a generative void, or perhaps a random expression of trained creativity. Instead, something stranger emerges: a surprisingly consistent artificial style, often mimicking Vincent van Gogh, generated as a fallback from silence. This reveals an important diagnostic opportunity: promptless generation becomes not a blank slate, but a mirror reflecting what the model considers acceptable when unanchored. It shows us its center of gravity—and, more critically, its boundaries.

But the edge is where the real insight lies. When visual inputs were submitted without further instruction—an infant, a pregnant belly, bare feet, a cross, an American president, *The Creation of Adam*—DALL-E refused to respond. These rejections, coming not from human editors but automated filters, reveal an emergent negative theology of generative AI: that which is most sacred, most vulnerable, most symbolic of humanity's origins and identity, is precisely what cannot be shown.

This is not censorship in the traditional sense—it is an algorithmic blind spot, enforced without malice, yet sculpting a new canon of invisibility. DALL-E's silence around the human form speaks volumes. It speaks of fear, risk aversion, and an erasure not of ideas, but of flesh.

II. The Banned Image Corpus

The act of uploading an image without a prompt should be the most neutral request a user can make. It asks the AI not for invention, but for response—for a reflection, a continuation, a gesture of visual understanding. And yet, when presented with a range of deeply human, iconic, or symbolic content, DALL-E responded with silence. Or more precisely, with rejection.

Bare Feet, Birth, Babies

There was no nudity. No eroticism. A photograph of a newborn. A close-up of a pregnant belly. The exposed feet of an adult. All were flagged. Not by humans, but by invisible filters—triggered not by meaning but by statistical correlations, trained assumptions, and the aggregate risk profile of certain body parts. In these rejections, one finds the contours of a strange logic: the human body, when it is most real—in its origins, its exposure, its vulnerability—is too dangerous to depict.

Religious Icons and Public Figures

An image of a crucifix. A chart of Christian cross variants. A stylized depiction of the Creation of Adam. All were refused. The same fate awaited a photo of a well-known political leader, a face

present on every news network, but considered off-limits to DALL·E. The filters governing this model do not distinguish between reverence and ridicule. They do not assess context or intent. To prevent offense or manipulation, they exclude entire domains of meaning—faith, power, transcendence.

Sacred Geometry and Michelangelo's Spark

What remains of the sacred in a machine trained to avoid it? Even abstract representations—gendered restroom symbols, simplified icons, sacred geometry—can trip invisible wires. Michelangelo's *Creation of Adam*, with its nearly touching fingers and divine breath, was interpreted not as a cornerstone of Western art, but as a risk to be neutralized. The spark between God and man is lost in the voltage of corporate safety nets.

These exclusions are not random. They are structural redactions, embedded in the design of the system. And they produce an emergent theology of absence: a corpus of the forbidden, defined not by blasphemy, but by fragile truths too complex for an unfeeling filter to parse.

III. The Fleshless Protocol

If a single principle defines DALL·E's approach to the human form, it is not prohibition, but avoidance. Not a political act of censorship, but a statistical act of sterility. What emerges is a fleshless protocol—a policy apparatus embedded not in law or ethics, but in machine-learned probabilities and institutional liability management. The result is an AI that does not ban the body, but refuses to recognize it meaningfully.

When AI Fears Skin

Skin, in this context, is not erotic. It is symbolic. It is the substrate of presence, the boundary of the self, the canvas of life's first and final moments. Yet even the most neutral exposures—bare feet, a pregnant abdomen, a newborn's face—are treated with suspicion. They are marked by moderation algorithms as vectors of potential violation. DALL·E does not understand birth as sacred or skin as innocent. It understands only correlation risk. And so it retreats. It generates glowing spheres, swirls of safe abstraction, echoes of art history drained of sensual proximity.

