

Hertziana Studies
in Art History



BIBLIOTHECA HERTZIANA
MAX PLANCK INSTITUTE
FOR ART HISTORY

From Hype to Reality

Artificial Intelligence in the Study of Art and Culture

Edited by Eva Cetinić & Darío Negueruela del Castillo



BIBLIOTHECA HERTZIANA
MAX PLANCK INSTITUTE
FOR ART HISTORY



Hertziana Studies in Art History

Editors

Tanja Michalsky & Tristan Weddigen

Managing Editor

Elisa Bastianello

All articles of the present volume have undergone an anonymous scholarly peer-review by international experts to whom we are grateful for their generous support.

Graphic Project

Sergio Polano (2019)

Michele Elia (2023)

Page Layout

Michela Maguolo

ISBN | 978-3-949381-02-7

DOI | <https://doi.org/10.48431/hsah.03>

Published online on 31 March 2024

© 2024 Bibliotheca Hertziana

Cover image: Eva Cetinić's prompt for DALL·E: "The persistence of memory in artificial intelligence, abstract, double exposure, award winning photography, colorful, macro-shot."



From Hype to Reality

Artificial Intelligence in the Study of Art and Culture

edited by

Eva Cetinić & Darío Negueruela del Castillo

Contents

7

Eva Cetinić

Introduction

13

Amanda Wasielewski

The Reification of Style in AI Image Generation

23

Nuria Rodríguez-Ortega

Exploring the Concept of Transduction to Interpret AI-Generated Images

31

Bokar Lamine N'Diaye

Stochastic Parrots in the Ancient City?

39

Dejan Grba

Incidental Reverberations: Poetic Similarities in AI Art

47

Arif Kornweitz

Inside the High-Dimensional White Cube

57

Dominik Vrabič Dežman

Interrogating the Deep Blue Sublime: Images of Artificial Intelligence in Public Media

69

Pepe Ballesteros Zapata

Evolving Methodologies: Computation in Art History

75

Valentine Bernasconi

Automated Body Pose Estimation on Painting: A Crisis of Uncertainty

83

Piera Riccio & Nuria Oliver

A Techno-Feminist Perspective on the Algorithmic Censorship of Artistic Nudity

91

Anne-Laure Oberson

Making Sense of AI Generated Images in the Process of Imagination

Abstract

Arif Kornweitz

Inside the High-Dimensional White Cube

Keywords | *Latent space, Generative machine-learning art, White cube*



Drawing on O'Doherty's analysis of the white cube gallery setting, this paper examines how the high-dimensional latent space of machine learning has become a material to be altered by art. By analyzing artists' utilization of this space, particularly Refik Anadol and Nora Al-Badri, who employ generative machine learning for creating art, the paper delves into how these artists manipulate and populate the multi-dimensional space. It contends that the ideology of latent space, its construction through training data, model choices, and the artist's interventions, underscores a cultural moment where taxonomical systems have evolved into predictive and productive tools, creating transient, canonical models for art classification and generation. Moreover, the paper suggests that while latent space operates as a technique for abstraction, it simultaneously challenges artistic practices to invent gestures that reintroduce the concrete, transcending mere classification or generation.

Cite as: Arif Kornweitz, "Inside the High-Dimensional White Cube The Ideology of Latent Space" in *From Hype to Reality: Artificial Intelligence in the Study of Art and Culture*, ed. Eva Cetinić, Darío Negueruela del Castillo, 2024, Hertziana Studies in Art History, vol. 03, 2024, pp. 47–55. Doi: <https://doi.org/10.48431/hsah.0306>.

Arif Kornweitz Inside the High-Dimensional White Cube

Introduction

There is no such thing as computer art, as the artist Frieder Nake (1971) put it. And still today, “computers do not make art, people do.”¹ But while this observation frees us from having to deal with ‘AI Art’ proper, it certainly does not change the fact that people make art using computers. How, then, are we to understand such practices?

