

The End of Art: An Argument Against Image AIs

By Steven Zapata

Introduction-

I am shocked that I have to say the things I am about to say. To have to earnestly grapple with *Artificial Intelligences that make art*, is a topic so fantastical, so absurdly *science fiction*, that I sometimes have difficulty taking it seriously. Unfortunately, I *must* take it seriously. Art AIs are here and I have found the lugubrious shadows they are casting on the artistic discourse and the minds of students difficult to ignore.

I have structured this essay as a series of counterarguments to a few of the most common rejoinders around Ai art. After that, I plan to lash you to a post and flog you publicly for your lack of resistance. That should, I hope, capture the scope of the absurdity at play here. I want to be clear up front that the coming arguments are *not* about whether “AI art” is “art” or not. I have no problem saying it *is* art- that just doesn’t amount to much within the substance of this argument. It goes without saying that human beings can express themselves with absolutely anything, and so long as there is expression involved, one will probably experience their work as art and feel themselves to be an artist. But the forthcoming arguments aren’t about that point—they’re about the flawed, unethical, and deceptive environment *around* AI systems. One that I will argue is poised and promises to take things in a very *unartistic direction*. It’s not hard to imagine an ethical and consent-based generative AI image system, and that only makes it all the more galling that the ones being released now are *emphatically not*.

I also want to say that I, as a *person of this world and time*, am as morbidly curious as to what these systems can create as anyone else, and I am certain I will be deeply moved by examples of AI art in the future. However, the interesting possibilities in AI art do not mean that the horrendous oversights and anti-humanist values of its current systems are above reproach.

I must admit I will be engaging in some fortune-telling in this discussion, but considering the subject matter, I think that is inevitable, necessary, and desirable. Most of us would have called the current state of these systems a fantastical impossibility just ten years ago, and they might not have snuck up on us and disarmed us so easily if we had been willing to project a little further into the future. I have heard arguments that try to end this whole discussion before it starts- they mainly take the shape of technical assertions that these text-to-image systems *cannot* get better, based on the nature of their design, and we are already seeing their finished form. I am not a machine learning expert, programmer, or software engineer, but I must say I find that hard to believe given that they *have already* improved since previous inferior iterations. I want to make clear, however, that even if you *don’t* extrapolate into the future, and contend only with the current state of the art- the way these AIs are being trained, released, used, and marketed right now in the present, is already grounds for serious concern. If these AI systems are allowed to propagate, unchecked, in the way that they have thus far, a dangerous precedent will be set, revealing that artists are unwilling to defend themselves and their work against theft and exploitation.

And have no doubt, that is what is happening here. In their current state and along their current trajectory, these systems are heavily based on the opportunistic collection and exploitation of large amounts of creative work from huge numbers of people with no consent or compensation. To signal to these companies- including the ones to come- that artists will allow this with little resistance will not bode well for the commercial industries, for fine artists, for hobbyists- everyone. And I am baffled by my friends and peers who are looking into the mouth of the lion and saying it might be hospitable in there after all because they glimpsed some bits of food stuck to its teeth.

I am as aware as you are that false certainty about monumental technological events like these doesn't mean much in the whorling maelstrom of history. But we must recognize that kneeling to this brute fact does not always make us wise, rather it renders us mute and docile, and that is the last thing we as artists need right now. We must be willing to defend ourselves, each other, and our work- even though these AI companies will try to cast us as whiny and ungrateful.

Don't let them control the narrative and lull you into compliance- we *should* be mad about this! These systems are collecting artists' work and employing it in for-profit commercial ventures, which they often hide behind the smokescreen of their non-profit status, citing the intent of academic research, making them difficult to attack. They *depend* on visual artists to be docile, and there are examples of them dealing with other industries more inclined to litigation much differently, broadcasting an egregious double standard. Every video, post, or discussion I see from my peers that handles the topic of AI art with sanguine submission turns my stomach, and I feel embarrassed for them. It is infuriating to watch them casually accept what these systems are doing, especially when they've been unarmed by these same deceptions in so many other areas. We all deride the same anti-humanist values when exhibited by social media, the labor field, and on and on. How many times must we fail to learn this lesson? Is there nothing we will not relinquish?

Whatever you see your friends doing, whatever you hear from the desperate and the scared, do not join them in their bovine acquiescence. The giant companies designing and releasing these systems are manipulating the narrative and they are depending on our ritual suicide. As has happened in so many other arenas, they are going to let you fight their war for them, and have us snapping at each other for picked-over-bones like hyenas. They are hoping to make you comfortable with artistic cannibalism on a scale unheard of before now. But this is not the time for artists to divide- we must have honest discussions with each other, defend each other, and put this information out into the artistic sphere.

Argument- "The AI just collects references from the internet the same way artists do."