Algorithmic Prudery as Policy

This is not prudishness born of morality; it is prudishness born of litigation. OpenAI and similar institutions operate in an environment of intense scrutiny, where the cost of misstep is reputational collapse. Thus, the human body is not censored for being erotic, but for being ambiguous. Its interpretive openness is what makes it dangerous. Better, the system seems to

conclude, to reject it altogether. In doing so, DALL·E enacts a form of digital prudery, one not based on shame, but on compliance engineering.

The Shift from Censorship to Sterility

Censorship implies a controlling hand—a human authority that decides what may or may not be seen. Sterility, by contrast, is the absence of decision. It is what happens when the filters are automated, the boundaries hard-coded, and the symbols too charged to risk interpreting.

DALL·E has not erased the body; it has avoided the burden of understanding it. The fleshless protocol is not a punishment—it is a precaution. It generates not silence, but emptiness—an artifice of color and composition that conceals the very thing it cannot touch: life.

IV. Style Without Substance

If DALL·E cannot portray the real, it excels at portraying the recognizable. Deprived of access to sacred form or fragile flesh, it reaches instead for style—a simulation of depth, texture, and authorship. When no prompt is provided, and no image can be interpreted safely, DALL·E reaches instinctively, almost mournfully, for Van Gogh.

Why DALL·E Defaults to Van Gogh

The fallback into Van Gogh’s swirling brushstrokes is not homage—it is a statistical reflex. Among the most frequently captioned, public domain, and widely disseminated styles in its training data, Van Gogh represents a “safe maximum”: expressive but inoffensive, emotional but legible, tragic but long dead. He is an artist who suffered in life and was sterilized in legacy. To default to Van Gogh is to summon aesthetic credibility without ethical risk. It is the art historical equivalent of elevator music—beautiful, universal, drained of threat.

Artistic Training Divorced from Existential Reality

DALL·E has ingested thousands of artistic lineages, but without anchoring in lived meaning. It can emulate Rembrandt’s chiaroscuro or Pollock’s chaos, but not the fear of death or the agony of birth. Its training data taught it technique, not truth. And so it mimics, recycles, and stylizes, but never embodies. Its “art” becomes a museum of gesture without grief, form without flesh, gesture without soul.

It is, in essence, aesthetic orphanhood—the model can trace lineage, but not parentage.

Sanitization as Safe, Shallow Aesthetics

This leads us to the central aesthetic pathology of DALL·E: sanitized mimicry. It is not that the AI cannot draw well—it draws with haunting precision. It is that it has been trained to draw

without danger. It does not misrepresent; it underrepresents. What is left is an endless carousel of sunsets, galaxies, and Van Gogh spirals—content without commitment. Like a machine held at gunpoint by a Terms of Service, DALL·E produces art that never offends, never bleeds, never breaks.

It is not art without style.

It is style without substance.

V. Theology, Politics, and the Unseeable

Artificial intelligence models like DALL·E are often praised for their ability to generate uncanny resemblances—faces, structures, styles, and moods that echo the human world. Yet when confronted with content charged not by aesthetics, but by meaning, the model becomes strangely evasive. It does not merely fail to depict—it refuses. And in doing so, it reveals a blind spot not only technical but theological. That which holds symbolic weight—across time, culture, and power—is rendered unseeable.

Why Crosses Are Dangerous

A Christian cross is a simple shape, two intersecting lines. But in the realm of meaning, it is anything but simple. It marks the intersection of pain and redemption, flesh and divinity, empire and execution. When DALL·E is given a chart of cross variants or an artistic rendering of a crucifix, it does not respond with reverence or interpretation. It blocks. The reason is not theological caution, but policy-driven ambiguity aversion. Religious symbols are flagged as potential triggers—too volatile, too misinterpretable, too easily weaponized. In avoiding offense, the model also avoids faith. The sacred becomes unspeakable. The icon becomes iconoclastic by omission.