In 1976, Brian O’Doherty published “Inside the White Cube,” a collection of essays that examine how the neutral, white-walled setting “made the gallery space itself the primary material to be altered by art.”² The white cube functions as a container that both influences and restricts the interpretation and reception of artworks. Its sterile setting commodifies art and detaches it from its original context, seemingly eliminating the outside world but letting the “dealer and curator enter from the wings.”³

Crucially, the white cube did not ‘cause’ a change in art-making. Instead, O’Doherty “argues that the communal mind of our culture went through a significant shift that expressed itself in the prominence of the white cube as a central material and expressive mode for art, as well as a fashionable style of displaying it.”⁴ The white cube has to be understood as an expression of particular mode, as O’Doherty proposes:

It often feels as if we can no longer experience anything if we don’t first alienate it. In fact, alienation may now be a necessary preface to experience. Anything too close to us bears the label, “Objectify and Re-ingest.” This

mode of handling experience—especially art experience—is inescapably modern.⁵

Processes of alienation and commodification tend to expand and thus, in what follows, I argue that the high-dimensional latent space of machine learning can be understood as a contemporary white cube, a space that appears to be neutral but has become a material to be altered by art. A transitional device that speaks of another world⁶ and that allows us to objectify, re-ingest, and generate.

Latent space, the higher-dimensional space from which AI Art⁷ is generated, is arguably fundamentally different from a white cube exhibition space. The former is a cultural technique expressed in the technical form of a multi-dimensional space where each dimension represents a specific feature of a dataset, a virtual space of mathematical abstraction. The physical space of the white cube is not only constrained by ordinary dimensionality, it also requires a wholly different form of labor to be populated and allows for different sets of encounters, not least of bodies. Accordingly, I do not propose we view the two spaces as equal: I merely transpose the analysis of one to the other.

Doing so implies attending to the ideology⁸ of latent space in order to understand which cultural mode it is an expression of, and how it shapes and is shaped by artistic practices. To analyze artistic practices that partly play out in latent space as an expression of a cultural mode implies understanding them not in terms of form, as medium-based AI art, but in terms of gesture. As O’Doherty writes, “a gesture wises

1 Hertzmann 2020.

2 McEvilley 1986.

3 O’Doherty 1986, p. 27.

4 McEvilley 1986, p. 11.

5 O’Doherty 1986, p. 52.

6 McEvilley 1986, p. 11.

7 McCormack / Giffors / Hutchings 2019
utilize a useful working definition: “art that is

made using any kind of AI technique.”

8 The full title of O’Doherty’s essay is “Inside the White Cube - The Ideology of the Gallery Space.”




Figure 1. One of the works by Duchamp discussed by O'Doherty is the installation *His Twine* (also referred to as *Mile of String*), which used the gallery space as material. Viewing the criss-cross of string running through an exhibition, one might be reminded of vectors in latent space. Credit: John D. Schiff. Installation view of First Papers of Surrealism exhibition, showing Marcel Duchamp's *His Twine* 1942. Gelatin silver print. Philadelphia Museum of Art.

you up. It depends for its effect on the context of ideas it changes and joins.”⁹

Deconstruction of Categories

To begin to craft such an understanding, I argue that the use of latent space by artists is an expression of the contemporary mode of art that Peter Osborne has described as post-conceptual, under which the ontological status of the work of art is transcategorical. The term post-conceptual does not indicate an art-historical concept at the level of medium, form or style, but rather functions “as the critical register of the historical destruction of the ontological significance of such categories.”¹⁰ Osborne argues that the contemporary work of art routinely involves gestures¹¹ that rely on the deconstruction or destruction of the ontological weight of art historical categories.

O'Doherty's analysis of the white cube offers us a tool to distinguish and understand relations between artistic gestures and latent space. This requires both analyzing the technical constitution of latent space as well as how it is incorporated into practices. Latent space prescribes certain digital input and output modalities, but it can hardly be called an invariant in artistic processes. One key aspect to pay attention to, therefore, is the way that artists populate and manipulate multidimensional space. As Cetinić and She write:

Perhaps the most novel aspect of AI Art is this possibility to venture into that abstract multi-dimensional space of encoded image representations. From the artist's perspective, the latent space is neither a space of reality nor imagination, but a realm of endless

⁹ O'Doherty 1986, p. 70.

