

# THE BAFFLER

Harrison, Will, "Another Bullshit Night in Slop City," *The Baffler*, June 11, 2025



© Jamison Lung

*Medium Hot: Images in the Age of Heat* by Hito Steyerl. Verso, 192 pages. 2025.

EVERY DAY THERE ARE MORE OF THEM. Every day the horde of bots grows larger, and every day they relentlessly scrape the world wide web, gobbling up terabytes of data to train the insatiable artificial intelligence models that have already begun to govern our way of life, whether they are ready to or not.

They've been watching us for some time now, the machines. Back in September 2019, before the Ghiblification of every image under the sun, before everyone and their mother started volleying around the word *agentiv*, before Chat-GPT started melting down about the immaculate conception and helping Ivy League undergrads sail through their classes, there was *ImageNet Roulette*. Part of an exhibition called "Training Humans" organized by researcher Kate Crawford and artist Trevor Paglen, *ImageNet Roulette* was an attempt to better understand the mechanics behind the so-called training sets used for

machine learning, namely the canonical dataset ImageNet, which contains more than 14 million photographs that have been scraped and then labeled and then divided into more than 20,000 categories by a legion of underpaid human data workers. *ImageNet Roulette* allowed users to upload pictures of themselves to a website where an AI model trained on ImageNet could then label what it saw. The results were revelatory, prescient, and very, very troubling. People loved it. At the height of its popularity—a few days before being taken offline—the site spit out as many as one hundred thousand labels an hour. Everyone’s face was bordered by a lime green box headed by a short stack of terms. Some entries were harmless and banal: “assistant professor,” “associate professor,” “newsreader,” “cheerleaders,” “Boy Scouts”; others were more colorful and a bit odd: “swinger,” “snob,” “slob”; and many reflected the overt prejudices that flesh-and-blood image labelers had imbued into the training set: “slattern,” “slut,” “Jihadist.” On this front, Crawford and Paglen’s project was a smash success: it exposed the unchecked xenophobia that machine learning had been built upon and prompted the researchers behind ImageNet to scrub more than half of the 1.2 million photographs in the dataset’s “people” category.

**Slop is the sinister child  
of scandal and banality,  
it is Pinterest board  
Cronenberg, it is the  
exact midpoint of a bell  
curve.**

But sometimes the truth is not enough, or merely gets lost in the torrent of information, in the never-ending game of telephone. “There are models on top of models, and training sets on top of training sets,” as Christo Buschek and Jer Thorp put it in their 2024 essay “Models All the Way Down.” As time hurtles by, and countless tech companies race to launch and tailor their products so as to make them as habitual and indispensable as possible, the “omissions and biases and blind spots from these stacked-up models and training sets shape all of the resulting new models and new training sets.” By the time

Crawford and Paglen’s work prompted ImageNet’s Stanford-based research team to conduct their purge, it had undoubtedly been injected into the algorithm already, used to fuel countless neural networks.

Six years on, this episode feels rather quaint. LAION-5B, one of the preeminent training sets of the moment, contains 5.85 billion images to ImageNet’s 14 million. Allegedly built for noncommercial research purposes, LAION-5B quickly proved irresistible in its enormity, and was used to train popular image generators such as Stable Diffusion and Midjourney. It was also assembled in purely automated fashion, without the intervention of clumsy, rushed, underpaid, and biased humans. It was instead infused with the bias of capital, pulling over 155 million images enriched with metadata from Pinterest and 140 million from Shopify (as well as over 1,000 images classified as Child Sexual Abuse Material, which led to the dataset being made temporarily unavailable for download in December 2023). “Neural networks,” which transport raw data through a dense system of interconnected nodes to create supposedly optimized outputs, “mimic a market logic, in which reality is permanently at auction,” as Hito Steyerl puts it in her essay “Mean Images,” which was originally published in *New Left Review* in 2023 and has now been included in her new essay collection *Medium Hot: Images in the Age of Heat*.

Steyerl, a Berlin-based filmmaker and essayist, is no stranger to the digital landscape or the artworld zeitgeist. Her always-expanding, often confrontational, conceptually omnivorous body of work has earned her a position in the top twenty of *ArtReview*’s Power 100 rankings every year since 2015, and *Medium Hot* unsurprisingly extends its web in all directions, forming a vast, teeming network of themes: the veracity of the image in the age of deepfakes and visual brain rot; the connection between surveillance, facial recognition, phrenology, and statistics; the literal heat that machine generated images

release into the atmosphere; the intentionally secretive data centers that “drain the planet’s energy in order to create a stable thermal environment—not for people but for information”; the compulsive, ongoing race to produce artificial general intelligence, or AGI; the tech monopolies that hijack “access to public utilities like information and communication” and sell people “back the products wrested from their own expropriation”; the unseen human labor behind every automated action; and the way that art or aesthetic novelty is used as a cover for all of this.