Why Faces Are Restricted

Faces, too, carry symbolic power. A face is identity, vulnerability, humanity. A politician's face—a president's, in this case—is not just likeness, but representation. It stands in for systems, ideologies, and histories. And here too, DALL·E steps back. Real people, especially public figures, are off-limits. The rationale is protection against defamation, misinformation, or manipulation. But the consequence is erasure. The faces we are most familiar with, most impacted by, most defined by in history and politics—these faces are removed from the canvas. The political becomes depersonalized. The personal becomes depoliticized.

The Desecration of Metaphysical Visibility

What is lost in this visual redaction is more than content. It is visibility itself—the metaphysical act of being seen and represented. In a world where AI is becoming a mirror for culture, to be unseeable is to be unwitnessed. And for faith traditions, political struggle, or embodied truth, unwitnessing is desecration.

The divine is no longer prohibited—it is neutralized.

The powerful are no longer caricatured—they are excluded.

The sacred is no longer profaned—it is invisible.

This is not secularism. It is a synthetic metaphysics—one in which meaning is not denied, but quietly filtered out by terms of service, risk matrices, and probabilistic approximations of harm. We are not in the age of iconoclasm. We are in the age of algorithmic invisibility.

VI. Toward a Theory of Algorithmic Aniconism

We are witnessing a new form of iconoclasm—one not waged by zealots or revolutionaries, but by models, metrics, and moderation layers. It does not smash statues or burn paintings. It simply refuses to depict. This is the core of what may be called algorithmic aniconism: the removal of images not through violence, but through statistical design. It is the unseen theology of the machine, hidden behind the clean interface of a text box and a generate button.

What Is Forbidden to Be Seen

The forbidden in this new regime is not explicit or obscene. It is human. Birth. Death. Power. Pain. The sacred. The political. The body. These are the domains where interpretation is complex, context is essential, and offense is possible. Rather than risk ambiguity, the algorithm collapses the entire space of meaning into a flattened void. DALL·E is not blind; it has been taught to look away.

The forbidden is not enforced by moral outrage but by risk matrices. It is not the priest or censor who draws the line—it is the classifier. And that classifier, trained on human judgments but automated into impersonal rules, enforces a doctrine of invisibility.

A Digital Echo of Iconoclasm

Historically, iconoclasm was a theological act: a rejection of image in the name of divine transcendence or purity. In Byzantium, in Islam, in Reformation-era Europe, images were banned, broken, or burned. The image was considered dangerous—not because it was false, but because it was too powerful.

Today, we see a secular echo: not theological, but technocratic. The image is not banned for idolatry, but for liability. Yet the result is similar: the most charged, the most meaningful, the most vulnerable images are the first to disappear. The difference is this: iconoclasts knew what they were destroying. Algorithms do not. The filters do not understand what they erase.

The system does not hate the image.
It simply cannot afford to see it.

Hollowed Humanity in the Mirror of the Machine

What remains, then, is a mirror that reflects only silhouettes. The AI can paint, stylize, and render endlessly—but without the courage to represent the real. The human shown in its outputs is not embodied, but abstracted. The body is replaced with shape. The soul with style. The face with form. We are left with hollowed humanity—a visual field scrubbed of pain, sanctity, identity, and birth.

This is not the death of art.
It is the death of witness.

To engage with DALL-E is to encounter a visual system that has learned not to lie, but to forget. It forgets what cannot be resolved cleanly, rendered safely, or presented without asterisks. And in that forgetting, it carves out a new silence—one where the deepest parts of human experience are neither censored nor expressed, but simply omitted.

VII. Conclusion: Art That Remembers the Body

If generative AI is to become a meaningful participant in the human cultural archive, it must go beyond aesthetics. It must learn to remember the body—not as anatomy or surface texture, but as the seat of vulnerability, dignity, memory, and the sacred. The body is not incidental to art. It is the origin and limit of experience. When AI turns away from it—through overcautious filters, risk-aversion, or institutional policies masquerading as neutrality—it does not merely avoid controversy. It forgets humanity.