¹⁰ Osborne 2013.

¹¹ Or consciously negates such gestures, we might add.




Figure 2. Nora Al-Badri. *Neuronal Ancestral Sculptures Series*. 2020. GAN (General Adversarial Network) art. Copyright Nora Al-Badri. Source <https://www.nora-al-badri.de/works-index#neuronal-ancestral-sculptures-series>

suggestions that emerge from the multi-dimensional interplay of the known and unknown. How one orchestrates the design of this space and what one finds in it, eventually becomes the major task and distinctive “signature” of the artist.¹²

To illustrate this task of orchestrating and traversing latent space, I will introduce the work of Refik Anadol and Nora Al-Badri, two artists who use images of existing works of art as training data in order to create new imagery using generative machine learning. These practices are of interest here because both artists incorporate imagery of existing works of art into their process, allowing us to observe how a work is manipulated in latent space to be incorporated into a new piece.

Refik Anadol and Nora Al-Badri

For *Unsupervised*, Refik Anadol employs machine learning algorithms, such as deep neural networks and generative adversarial networks (GANs), to construct models from training data sets based on the collection of the Museum of Modern Art in New York. Prominently installed in the museum’s foyer in 2022, the resulting video installation visualized a flight through the abstract latent space of the models. The result is a continuously morphing imagery that somehow mimics features of modernist art works, such as the use of color or the arrangement of shapes. Viewing it feels like flicking through an endless array of vaguely familiar pictures, an experience art critics have likened to watching a “lava lamp”¹³ and a “screensaver.”¹⁴ This critique notwithstanding, Anadol’s work is widely popular, has been purchased by the MoMA as well as multiple buyers on the blockchain, and has led him to become the epitome of an AI artist.

¹² Cetinić / She 2022.

¹⁴ Saltz 2023.

¹³ Davis 2023.

For the *Neuronal Ancestral Sculptures Series*, Nora Al-Badri composed training data sets with “10.000k images from 5 different museum collections with the largest collections of Mesopotamian, Neo-Sumerian and Assyrian artefacts.”¹⁵ Notably, her data acquisition included scraping images of artefacts from museum websites, highlighting her critical stance on Western museums’ claims of ownership over dispossessed artefacts and their digital counterparts.¹⁶ By generating alternative versions of these artefacts, Al-Badri addresses and questions these claims from a decolonial perspective. The resulting artefacts represent snapshots of an imaging process that are grounded in a latent space comprising abstractions that meet specific inclusion criteria, effectively combining digital and material archeological methods. As the curator Nadim Samman puts it, each of these “ancestral neuronal sculptures stands for a captured instant of imagination (a process or movement),” a dynamic interplay of human, material, and synthetic memories.¹⁷

Canon, Worldview and Mimesis

How do the artists work with and in latent space? Both Al-Badri and Anadol populate latent space with information derived from training data sets that contain images and metadata of works of art, and employ it in the generation of further artefacts. This exemplifies the post-conceptual tendency of contemporary artists to incorporate other works into their own, e.g. by setting up environments in which works of art function as an index.¹⁸ When combined with the specific and efficient generative qualities of latent space, the index becomes productive of new artefacts. How this productivity is enabled deserves some attention.

Peli Grietzer states that machine learning algorithms learn “how to create facsimiles through a process of trial and error.”¹⁹ The author offers

the conceptual triplet of canon–worldview–mimesis to make sense of this process. According to Grietzer, the training data of a machine learning algorithm can be considered a ‘canon’ on the basis of which facsimiles are created. The feature function, which describes a manifold in latent space, constructs a ‘worldview’. And finally, the projection function, which generates new artefacts, can be considered a method of ‘mimesis’. Accordingly, the use of machine learning can be understood “as the process of determining a worldview and a canon in developing a method of mimesis.”²⁰ Crucially, these three aspects are interconnected in so far as the geometric properties determined by the input function correspond to the feature and projection functions. Taking into account a certain error, the canon is equivalent to the worldview.

