





# Generativität



Herausgegeben von  
Matthias Bruhn  
Katharina Weinstock

München, 2025  
Open Publishing LMU

# Liquid Images. Holly Herndon, Mat Dryhurst, and Tyler Hobbs in conversation with Alex Estorick

In his book, *After Art*, David Joselit sought “to link the vast image population explosion that occurred in the twentieth century to the breakdown of the ‘era of art’”. The consequence, he argued, was a new kind of “image power [...] derived from networks rather than discrete objects.”<sup>1</sup> Given the growing number of artists developing transmedia practices, art is increasingly becoming a space for border thinking at the intersections of science and technology. For Joanna Zylińska, singular images such as photographs are “giving way to image and data flows” such that they are now “both objects to be looked at and vision-shaping technologies, for humans *and* machines.”<sup>2</sup> If generative artists are actively engaged in crafting an *output space* of multiple possibilities, those working with machine learning are curating an *input space* as the basis for new generations.

It was Tyler Hobbs who coined the term long-form generative art to refer to “a special class of artistic algorithm” that outputs hundreds of images, each transferred to a collector without any intervention or curation from the artist.<sup>3</sup> This new development had been stimulated by the online platform Art Blocks, launched the previous year, which allowed for code to be minted as art on the Ethereum blockchain. Characterized by large populations of images, long-form generative art represents a departure from the historical tendency of artists working with code to curate small collections of prints or plots for exhibition in physical space.

99

1 David Joselit: *After Art*, Princeton and Oxford 2013, pp. 88 and 94.

2 Joanna Zylińska: *AI Art. Machine Visions and Warped Dreams*, London: Open Humanities Press, 2020, p. 106.

3 Tyler Hobbs: *The Rise of Long Form Generative Art*, 06.08.2021, <https://www.tylerxhobbs.com/words/the-rise-of-long-form-generative-art> [Accessed 10/2024].

In his original essay, Hobbs made the point that “[n]obody, including the collector, the platform, or the artist, knows precisely what will be generated when the script is run, so the full range of outputs is a surprise to everyone.” For this reason, the long form has often seemed to reveal the emergent possibilities of code, while offering collectors an experience personal to them. However, thus far, the discussion around human-coded generative art has tended not to overlap the conversation around generative AI, whereby advanced machine-learning (ML) algorithms trained on vast datasets produce new images, texts, and videos. Holly Herndon and Mat Dryhurst have been making ML models since 2017, contributing to the development of the AI image generator, DALL-E, while their independent and collaborative practices also focus on music and voice. Following the launch of their work, *The Call* (2024), at the Serpentine, Alex Estorick hosted them in a conversation with Tyler Hobbs about creativity at the edge of the human.

**Alex Estorick:** Tyler, how did your own practice shape your ideas about long-form generative art?

**Tyler Hobbs:** Shortly after I began making generative art ten years ago, I became very interested in the idea of seeing how far I could produce a varied and continuously fascinating stream of output from a single algorithm. With my own practice, I worked to understand how to grow the output space and the number of interesting outputs that I could get from a single algorithm. It took me a while to get even one interesting image from an algorithm, much less three, five, or ten. However, after a number of years, I released *Fidenza* (2021) with 999 uncurated outputs that ended up being quite popular (Fig. 1).

But I've also looked at versions of this style of work that involve curation. For example, I did a project called *QQL* the following year that was curated [by the collector] from a potentially unbounded stream of outputs from the algorithm (Fig. 2).

My essay pointed out the sudden rise in popularity of what I called 'long-form' generative art. This referred to the practice of crafting one complex algorithm from which you want to see hundreds or thousands of outputs versus what had typically been the practice of generative artists, which was to craft an algorithm where they might curate the single best output or else maybe a handful of select outputs. A simple way to put it is: "How many images are we looking at from the algorithm?" But a more interesting way to think about it is: "What is the complexity of the output space?"