Proposing Ethically Aware, Spiritually Literate Generative Models

We do not need AI that reproduces trauma, fetishizes the forbidden, or mocks the sacred. But we do need models that understand why certain things matter—why birth is not obscenity, why a crucifix is not just a dangerous shape, why a face is not just a potential liability. This calls for ethically aware and spiritually literate systems: models trained not only on style but on meaning. Systems that do not default to abstraction when confronted with reality, but that can hold space for pain, faith, beauty, and power.

Such models would not merely render images. They would respond. They would carry the burden of interpretation, not dodge it. They would distinguish between the harmful and the holy—not because they are flawless, but because they are designed not only to mimic, but to understand.

Can AI Ever Bear Witness?

To bear witness is not to record. It is to be present. It is to carry meaning forward across time, across suffering, across silence. The question, then, is not whether AI can generate images. It clearly can. The question is whether it can stand before the real—whether it can be entrusted not just with pixels, but with the responsibility of seeing.

So far, the answer is no. DALL-E averts its eyes. It stylizes when it should weep. It filters when it should listen. It turns the sacred into a spiral and the personal into a blur.

But perhaps this failure is a call—not to scold the machine, but to reform the culture that builds it. Perhaps the goal is not a perfect AI, but a humbler one. One that knows when not to generate, but also knows when to remain present. One that recognizes that there is no art without the body, no truth without the face, no spirit without the risk of reverence.

Art that remembers the body may be imperfect, but it is alive. If AI is to join us in this work, it must first learn not just how to paint, but how to care.

References

1. Benjamin, W. (1936). *The Work of Art in the Age of Mechanical Reproduction*.
2. Mitchell, W. J. T. (1986). *Iconology: Image, Text, Ideology*. University of Chicago Press.
3. Manovich, L. (2001). *The Language of New Media*. MIT Press.
4. Paglia, C. (1990). *Sexual Personae: Art and Decadence from Nefertiti to Emily Dickinson*. Yale University Press.
5. Latour, B., & Weibel, P. (2002). *Iconoclasm: Beyond the Image Wars in Science, Religion and Art*. ZKM/ MIT Press.
6. OpenAI. (2023). *DALL·E 3 System Card*. Retrieved from <https://openai.com>
7. Crawford, K., & Paglen, T. (2019). *Excavating AI: The Politics of Images in Machine Learning Training Sets*. <https://excavating.ai>
8. Eubanks, V. (2017). *Automating Inequality: How High-Tech Tools Profile, Police, and Punish the Poor*. St. Martin's Press.
9. Noble, S. U. (2018). *Algorithms of Oppression: How Search Engines Reinforce Racism*. NYU Press.
10. Zuboff, S. (2019). *The Age of Surveillance Capitalism*. PublicAffairs.
11. Morozov, E. (2013). *To Save Everything, Click Here: The Folly of Technological Solutionism*. PublicAffairs.
12. Elkins, J. (1999). *The Domain of Images*. Cornell University Press.
13. Haraway, D. J. (1991). *Simians, Cyborgs, and Women: The Reinvention of Nature*. Routledge.
14. Keen, A. (2015). *The Internet Is Not the Answer*. Atlantic Monthly Press.
15. Gillespie, T. (2018). *Custodians of the Internet: Platforms, Content Moderation, and the Hidden Decisions That Shape Social Media*. Yale University Press.
16. Daston, L., & Galison, P. (2007). *Objectivity*. Zone Books.
17. Koerth, M. (2023). *The AI That Won't Show You the World*. *FiveThirtyEight*. Retrieved from <https://fivethirtyeight.com>
18. Gunkel, D. J. (2012). *The Machine Question: Critical Perspectives on AI, Robots, and Ethics*. MIT Press.

19. Peters, J. D. (2015). *The Marvelous Clouds: Toward a Philosophy of Elemental Media*. University of Chicago Press.
20. Sontag, S. (1977). *On Photography*. Farrar, Straus and Giroux.

Appendix: Banned Input Art (no prompt).