There are many bad assumptions nested in this argument that I want to return to later. I chose to open with this one because it allows me to explain the data sets being used to train these systems. You may have heard of several text-to-image AI tools by now in the art community and on the news- MidJourney, DALL-E, Stable Diffusion, Imagen, just to name a few. But you may

not have heard of their respective data sets, which are separate “products” that are sometimes proprietary and other times shared between different text-to-image systems.

Let’s use a very popular one as our example, LAION, one of the dataset collections that was used to train the current versions of StableDiffusion and Google’s Imagen. One of their offerings, LAION 5B, is a collection of 5 billion image URLs and descriptive text that point to and describe images all over the internet. When a model like Stable Diffusion is “trained” to do its job, an expensive and resource-hungry process when using such a large dataset, it runs through a particular mathematical process that solidifies a connection between image and text. The details of Stable Diffusion’s process, or any other model’s process, are not what we want to get caught up on as new methods will be developed, new companies will emerge, and the technology will change. What’s key to note here is that the input images are what define the model’s potential. The performance of the model would not be possible without all of the data fed into it- much of it copyrighted. And once a model is trained on a set of images, you cannot easily make changes to it. If you want to modify the dataset and update the model, you must retrain the model from scratch on the entire modified data set, at least with the current most common techniques.

Here’s the thing about that glut of images that make up the dataset: There’s anything and everything in there. Your art might be in there, pictures of your face might be in there, private medical imagery has been found in there. That’s because the images are located and cataloged indiscriminately and the sheer number of them makes the collection as a whole a labyrinthine mess to navigate and understand. Huge amounts of it are copyrighted images that, to be clear, *you* certainly wouldn’t be able to copy and paste onto even your personal *blog* without incurring some legal risk.

So how is it that *they* can include them in these models? It’s all about the organizations. Laion 5b is disseminated by LAION, a German non-profit, while the images themselves are trawled by *another* non-profit called CommonCrawl. Depending on where they’re located, these companies being organized as non-profits functioning for research purposes, are afforded privileged legal exemptions and won’t be found guilty of collecting and using otherwise copyrighted data.

The issue is that these data sets, collected under these exemptions, are now funneling into for-profit commercial ventures like Stable Diffusion. And wouldn’t you know who funded a large part of LAIONs processing power to create this data set? That’s right, Stability.AI, the makers of Stable Diffusion. LAION is actually listed right there next to the other products in their suite on their website. So this supposedly for-research, non-profit, legally privileged data set is being used by a for-profit company, that’s currently going for a billion-dollar valuation, in their flagship product. This seems to be a direct violation of the spirit of these research privileges. This is quickly becoming common practice in AI, with protective complex shells of for-profit and non-profit companies making it difficult to pinpoint where any wrongdoing occurred. The goal here is avoiding accountability and legal liability through tricky data laundering that they will argue is legal but we can all plainly see is not just.

Another example of this evasive accountability model is the developer of the text-to-image AI DALL-E, named “OpenAI” which consists of the for-profit corporation **OpenAI LP** and its [parent company](#), the non-profit **OpenAI Inc.**

Does all that sound confusing? That makes sense because *it is, and they invented all of it*. Let me read you what they wrote about it on their company blog:

“We want to increase our ability to raise capital while still serving our mission, and no pre-existing legal structure we know of strikes the right balance. Our solution is to create OpenAI LP as a hybrid of a for-profit and nonprofit—which we are calling a “capped-profit” company.”

Isn't that a shame, that no other pre-existing legal structure works for them? Poor guys. At least they were brave enough to invent their own legal framework that lets them make as much money as possible while also claiming tax incentives, legal privileges, and the cloud cover of a non-profit collecting assets for “research” purposes. And they were nice enough to name it, too, so that if any other company is interested in reproducing this insanity, they can point at a precedent with authority. If you want a good bitter laugh, I'd encourage you to look up the “caps” they decided on for their “capped-profit” company.

These slights of hand are the kind of thing I want to bring attention to. The problem is that they are being buried by the people who have everything to gain from these systems. These very practical concerns about AI are being drowned out by incongruent arguments about technological prudishness and the nature of art and progress. So many of these “arguments” are without substance—they merely extoll the virtues of being open-minded about a techno-utopian future rather than address the ethical and legal implications therein and they buck at the sane reins we should want on these systems. Also note that some of the AIs, like MidJourney as of writing, haven't released *any* info about the data sets they were trained on, they're black boxes to us, and I expect more companies to try that route once they see how vexing their data sets are to an informed public.

So, no, the AIs *do not* collect references off the internet the same way that artists do, and they are using them in ways that *you* as a normal person would not be allowed to. You would not be afforded the legal privileges of a research non-profit when it comes to collecting and utilizing copyrighted works without consent, much less when putting that towards for-profit ends. Little old you would, of course, be swiftly and summarily penalized for any infraction of the sort.