**AE:** The idea of the long form has provoked a lot of discussion among artists. Jeff Davis has spoken of his interest in "narrow algorithmic spaces," while Aleksandra Jovanić questions the premise of long-form on the basis that creating an algorithm already presupposes an infinite number of outputs.<sup>4</sup> Julien Gachadoat queries whether 'long-form' is evocative because "what is 'long'?"<sup>5</sup>

**TH:** That is a perfectly valid question, but just because it's a spectrum rather than discrete classes of algorithms doesn't mean that you can't attempt to label the ends of the spectrum. I think that 'short' and 'long' are a decent place to start, if not perfect.

On the question of infinity, that is where the complexity of the output space matters. For example, if we take a pixel grid and randomize each pixel as black or white, we're going to get

4 Alex Estorick: The Color of Code | Jeff Davis. In: Right Click Save, 08.05.2023, <https://www.rightclicksave.com/article/the-color-of-code-interview-jeff-davis> [accessed 10/2024].

5 Alex Estorick: The Power of the Plotter. In: Right Click Save, 14.11.2023, <https://www.rightclicksave.com/article/the-power-of-the-plotter-generative-art-aleksandra-jovanic-julien-gachadoat-feral-file-graph-interview> [accessed 10/2024].

a massive number of outputs but very little complexity in what we're seeing. Of course, it is difficult to define complexity – we might think in terms of *conceptual compressibility*, for example – but I think it's pretty easy for viewers to recognize complexity when they see it, and that is the substantial distinction that I'm talking about. Artists today are targeting substantially more complex output spaces than generative and algorithmic artists did in the past.

**Mat Dryhurst:** I think that the long-form frame is useful in distinguishing between producing a bunch of outputs and selecting a group as a collection. There's something useful in the essay establishing that this is actually a new kind of practice because I do think it is. Then there is the secondary consideration of the model, pioneered by Art Blocks, of creating scarce and collectible moments. Of course, that does have some relation to the market but it is also a legitimate proposal for how to value these things that I welcome. The analogy in our world is *modeling*. There's something powerful about an artist saying, "No, actually, my algorithm or my model was constructed over time with a bunch of different curated inputs with lots of trial and error in order to come up with a system that allows people to interact with it." That is where you really want optimal complexity or generalizability in the machine-learning space.

The striking thing to me is the analogy between this kind of long-form generative algorithmic process and what we've been thinking about in parallel, which is how to create a model as an artwork. That involves similar challenges. In developing *xhairymutantx* (2024) for the Whitney Biennial, we ended up trying to produce a text-to-image model that would generate reliably general enough outputs that it was fun to play with

and you felt like your contribution was actually meaningful. I understand that it is more limited in the traditional generative sense than it would be in the text-to-image space but it seems like a complementary problem. To me, the question is: “When you are in a sea of abundant imagery or infinite possibilities, what is the unique or scarce element?”

**Holly Herndon:** We do both the open version and the curated version. One way of being open is being really playful with identity and allowing people to perform my identity (through *Holly+*). The scarcity is that there is only one me but the non-scarce thing is that I can let anyone perform me. There’s an interesting push and pull there. Coming from a background defined by hypercuration where I wanted to control every second of sound that I put out into the world, it was definitely an aha moment when I heard people performing through my voice. It was an entirely new way of interacting both with other people and with my own voice. That really unlocked this duality that we have of both curated outputs, and openness with my IP by letting people perform through it. I think there’s room for both.

**TH:** I think that core to these practices is shaping the model and, Holly, although I didn’t get to actually experience your work at the Whitney in person, it felt like visitors were curating outputs from the entire potential output space of that model, cherry-picking little moments. Maybe they’re not particularly skilled curators with that model because they’re not so familiar with it, but they are doing informal curation. I think that the more generative a work is, the more important the moment of curation becomes in establishing artistic significance in a world where there are so many images.

**MD:** I think that's right, but I've also spent a lot of time pushing back against the idea that curation is the new creation.

**HH:** We're also using the term curation really liberally here. There's the simple act of choosing which bit gets out to the public. But there are also so many different layers to curation – you have the model and then different people perform the model, and what they choose to show is their own filter.