“Mean Images,” which acts as the backbone of the collection, investigates the “hallucinated mediocrity” of machine generated visuals, which are seemingly spontaneous and yet bound to numerical probability, materializing as the averaged-out reflections of colossal and flawed datasets like LAION-5B. They come to us secondhand and are composed of millions, if not billions, of other images that have been stolen and crushed up and centrifugally swirled together. These derivative excretions—what we would colloquially call “slop”—“are documentary expressions of society’s views of itself” but also cede contact with reality, eliminating any “possibility of ever capturing something which is not yet already known.” They are “mean” in every sense: malevolent, meager, mediated, miserly, minor. “They converge around the average, the median,” circumventing observation, epiphany, and anything organic. Even the hyperrealistic graphics produced by Google’s newest video generation model, Veo 3, are nothing more than uninventive, letter-for-letter acts of mimetic reproduction. We have moved from the decisive moment of the photograph to on-demand jpegs of cats in diapers.

Images have always been somewhat mystifying intermediaries between the world and humans, even before this especially artificial era, constantly shirking objectivity in favor of deception, illusion, and outright fantasy. But how has this mediation changed or advanced now that so many images are *themselves* mediated by generative AI? “Human beings forget they created the images in order to orientate themselves in the world,” writes theorist Vilém Flusser in his landmark 1984 text *Towards a Philosophy of Photography*. Since they are no longer able to decode them, their lives become a function of their own images: Imagination has turned into hallucination.” Amid the exponential rollout of petrol-guzzling AI infrastructure—an era that will likely be studied as the exact moment humanity gave up on carbon neutrality and drove 140 miles per hour off a cliff—the second-rate psychedelia that has been hallucinated by models is being used as the missing rationale for all the hype around machine learning. Artists have been asked to help us “orientate” ourselves by shilling for the “tools” of so-called artistic democratization that seek to make them redundant.

If art is being used as a Trojan horse of sorts—as a means of onboarding customers into an extractive technological environment—then an aesthetic decoding of machine-generated images becomes imperative. In “Mean Images,” Steyerl writes about finding photographs of herself inside LAION-5B before asking Stable Diffusion “to render ‘an image of hito steyerl.’” The resulting portrait—made through a mysterious process that involves drowning training data in noise before reversing course and removing sound completely—is indeed quite “mean”: Steyerl’s mechanically reproduced face appears gaunt and creased, and possesses a texture more reminiscent of silicone than flesh. It looks as if a simple turn of her digitized head might elicit a smear of checkerboarded pixels. A few pages later, Steyerl reflects on her inclusion in MS-Celeb-1M, a Microsoft-owned dataset that was used to train yet another dataset called Racial Faces in the Wild, which was itself used to “reduce bias in facial-recognition software,” a ridiculous and ill-fated idea that feels like an AI model’s impression of liberal equality, a DEI-ification of inherently racist police state surveillance (which is exactly what the research was used to enhance). “By now,” Steyerl notes, “the majority of the faces that have appeared on the internet have probably been included in such operations.”

Indeed, the conveyor belt of progress is moving so rapidly that nothing can ever be fixed or taken back or even reconsidered, and our sudden dependence on machine learning has already begun to induce a wholesale cognitive decline that will make this type of intervention even more unlikely. While hyper-scalers like Microsoft, IBM, and Google are themselves warning of the “[amplified risks](#)” of soon-to-be-ubiquitous AI agents, it’s already too late—your mother’s face is in the same dataset as several snuff films and none of them can be removed, none of them can be “ablated,” they’ve already been ingested by in-bred models trained upon the prejudice and violence of their predecessors, and those same models will feed agents that will then be responsible for acting on her behalf if she were to turn to the internet to learn how to apply to a new job or [respond when she gets into a car accident](#). Is it crazy to instinctually fear these machines, which seek to execute pivotal tasks for us while actively reducing our critical thinking skills while also being owned by power-hungry despots? Throughout *Medium Hot*, Steyerl references self-styled god-emperors like OpenAI’s Sam Altman who celebrate the acceleration of machine learning-induced societal disorder and entropy, believing this corrosion to be the innate will of the universe. There is suddenly no divide [between governments and tech lords](#), between [tech lords and war](#).

If, as Flusser claims, our “lives become a function of [our] own images,” then we can begin to understand why our present era feels so maximal and yet so bland, it being shaped in no small part by a visual aesthetic built on extraction, violence, cruelty, and market-ready optimization, which elicit little more than fleeting cortisol spikes and chronic narcoleptic fatigue. If images are machine generated, then popular culture is a machine magnification. Everything is a pre-packaged populist product, a jumble of overfamiliar and uncontroversial patterns, a nutrient-deficient remix. Images were once optical; now they are statistical derivatives, averages of averages, secondhand slop.