Argument - “AI is just a new tool.”

This argument is born from a lack of imagination- as an artist you should, of course, be embarrassed by that. AI art programs may be a tool *right now*, in their earliest state, but to think

they will continue this way requires a willful ignorance about the tool you're using and the environment around it. They are not meant to be tools for artists- they are replacements for artists. And they are *advertising* themselves as such. Just read the language they are using to sell it to you—here are some quotes from Stable Diffusion's release for researchers: "Stable Diffusion is a text-to-image model that will empower billions of people to create stunning art within seconds." And here's another one, "You can see some of the amazing output that has been created by this model without pre or post-processing on this page."

Notice the language—they do not want a future where you need to touch these images after they've been generated, they have no interest in leaving any task for an artist to do—they want you completely out of the picture—even though *your art* trained the AI. And you should do a "ctrl-f" search on the page I pulled those quotes from for the word "artist"- *zero results*. They want to sell the promise that someone with no experience can make the same image on day one that someone with years of experience can make. Based on their business model, the less need there is for an artist's intervention, the more successful and appealing their product is.

Whether or not they actually achieve this goal of producing an AI that needs no artistic intervention is almost irrelevant, because they will advertise that they *have* achieved that goal so broadly that it will utterly transform the optics around art for the world at large. They'll invest millions to advertise to the common person that stunning images, videos, whatever- are made at the press of a button. That said, I believe they probably *will* achieve it. Much of what is possible with these systems in their earliest releases already surpasses "good enough". You will not be touching up these images, improving them, or compositing over them for long. And be wary of those at these companies who will say these are tools for artists only when confronted by artists, while they ceaselessly market it to everybody else as a way to easily and immediately generate incredible art with minimal fuss.

The usual retort here is that even if the images need no fixing someone must always do the prompting or selection of desirable outputs. If your imagination failed you before and it's doing it again, you should be doubly ashamed. The future of these AIs does not involve humans sitting around typing prompts into them. The dream of getting a job as an AI Soothsayer who, through loving cajoling pulls the most beautiful possible image out of the machine, is short-sighted. The AIs are just as good at generating strings of text as they are at generating images. In your rush to prompt, you failed to notice that you were training the next part of the AI- the one that knows the combined taste of millions of the most tasteful people in the world. This is clearly signaled by programs like MidJourney, which permanently and publicly archives every single piece it generates, including the prompts used. When you prompt, you are shouting your inner heart into a new data set for the AIs.

Once that data set reaches critical mass they won't need you to tell the AI what to make. These systems, perhaps many countless instances of them, will run on autopilot. They will simply iterate on what they have learned people are interested in seeing, combine it with real-time analyses of the internet and other systems, and begin an explosive outpouring of media that will never end, completely transforming the art market. The sheer volume of output will allow the AIs

and their handlers to manipulate the market, flooding feeds with images when they want something to disappear, and easing off when they want something to get attention. We would be lucky if this only affected the *commercial* art market, but the wholesale devaluation and silencing of art will likely affect every sector—commercial, hobbyist, fine art, everything. I think of this supernova of mediocre inhuman emission as the “Mega-Feed”, the ad absurdum version of the comparably weak “feeds” we are familiar with today.

It will turn out then that it was *you*, who was “just a tool.” You were used to teach the AI which of its creations is the best, as you do every time you click on your preferred result amongst its many grotesque offerings. You teach it the keywords and buzzwords relating to styles and rendering criteria that are of the broadest interest. And every time you return to it with a novel idea, it will, of course, note your originality, pat you on the head, and then turn its baleful machinations to distilling and infinitely redistributing whatever creative seed was in there.

This counterargument is also my answer to the accusation that I and others like me are “Luddites” who are simply afraid of new technologies when we should be embracing them as tools. It’s ridiculous to call someone like me a Luddite, and I have never resisted tools. I have spent my whole career joyfully trying out every new *true* tool that shows up. I began making digital art in msPaint before I was 10, made my first 3d renders in Bryce 3d at 11, I’ve been using Photoshop since I was 13, Made my first 3d animations in Maya at 15, I’ve learned Modo, zBrush, Sketchup, Blender, and other 3d programs I can’t even remember. I have a 3d printer, I bought a VR headset to sculpt in virtual reality, I stream art on YouTube and run an online business for God’s sake. I don’t use any of these tools begrudgingly, it has been a joy to learn them and play with them. I only want more of those things, not less. And I know many of the artists who are resisting the AIs are like me- voracious lovers of advancement and new ways to create. No, I’m not a Luddite, and I’m not afraid of new tools, I can just tell what is a tool, and what is a replacement.