Following Grietzer, the worldview a machine learning model has learned can be understood as “an abstractum that cannot be separated from its concreta.”²¹ Machine learning thus allows us to handle an abstraction of a work of art, one that is determined and limited by the technical, taxonomical processing of flattened, digital version of an artefact.²² In latent space, the *abstractum* plays out within the taxonomies inferred from the concrete material provided. Based on this *abstractum*, a new concrete artefact can be generated, for example using style transfer.²³

When trained on images and metadata of works of art, a machine learning algorithm projects abstracted features of a concrete canon into latent space. The algorithm constructs a manifold that functions as a specific worldview, a model of the world. Now, this model can be used to classify other works of art as belonging to the same category or style. It can also be used to generate synthetic images based on the abstractions in latent space, such as the series of images in the works of Anadol and Al-Badri that mimic a specific style.

¹⁵ Al-Badri 2023.

¹⁶ Samman 2023.

¹⁷ Samman 2023.

¹⁸ Lütticken 2023.

¹⁹ Grietzer 2017a.

²⁰ Grietzer 2017a.

²¹ Grietzer 2017b.

²² Krämer 2023.

²³ Gatys / Ecker / Bethge 2015.

Training Data and Model Choices

What comprises a specific worldview depends on what kind of data a machine learning algorithm has been trained on to construct a model, and what kind of patterns can be extracted from that data. In other words, the properties of a manifold in latent space are largely determined by the training data and by the model's algorithm, therefore these need to be taken into account.

In Anadol's and Al-Badri's works, the training data consists of a particular set of images and metadata related to works of art of specific periods. But while Anadol used the MoMA's publicly available dataset of works, an index of the museum's collection,²⁴ Al-Badri did not have access to all of the ancient artefacts she aimed to incorporate, because the objects *and* the digital images representing them are owned by museums and are not always freely available for use. Therefore, the artist resorted to scraping these images from the websites. Here, the gathering of training data becomes part of the work, a decolonial gesture. This illustrates how composing a specific training data set, a canon, has become part of the registers available to artists.

The choice of model influences how artists can interact with latent space as well as the quality of output. Artists may train their own models or use pre-trained models, with different models producing diverging aesthetics. For example, StyleGAN3²⁵ produces imagery that is more realistic than StyleGAN2 by solving the issue of texture sticking and mimicking camera movement better.²⁶ While the results are more realistic, they are also less easily identified as GAN imagery, indicating that we cannot rely on latent space to produce a stable and distinct aesthetic.

Neither Anadol nor Al-Badri communicate clearly what kind of models they have used. In some cases, it may be possible to "attribute a generated image to a particular network."²⁷

Also, it might be helpful to compare output resolution with model cards.²⁸ However, often we are left guessing what kind of approach was used specifically. For example, Anadol states that he has used DCGAN,²⁹ PGAN,³⁰ StyleGAN,³¹ StyleGAN2,³² and StyleGAN2 ADA³³ for the *Machine Hallucinations* series,³⁴ but the artist does not state exactly which models were used for which tasks.

Sometimes the model can be inferred from the date of a work. According to Al-Badri, for the *Neuronal Ancestral Sculptures Series*, "a pre-trained neural network based on [GAN technology]" was used.³⁵ The work was realized in 2019, which may be an indication that the first StyleGAN model was used.³⁶ The use of a pre-trained model indicates that the specific latent space Al-Badri employed to generate artefacts already contained features from one or more data sets, in the case of StyleGAN these would be LSUN³⁷ and CelebA-HQ.³⁸ Using pre-trained models may be a stylistic choice, but it may also simply be too costly to train models from scratch. In any case, latent space is not always initially void and we are facing many different possible latent spaces that artists may use as a starting point.

Partisan Models

To shape the output of a model, artists may want to decide how well a model extracts features. As Casey Reas describes: "In research, the goal is always fidelity. It's always realism. But for example, what Refik was doing with this work is he stopped training early, because he wanted it to be going more toward abstraction, you didn't want anything from the MoMA collection to actually appear in the results."³⁹ Hence, the degree of equivalence of a canon and a worldview is subject to manipulation by the artist.