What is slop, exactly? Slop is the sinister child of scandal and banality, it is Pinterest board Cronenberg, it is the exact midpoint of a bell curve. It is a punishing, cynical aesthetic that has infected the very texture of our lives in the form of brief videos disseminated on rotted-out, post-peak, tumbleweed-swept social media platforms. Slop is a pregnant Lionel Messi breaking down in tears; it is Donald J. Trump lustily rubbing himself in Cheeto dust; it is Jesus and the Pope queening out in heaven, turning water into wine; it is a shark wearing Nikes. Slop is a data-trained model’s idea of what you crave; it is always on-demand, at once ordinary and extreme, opulent and pitiful, an extravagant expression of pornographic mediocrity—a viscous, acrid gelatin.

**Proof alone will not be enough to break the cycle of Habsburgian datasets and models and oncoming agents.**

In her essay “Who Does the World Belong To? AI, Art and Common Sense,” Steyerl reflects on another primary reason for art being foregrounded in the conversation about generative AI. “For advanced AI to acquire a sense of orientation in the world, it needs to master common sense,” she writes. “Art or aesthetic judgement are simply instrumental milestones . . . central to this development but at the same time completely irrelevant.” But even though large language models can easily “write” somewhat

convincing critical missives—an advancement that may indeed help hasten the eventual creation of AGI—they often get lost [during multistep conversations](#), and the “art” produced en masse by their visual counterparts, such as Stable Diffusion and Midjourney, is still devoid of any sense, common or otherwise. Slop has no aesthetic conviction, no intellectual underpinning, and no allegiance to anything but tight turnarounds and the thick-headed literalism of minimally curated datasets.

# ALTMAN SIEGEL

1150 25TH ST. SAN FRANCISCO, CA 94107

tel: 415.576.9300 / fax: 415.373.4471

www.altmansiegel.com

At one point, Steyerl relates the “consequences of digital image production,” environmental and otherwise, to Marx’s notion of “social metabolism,” a “preindustrial circular economy” that involved using human waste to fertilize fields that grew food that fed humans. “After its circulation ceased, waste started to pile up in the cities, where it caused epidemics and pollution,” she writes. This “metabolic rift” has now been replicated in our machine learning moment: digital data waste is hoarded and stockpiled, spreading a disease of sameness that “toxifies physical but also intellectual and political climates,” a diarrhetic feedback loop that begets little more than suffocating inertia and exhaustion.

Reading *Medium Hot* felt like downloading an especially adaptive dataset and plugging it directly into my neural pathways—in almost no time, I’d become a small language model of my own. But viewed from another angle, Steyerl’s open-ended style is a maddening attempt to check all the boxes while following through on very few of them. Perhaps this is why it is best read as a genuine collaboration between author and reader, an assignment both exhilarating and head-spinning. I found myself alternately frustrated with the collection’s shortcomings and desperate to share it with everyone I knew. Though Steyerl has successfully produced an urgent primer for a topic I had long avoided researching, she is too content to position it as an artifact, something looked back upon instead of harnessed as a catalyst for actual change. “Perhaps they will become time capsules for future historians,” Steyerl writes of the eleven essays that make up this collection, a goal that is reflected by the book’s surprisingly accessible style. “What did people think when all this started? How did they deal with the threat of their own zombification . . . of being made obsolescent or undead by tools of statistical conformism?”

I began to wonder if this was an outcome that the Silicon Valley stooges had not necessarily encouraged but could cope with: a consolation prize, a crumb of catharsis for the creative set. Proof alone will not be enough to break the cycle of Habsburgian datasets and models and oncoming agents, as the case of *ImageNet Roulette* makes clear, and no lone essay collection could ever single handedly counter the immutable instability of build-at-all-costs capitalism. But in a time when humanity’s “self-alienation has reached such a degree that it can experience its own destruction as an aesthetic pleasure of the first order,” as Walter Benjamin wrote at the tail end of “The Work of Art in the Age of Mechanical Reproduction,” platformed artists like Hito Steyerl should respond with appropriately scaled urgency. One can imagine Steyerl’s exhaustive text existing in another medium entirely, as an open source document, a constantly adapting archive, a polemic with an overt call to action. As Steyerl herself predicted, *Medium Hot* is already a “time capsule” upon its publication—its “success” has been outsourced to its horrified readers, who must independently look through and past the rising torrent of slop and find a way to prevent this book from becoming an averaged-out facsimile itself.