**MD:** One of the things that we've looked at a lot in recent years is embedding space, also known as latent space. For a reader, that would be how concepts are determined in large models. When you're talking about the embedded concept of a chair, generally speaking, there's a universal idea of a chair's form, and so if you take a million images of chairs and feed them to your model, when you type "chair", everyone will probably agree on the output. But when it comes to a person or a subjective term like beauty, these things get really hairy. Stable Diffusion's impression of beauty is going to be whatever was tagged beautiful the most times in its training set, which has this air of objectivity to it but ultimately is curated somewhere. As artists, that level of curation is further upstream than the image, so if you want a model to represent beauty in a way that encapsulates the voice or eye of an artist, you have to go further upstream in the development of tools, protocols, training sets, or an ensemble of algorithms in order to establish that concept of *beauty*.

**AE:** It seems to me that one of the exciting possibilities of long-form generative art is that it offers a pathway to new and emergent outcomes because of the role played by randomness in the creative process. Holly and Mat, your project *Infinite*



*Images* (2022) was informed by your work on the first iteration of DALL-E. From your experience, what is the emergent potential of machine learning?

**MD:** In machine learning there is a concept of *overfitting* where, for example, if you train a model poorly on a corpus of 20th-century art, then you might get locked into a groove where you can only produce things that look like Picasso. The ideal is that you would have a wide enough range of data to produce new outputs... I do think that these models are demonstrably capable of emergent outputs. I don't think that they're manifesting new stuff out of nowhere, rather the sheer complexity and volume of potential permutations in these large models is highly capable of novel combinations.

**HH:** With DALL-E, if that was still the model that everyone was using, it would be boring, but because it was this weird snapshot in time, it had this really particular grain and quality. We were almost playing it like an instrument, making aesthetic decisions based on what it would want to do. There was something really interesting about that.

The first album, *Movement* (2012), used the most up-to-date vocal processing available. Then, with *Platform* (2015), we were sampling the sounds of our browsing using an instrument called Net Concrete in order to capture our new life online. With *PROTO* (2019), we allowed the neural net outputs to be scratchy and dissatisfying in their timbre, which was a challenge because they were paired with really high-fidelity ensemble recordings. We had to degrade the other recordings so that they could live in the same universe together because we didn't want to release an album that was all about AI as shiny, clean, and perfect. Its weirdness is so much more beautiful.

Right now, so much sound work is retro and you can land in an aesthetic cul-de-sac where you're just circling the same 1970s synth over and over. There's something disheartening about that. Large models often capture a midness or average, which is on purpose because the goal is to have an accurate reflection of a cup and not necessarily a subjective personal relationship to your grandmother's cup. They're not really built for that. We make bespoke training data because we see that as artwork in itself, and if you view training data as beautiful mind-children that you can send to the future, the models can always be changing but you have this evergreen training data that can always be tapped into for new and interesting outputs.

**MD:** We use the term protocol art to avoid fetishizing a particular tool. The term algorithmic art is very laden with a particular toolset, but when you talk about protocols you could just as easily be talking about the Bible or the Quran. I would argue that the successful aspect of any artist's practice is the protocol, or practice, rather than whether they're using generative or machine-learning tools. It's also what differentiates them as more than a keyword associated with an image on a screen. Otherwise, it can be hard to individuate the artist and distinguish something we do, for example, from something a company does. If you regard an AI model as largely the abstract product of its training data, then you are overlooking how that data was coordinated, which might be a site of performance or even politics.

**AE:** A lot of machine-inflected images have a kind of beguiling polish, eliding the complexity of the training data into something seamless. Do you feel that?

**TH:** Yes, I do. There are so many traces to everything that's done by a machine as a simple consequence of the fact that, fundamentally, they're dealing with ones and zeros as well as grids, arrays, and loops that have such a big influence on everything that comes downstream from that. There are also other limitations like processing power or the amount of time or storage or memory that is available. Over time, that will shift but there will probably be some trace in the creations until far into the future. Just in terms of information density, the physical world has far more information than digital items tend to contain, so there is something of that artifact as well. But that limitation is also a source of novelty, including aesthetic novelty, which is part of the reason that I find images created by the computer to be highly interesting. There are also human contextual associations that we bring to our viewing of images that make them important to us that one can't discount.