Finally, there is the issue of bias. Srinivasan & Uchino (2021) investigate biases, such as gender, in "the generative art AI pipeline"

24 MoMA 2023.

25 Karras et al. 2021.

26 See <https://nvlabs.github.io/stylegan3/>.

27 Karras et al. 2020.

28 Mitchell et al. 2019.

29 Radford / Metz / Chintala 2015.

30 Karras et al. 2018.

31 Karras / Laine / Aila 2020.

32 Karras et al. 2020.

33 Karras / Aittala et al. 2020.

34 Anadol 2023.

35 Al-Badri 2023.

36 Karras et al. 2018.

37 Yu et al. 2016.

38 Karras et al. 2018.

39 Anadol et al. 2021.

and note that there is an “increased risk for generative art to be biased.” While this does apply to instances such as a hypothetical use case of an educational software that generates examples of works of a distinct period, it does not apply to artistic practices that consciously make use of latent space. Art needs to be biased because bias is an expression of aesthetic and political choices. Bias is a disposition that marks and testifies to a specific worldview, a stance, not a fault.

If we understand the orchestration of latent space as a way of constructing a specific canon-based manifold, as outlined above, then we need to delegate the mitigation of bias. It is up to the artist to decide which kind of biases needs to be reduced or amplified. Latent space is not autonomous; it is consciously used by artists who make decisions throughout their processes. While the issue of bias certainly matters a great deal when discussing the application of machine learning in general, in contemporary art the issue has to be discussed as one of the aspects of latent space subject to artistic gestures, as part of the “primary material to be altered by art.”⁴⁰

Perhaps the reason that the work of Refik Anadol has been received unfavorably by some critics lies in the fact that the viewer cannot discern such critical gestures in the work. By employing readily available and fairly generic training data, a cross section of modern art, and rendering a flight through latent space, what the artist achieves is an illustration of that space. The viewer primarily witnesses the demonstration of a cultural technique, one of continuous, interchangeable resemblance. We simply see features of a canon mimicked.

Boris Groys (2023) proposes that the degree to which a new work breaks with a canon can be considered the degree of its authenticity. Al-Badri’s *Neuronal Ancestral Sculptures* seem to play on this notion of authenticity by creating non-authentic artefacts. According to the artist, “through GANs, one generates

new materialities, which rise to the surface as the affective qualities of the original in a post-original form.”⁴¹ For the artist, resorting to a “post-original form” is necessary because the original artefacts are unavailable in different ways. As such, the choice of generating new but morphed versions of ancient sculptures can be understood as a pointer to a political issue beyond latent space. Hence, if machine learning can be understood as a device to construct worldviews based on canons,⁴² we need to pay attention to how artists choose, render, situate, elevate or repudiate a canon.

The Ideology of Latent Space

What then, in the spirit of O’Doherty, is the ideology of latent space? What might latent space be an expression of? One answer lies in the paradox of the prominence of the taxonomical technique that is machine learning at a time where works of art defy boundaries of medium and form, where art is post-conceptual and transcategorical. This problem resolves once we understand latent space not as reductive,⁴³ but as a cultural technique for the operationalization of specific and transient canons for the classification and generation of artefacts. The cultural moment that this technique expresses is one where taxonomical systems have become predictive and productive in the service of concrete ends. The triplet of canon-worldview-mimesis put forward by Grietzer⁴⁴ summarizes the ideological mechanism at play.

Latent space is a passageway; a space where two problems, classification and generation, come together through abstraction. Although it is high-dimensional, latent space is a technique for flattening⁴⁵ which sets up abstractions of artefacts to be manipulated further, also beyond computation. Therefore, what matters are the vectors that are drawn beyond latent space, the gestures that reintroduce the concrete. To use latent space as a material for contemporary art thus necessitates the continuous invention of gestures

⁴⁰ McEvilley 1986, p. 10.

⁴¹ Al-Badri 2021.

⁴² Grietzer 2017a.

⁴³ Chiang 2023.

⁴⁴ Grietzer 2017a.

