The Art
We Overlook
Mat Dryhurst

In 2015, Elman Mansimov published a conference paper about a model he and collaborators had developed called alignDRAW. The system could generate (blurry, half-formed) pictures from simple text prompts—a dog blob, a blurry bus, a mirage of a herd of elephants. It was fragile and strange, but unforgettable. alignDRAW was one of the first serious attempts at what we now take for granted: the ability to type words and receive images.

I think it's important to validate code experiments like alignDRAW as artworks of consequence. Mansimov's GitHub repository and paper ought to be understood in the tradition of process-based art. The point is not only the finished artifacts, but a view into the reasoning, methods, and limitations behind them. What's new about the networked age is that we can follow the development of these aesthetic and engineering experiments in real time, should we take the time to look.

Anyone closely following machine learning over the past decade knows that it's the academic publishing platform arXiv, and esoteric corners of Twitter, where you're most likely to encounter glimpses of the world to come. Our feeds have become stages for engineers who display more mastery and awareness of the defining medium of our time (software) than any traditional artist. These code experiments emerge from all manner of sources: anonymous accounts testing prospective applications in public, academics demonstrating papers, or a steadily growing ecosystem of creative engineers working within or alongside larger tech companies to quickly test new interactions with existing tools.

I've noticed how, over the years, the engagement these interactive code experiments garner online has steadily grown to dwarf engagement with traditional arts. At the dawn of social media, it was common for marketing from tech companies to describe their engineers as "rock stars." Now it's more likely for artists to position themselves as engineers. The cultural energy and attention has moved. Much of contemporary art validates itself through technical innovation and systematic exploration, while the best engineering validates itself by producing experiences that feel genuinely aesthetic, even sublime.

The sublime is often associated with mountains, storms, and natural wonders. For Immanuel Kant, the sublime was the mind stretched against its own limits. For Jean-François Lyotard, it was the unpresentable made present, the imperative of the artist. The contemporary sublime is prompted through engineering. It's the uncanny thrill of a system producing something we didn't think possible, conjuring a new degree of freedom we didn't know was available. As machine learning has matured and established itself in popular consciousness, a new model release from OpenAI is now the closest thing we have to a shared media event equivalent to last century's film trailer or album launch. There's now a popular audience for encountering and critiquing new and alien tools. Perhaps a thirst for a new century, coupled with an implicit understanding that experiments interacted with today will play a consequential role in our future lives.

The alignDRAW outputs were blurry, spectral half-dreams. As with many early experiments in AI-generated imagery, their incompleteness invited the observer to complete the image themselves. There's a gorgeous boundlessness to this early work, inherently participatory, as these initial images prompt the viewer to imagine where they lead. Many machine learning practitioners (myself included) will report that there's an underwhelming dullness to image and music models becoming more capable of producing passable media. The joy is in the pursuit, tinkering, and thrill of discovery. When we're invited to fill in the blanks, anything is still possible.

Ken Stanley and Joel Lehman argue in *Why Greatness Cannot Be Planned: The Myth of the Objective* (2015) that true innovation emerges not from optimizing for a particular target but from open-ended exploration. Stanley, the godfather of open-endedness in AI research, has spent his career demonstrating that remarkable discoveries happen when we follow intuition over a plan. His research articulates something the art world has understood for decades: Overemphasis on outcomes can stymie the very creativity that produces them. The bleeding edge of machine learning research is now trying to emulate how humans produce remarkable things, not through rigid, objective functions, but through curiosity, play, and following what seems most interesting in the moment.

This is exactly the territory contemporary art claims to occupy. It's the institutional alibi: We facilitate experiments that help us view the world differently, that reveal new possibilities, that engage with the conditions of the present. Open experimentation produces remarkable things simply by virtue of not over-optimizing or constraining what the outcome might be. The studio practice, the sketch, the failed attempt that leads somewhere unexpected—these aren't preludes to the real work, they *are* the real work. Many code experiments share this spirit. Have an interesting idea, put it out into the world, see what happens. This interactive feedback dimension matters enormously. The experiment isn't complete when it's published; it's complete when it circulates, when others fork the repository, when someone discovers a use the creator never imagined.

If you show me a contemporary art exhibition exploring AI, generative systems, or algorithmic aesthetics, I can point you to the arXiv paper or GitHub repository that preceded it, often by years. arXiv and engineers publishing code experiments are demonstrably upstream of gallery work. Many practices we now take for granted (text-to-image generation, style transfer, latent space exploration, prompt engineering as a creative practice) all emerged first as provisional

experiments shared in academic papers and repositories. They circulated through technical communities, were forked and expanded upon, and only later migrated into gallery contexts, often by people who went to art school and frequent the same social circles as those housed in art institutions, while the engineers who originated these ideas remain largely unacknowledged.

This represents a failure of the institutional art system to live up to its own stated values. One of the enduring challenges of art historical institutions has been tracing provenance, documenting influence, establishing who did what when. Sketches are lost, conversations go unrecorded, the chain of influence becomes murky and contestable. But much of that provenance is now visible on GitHub. Every commit, every fork, every conversation in the issues thread. The complete genealogy of an idea, timestamped and publicly archived. We can watch in real time as someone has an initial insight, shares it, watches others build on it, sees it mutate into applications few anticipated. The archive is public and precise.

Code experiments are entirely consistent with the logic and purpose contemporary art claims: process over product, methodology made visible, participation invited, new aesthetic territories explored. If contemporary art institutions claim to care about where genuinely new forms of seeing and making emerge, then they need to expand their purview to these areas. I feel a sense of injustice watching engineers who produce genuinely groundbreaking work go unrecognized while others (often people with MFAs and institutional access) receive credit for repeating these experiments years after the fact. If we take seriously the idea that contemporary art should engage with the defining conditions and technologies of its time, then we must take seriously the people who are inventing those conditions and technologies. Fortunately, a new, often technical, collector class has emerged that recognizes the artistic significance of these works. alignDRAW has been widely collected and exhibited by the NFT community, and it is no surprise that many machine learning experiments have been embraced and canonized on-chain.

Contemporary art institutions spend considerable energy trying to find their place in a rapidly changing, technologically fueled world. In the best cases, they attempt to support artists in producing their own experiments. In the worst cases, they pander to social media, trying to make exhibitions more superficially photogenic and shareable, satisfying the demands of often-debased platform incentives in lieu of playing a greater role in the development of the defining cultural conditions of our time. The opportunity is clear: Dig into these archives and begin to validate and canonize the code experiments that underpin the world around us. The archives exist and are public, but they require mining and debate. They need the kind of serious critical attention and institutional validation that contemporary art institutions provide.

I was speaking at a conference in Switzerland recently where curators and museum directors lamented that contemporary art appears to be less engaged with by the public. When I spoke, I said this interpretation confused me. Between the development and popular dissemination of machine learning and crypto, I have never known a time where art was being more hotly debated and contested by the wider public. The challenge is that institutions rarely capture that energy because they're overlooking its source: the tinkerers and visionaries quietly publishing experiments in software.

Just as alignDRAW's blurry images invited the viewer to fill in the blanks, code experiments require institutions to do that same interpretive work—that willingness to engage with something provisional and incomplete, that ability to recognize something promising before it has been validated. Isn't that supposed to be what contemporary art institutions do? Isn't that the thrill of it?