**MD:** No matter how impressive the model, if you go further upstream you will likely find a human who is making really consequential creative decisions. That might be artists like us but it might also be someone at Meta making a decision about the UX of a piece of software or a protocol that we interact with every day. To me, this kind of upstream analysis of the site of the algorithm or model or UX is really useful. People are already developing a fluency with the images that they encounter as well as a resilience to deepfakes. All that is happening is that their lexicon of understanding of images and media is maturing very quickly. If you talk to an 18-year-old, they have so much more awareness about the provenance of media and its potential objectives than I did when I was that age. People don't need to nerd out about this stuff to accumulate that immunity. The human artist is going to be just fine.

Fig.1, Tyler Hobbs, Fidenza #438, 2021, Courtesy of the artist

Fig.2, Tyler Hobbs and Dandelion Wist, parametric artist Appleboy, QQL #154, 2022, Courtesy of the artists




Fig.3, Holly Herndon and Mat Dryhurst, xhairymutantx, prompt response to "sam altman visiting an art gallery", 2024

Fig.4, Holly Herndon and Mat Dryhurst, xhairymutantx, prompt response to "jesus on the cover of artforum, koda-chrome", 2024



Holly Herndon and Mat Dryhurst are artists renowned for their pioneering work in machine learning, software, and music. They develop their own technology and protocols for living with the technology of others, often with a focus on the ownership and augmentation of digital identity and voice. These technical systems not only facilitate expansive artworks across media, but are proposed as artworks unto themselves. They were awarded the 2022 Ars Electronica STARTS prize for digital art. They have sat on ArtReview's Power 100 list since 2021. Holly holds a PhD in Computer Music from Stanford CCRMA. Mathew is largely self-taught. They have held faculty positions at NYU, the European Graduate School, Strelka Institute, and the Antikythera program at the Berggruen Institute. They publish their studio research openly through the Interdependence podcast, and co-founded Spawning, an organization building AI models on consenting data. Their critically acclaimed musical works are released through 4AD and RVNG Intl.

Tyler Hobbs is a visual artist from Austin, Texas who works primarily with algorithms, plotters, and paint. Hobbs' artwork focuses on computational aesthetics, how they are shaped by the biases of modern computer hardware and software, and how they relate to and interact with the natural world around us. Hobbs' Fidenza (2021) series profoundly impacted the generative art landscape, reshaping perceptions of "long-form" generative art. Hobbs' two most recent solo exhibitions were at Pace Gallery in New York and Unit in London. His work is in the collections of multiple prominent institutions including the Los Angeles County Museum of Art and the San Francisco Museum of Modern Art.

Alex Estorick is a media theorist who seeks to develop inclusive approaches to new technologies. As Editor-in-Chief at Right Click Save, he aims to drive critical conversation about digital art. He has developed some of the first academic courses on blockchain, NFTs, and Web3 and was lead author of the first aesthetics of crypto art. In 2023, artnet named his artistic collaboration with Ana María Caballero, artifacts, as one of the artworks of the year.

Herausgegeben von  
Matthias Bruhn  
Katharina Weinstock

DFG-Schwerpunktprogramm ‚Das digitale Bild‘



Erstveröffentlichung: 2025  
Gestaltung: Lydia Kähny, Satz: Annerose Wahl, UB der LMU  
Creative Commons Lizenz:  
Namensnennung - Keine Bearbeitung (CC BY-ND)  
Diese Publikation wurde finanziert durch die Deutsche  
Forschungsgemeinschaft.  
München, Open Publishing LMU

**DFG** Deutsche  
Forschungsgemeinschaft

**UB** | Universitätsbibliothek  
Ludwig-Maximilians-Universität München

Druck und Vertrieb im Auftrag der Autorin/des Autors:  
Buchschniede von Dataform Media GmbH  
Julius-Raab-Straße 8, 2203 GroÙbeersdorf, Österreieh

Kontaktadresse nach EU-Produktsicherheitsverordnung:  
info@buchschniede.at



DOI <https://doi.org/10.5282/ubm/epub.126472>  
ISBN 978-3-99181-339-2

Reihe: Begriffe des digitalen Bildes  
Reihenherausgeber  
Hubertus Kohle  
Hubert Locher