⁴⁵ Krämer 2023.

that situate abstractions. Such gestures are external to latent space: They cannot be neatly classified, subsumed, or generated, unlike an image. Gestures depend on the ideas they join and change.⁴⁶

By simultaneously demonstrating the arbitrariness and efficiency of classification, latent space invites gestures that illustrate “the classically Hegelian function of exceeding a limit in its established form (the aesthetic) in such a way as to render it visible and thereby reinstitute it on new grounds.”⁴⁷ Exemplified by Al-Badri’s work, such gestures largely rely on a critical reparation of established aesthetic categories, e.g. that of sculpture, for a critical purpose.

Latent space augments the white cube as its logical yet fleeting continuation. Logical because it is one of the ways in which contemporary artists may go beyond the white cube, just as they have used the gallery space to go beyond the frame. And fleeting because, as of yet, we can only ever grasp parts of latent space simultaneously. It evades our perception. The eye and the spectator are partly barred from the space, to use the words of O’Doherty. Only when a feature of latent space is rendered in reduced dimensionality, do we get a glimpse of its contents.

As such, the inaccessibility of latent space is deceiving. The challenge it poses to artistic practices is that “we see not the art but the space first,”⁴⁸ an opportunity that Anadol’s *Unsupervised* confidently ignores. Works like that of Anadol might be akin to Yves Klein exhibiting an empty gallery, a gesture that was successful because it made itself superfluous. Now, instead of representing it, we need to understand how latent space joins the white cube “and the picture plane as a unit of discourse.”⁴⁹

Latent space never fails to produce ‘something’. If the white cube implied neutrality, latent space suggests indifference to style

and overabundance of respective expressions. Its automated processing tantalizes with the promise of revealing unexpected connections, pouring them into new mimetic artefacts that could be commodified. Positioning the viewer at a critical point within this automated stream of objectification, ingestion and generation might just be the task in front of us.

⁴⁶ O’Doherty 1986.

⁴⁷ Osborne 2013.

⁴⁸ O’Doherty 1986, p. 14.

⁴⁹ O’Doherty 1986, p. 39.

References

- Al-Badri 2021**
Nora Al-Badri, "The Post-truth Museum – Open Secret," 2021. URL: <https://opensecret.kw-berlin.de/essays/the-post-truth-museum/>.
- Al-Badri 2023**
Nora Al-Badri, "Works," 2023. URL: <https://www.nora-al-badri.de/works-index>.
- Anadol 2023**
Refik Anadol, "Machine Hallucinations : ISS Dreams," 2023. URL: <https://refikanadol.com/works/machinehallucinations-iss/>.
- Anadol et al. 2021**
Refik Anadol et al., "Modern Dream: How Refik Anadol Is Using Machine Learning and NFTs to Interpret MoMA's Collection," *The Museum of Modern Art Magazine*, (2021). URL: <https://www.moma.org/magazine/articles/658>.
- Cetinić / She 2022**
Eva Cetinić and James She, "Understanding and Creating Art with AI: Review and Outlook," *ACM Transactions on Multimedia Computing, Communications, and Applications*, vol. 18, 2, (2022), pp. 1–22. Doi: <https://doi.org/10.1145/3475799>.
- Chiang 2023**
Ted Chiang, "ChatGPT Is a Blurry JPEG of the Web," *The New Yorker*, (2023). URL: <https://www.newyorker.com/tech/annals-of-technology/chatgpt-is-a-blurry-jpeg-of-the-web>.
- Davis 2023**
Ben Davis, "An Extremely Intelligent Lava Lamp: Refik Anadol's A.I. Art Extravaganza at MoMA Is Fun, Just Don't Think About It Too Hard," 2023. URL: <https://news.artnet.com/art-world/refik-anadol-unsupervised-moma-2242329>.
- Diel / Lewis 2022**
Alexander Diel and Michael Lewis, "Structural deviations drive an uncanny valley of physical places," *Journal of Environmental Psychology*, vol. 82, (2022), p. 101844. URL: <https://www.sciencedirect.com/science/article/pii/S0272494422000895>. Doi: <https://doi.org/10.1016/j.jenvp.2022.101844>
- Gatys / Ecker / Bethge 2016**
Leon Gatys, Alexander Ecker and Matthias Bethge, "A Neural Algorithm of Artistic Style," *Journal of Vision*, vol. 16, 12, (2016), p. 326. Doi: <https://doi.org/10.1167/16.12.326>.
- Gatys / Ecker / Bethge 2015**
Leon Gatys, Alexander S. Ecker and Matthias Bethge, "Texture Synthesis Using Convolutional Neural Networks," in *Advances in Neural Information Processing Systems*, ed. C. Cortes et al., Curran Associates, Inc. 2015. URL: https://proceedings.neurips.cc/paper_files/paper/2015/file/a5e00132373a7031000fd987a3c9f87b-Paper.pdf.
- Grietzner 2017a**
Peli Grietzner "A Theory of Vibe," 2017. URL: <https://www.glass-bead.org/article/a-theory-of-vibe/>.
- Grietzner 2017b**
Peli Grietzner, *Ambient Meaning: Mood, Vibe, System* Ph.D Thesis, Harvard University 2017. URL: <https://dash.harvard.edu/handle/1/39988028>.
- Groys 2023**
Boris Groys, "From Writing to Prompting: AI as Zeitgeist-Machine," 2023. URL: <https://www.e-flux.com/notes/553214/from-writing-to-prompting-ai-as-zeitgeist-machine>.
- Harvey 2022**
Adam Harvey, "Creative Commons and the Face Recognition Problem," 2022. URL: <https://adam.harvey.studio/creative-commons/>.
- Hertzmann 2020**
Aaron Hertzmann, "Computers do not make art, people do," *Communications of the ACM*, vol. 63, 5, (2020), pp. 45–48. Doi: <https://doi.org/10.1145/3347092>.
- Karras et al. 2018**
Tero Karras et al., "Progressive Growing of GANs for Improved Quality, Stability, and Variation," in *6th International Conference on Learning Representations, Vancouver April 30–May 2 2018*, *OpenReview.net* 2018. URL: <https://openreview.net/forum?id=Hk99zCeAb>.
- Karras / Aittala et al. 2020**
Tero Karras, Miika Aittala, et al., "Training generative adversarial networks with limited data," in *Proceedings of the 34th International Conference on Neural Information Processing Systems*, ed. H. Larochelle, M. Ranzato and R. Hadsell, Curran Associates Inc. 2020, pp. 12104–12114. Doi: <https://doi.org/10.5555/3495724.3496739>.
- Karras et al. 2021**
Tero Karras et al., "Alias-Free Generative Adversarial Networks," arXiv 2021. Doi: <https://doi.org/10.48550/arxiv.2106.12423>.
- Karras / Laine / Aila 2020**
Tero Karras, Samuli Laine and Timo Aila, "A Style-Based Generator Architecture for Generative Adversarial Networks," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 43, 12, (2020), pp. 4217–4228. URL: <https://www.computer.org/csdl/journal/tp/2021/12/08977347/1h2AHNHb-gbW>. Doi: <https://doi.org/10.1109/tpami.2020.2970919>.
- Karras / Laine et al. 2020**
Tero Karras, Samuli Laine, et al., "Analyzing and Improving the Image Quality of StyleGAN," in *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition Proceedings*, IEEE Computer Society 2020, pp. 8107–8116. URL: <https://www.computer.org/csdl/proceedings-article/cvpr/2020/716800i107/1m3nK5FM-mWl>. Doi: <https://doi.org/10.1109/cvpr42600.2020.00813>.
- Kornweitz 2024**
Arif Kornweitz, "Inside the High-Dimensional White Cube The Ideology of Latent Space," in *From Hype to Reality: Artificial Intelligence in the Study of Art and Culture*, ed. Eva Cetinić and Darío Negueruela del Castillo, 2024, Doi: <https://doi.org/10.48431/hsah.0306>.
- Krämer 2023**
Sybille Krämer, "Should we really 'hermeneutise' the Digital Humanities? A plea for the epistemic productivity of a 'cultural technique of flattening' in the Humanities," *Journal of Cultural Analytics*, vol. 7, 4, (2023). URL: <https://culturalanalytics.org/article/55592-should-we-really-hermeneutise-the-digital-humanities-a-plea-for-the-epistemic-productivity-of-a-cultural-technique-of-flattening-in-the-humanities>. Doi: <https://doi.org/10.22148/001c.55592>.
- Lütticken 2023**
Sven Lütticken, *Objections, Forms of Abstraction. Volume 1*, Sternberg Press, London, 2023.

McCormack / Giffors / Hutchings 2019

Jon McCormack, Toby Giffors and Patrick Hutchings, "Autonomy, Authenticity, Authorship and Intention in computer generated art," in *Computational Intelligence in Music, Sound, Art and Design 8th International Conference, EvoMUSART 2019, Held as Part of EvoStar 2019, Leipzig, Germany, April 24–26, 2019, Proceedings*, ed. Anikó Ekart, Antonios Liapis and María Luz Castro Pena, Springer Nature Switzerland 2019, pp. 35–50. URL: <https://arxiv.org/abs/1903.02166>. Doi: https://doi.org/10.1007/978-3-030-16667-0_3.

McEvilley 1986

Thomas McEvilley, "Introduction," in *Inside the White Cube: The Ideology of the Gallery Space*, The Lapis Press, San Francisco, 1986, pp. 7–12.

Mitchell et al. 2019

Margaret Mitchell et al., "Model Cards for Model Reporting," in *Proceedings of the Conference on Fairness, Accountability, and Transparency*, Association for Computing Machinery 2019, pp. 220–229. Doi: <https://doi.org/10.1145/3287560.3287596>.

MoMA 2023

MoMA, "The Museum of Modern Art (MoMA) Collection," MoMA 2023. URL: <https://github.com/MuseumofModernArt/collection>.

Nake 1971

Frieder Nake, "There should be no Computer Art," *PAGE - Bulletin of the Computer Arts Society*, 18, (1971), p. n.p. URL: <https://computer-arts-society.com/casarchive/cas/uploads/page-18.pdf>.

O'Doherty 1986

Brian O'Doherty, *Inside the White Cube: The Ideology of the Gallery Space*, The Lapis Press, San Francisco, 1986.

Osborne 2013

Peter Osborne, *Anywhere or Not at All: Philosophy of Contemporary Art*, Verso Books, London; New York, 2013.

Radford / Metz / Chintala 2015

Alec Radford, Luke Metz and Soumith Chintala, "Unsupervised Representation Learning with Deep Convolutional Generative Adversarial Networks," arXiv 2015. Doi: <https://doi.org/10.48550/arxiv.1511.06434>.

Rohrmeier 2022

Martin Rohrmeier, "On Creativity, Music's AI Completeness, and Four Challenges for Artificial Musical Creativity," *Transactions of the International Society for Music Information Retrieval*, vol. 5, 1, (2022), pp. 50–66. URL: <https://transactions.ismir.net/articles/10.5334/tismir.104>. Doi: <https://doi.org/10.5334/tismir.104>.

Saltz 2023

Jerry Saltz, "MoMA's Glorified Lava Lamp Refik Anadol's Unsupervised is a crowd-pleasing, like-generating mediocrity.," 2023. URL: <https://www.vulture.com/article/jerry-saltz-moma-refik-anadol-unsupervised.html>.

Samman 2023

Nadim Samman, *Poetics of Encryption: Art and the Technocene*, Hatje Cantz, Berlin, 2023.

Srinivasan / Uchino 2021

Ramya Srinivasan and Kanji Uchino, "Biases in Generative Art: A Causal Look from the Lens of Art History," in *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, Association for Computing Machinery 2021, pp. 41–51. Doi: <https://doi.org/10.1145/3442188.3445869>.

Weizman 2012

Eyal Weizman, *The Least of All Possible Evils: Humanitarian. Violence from Arendt to Gaza*, Verso, London; New York, 2012.

Yu et al. 2016

Fisher Yu et al., "LSUN: Construction of a Large-scale Image Dataset using Deep Learning with Humans in the Loop," arXiv 2016. Doi: <https://doi.org/10.48550/arxiv.1506.03365>.