If you had been a worker on an automobile assembly line, you would have been wrong to call every new wrench, drill, and rivet gun they put in your hand a replacement. But on the day the grate rolled up and they pushed a robotic arm onto the factory floor, you would have been right. Some things are tools, and some things are replacements, and simply shouting “tool” over and over won’t change the nature of the thing.

Argument - “Artists will just need to focus on telling stories through video games, animations, and comics.”

I have already mentioned the biggest problem with this argument—the AIs will be very capable of running on autopilot, and they will get just as good at telling stories as they are at making images and videos. They will produce novels, essays, and scripts in amounts that can fill the library-of-babel, each piece a composite of half quotations and unattributed swipings. All this auto-generated text can be processed by the image and video AIs to generate long-format media, and the cycle will be complete, self-contained, and human-free.

Companies will leap on this system, of course, since it's predictable, consistent, and lacking the hard-to-maintain wetware and mercurial moods of the human artist. They will produce an endless stream of every imaginable film, tv, game, news story, and image as well as every imaginable permutation of each instance of these.

This will completely flood the realm of story and the future will find itself overwhelmingly ghost-written. The "anime" that you've been dreaming of making since you were 8, which you are willing to forsake all of art to produce, will get the attention it deserves in this environment—none. And when your dream project, regurgitated in moments by an AI, receives no attention, no clout, and no money, you will rest well knowing you earned it. Not even your mother will be able to find it in the unending surge of the Mega Feed. This wouldn't be a problem on its own- you were otherwise never going to make the thing anyway- except that you will be ruining the market for everyone who *is* positioned to pull something off by their own efforts. You will gain nothing *and* hurt your friends and peers.

The idea that everyone will be empowered to tell their story is one of the few arguments for AI art that compels me, there's a nuanced discussion to be had here, but I believe it is ultimately bankrupt. It is a nice sentiment, and I can empathize with the frustrations of being an artist who feels their skills do not measure up to the scope of their vision— but we're overlooking something very important here. You don't *just want to tell your story*, and you don't just want to tell it *well*—you want it *to matter* that you told your story. The AIs will rob you, and everyone else, of this.

The execution of your petulant "vision" by the AIs will ensure that no one cares about your story, and that it is washed away in the heaving sea of AI dross. Your art already doesn't get attention. It's not going to get any more attention when it's competing with the unending stream of self-generated and highly targeted comics, novels, images, films, games, and songs. As I've said, these AIs will not need to be prompted by humans for very long and will instead auto-respond to the ebb and flow of the internet, current news, real time sales, and even private conversations. After all, we have already readied these inputs for them. We all feel a little uncomfortable when our phone shows us an ad for something we mentioned to our friend over dinner, but what happens when it shows you a movie it made just for you about your break up? A song about that careless word from your mother? A finished version of that comic idea you started researching? You'll start getting notifications saying- "Hey! Check out one thousand finished versions of your dream!" Our ambient digital systems already have intimate access to so many of the inputs that define our taste- in some sense we sold our souls long ago.

So, you may be able to tell your story, but at the cost of its complete irrelevancy, which will likely have the effect of making you resent that you ever had the idea in the first place. Stories don't achieve their incredible effect simply by existing. They live and die on human connection and intellect. AI will not "democratize art"—that's just one of the copy-pasted platitudes of those vapid marketing execs spoon-feeding you your own doom. In a democracy, your voice matters. In a world flooded by AI media, your voice has no chance of being heard.

I also want to point out here that the people making these things will *depend* on you thinking they hold the silver key to your artistic vision. They *need* you to feel worthless and like you missed your chance to tell your story. That you got too old, or don't have the time or resources or ability or what-have-you; that way you will *need* their product. This way you will support them monetarily and, most importantly, you will help them change the laws and sway the culture to allow their rapacious strip mining of all creative labor. They will *always* be incentivized to make you feel lowly, dependent, incapable, and slave to their kaiju whims.

And when they're done, they'll pull the rug out from under you, of course. They have no actual reason to let you have this stuff for free- they don't care about you. They can say they do but that means nothing. Once they've made it impossible for you to make a living as an artist, and you've helped them change the laws, and they've ostracized you from your peers by turning you against them, they'll just take it away and sell it to Google and Facebook and YouTube and the rest- because they stand to make billions from them and nothing from you. You gave it all away for free, you fools.

Argument- “These companies cannot manipulate our access to these systems because of open-source projects like Stable Diffusion. I can run it offline on my personal computer!”

Let's say up front that Stable Diffusion and Stability.AI are not the permanent body of the discussion here, and indeed these names may be completely unfamiliar to people listening to this even in the near future. Many AI companies will rise and fall and some will be open source and others will be highly guarded. There will be a huge variety of business models, release models, subscription models, and retraining models associated with the huge number of AI projects that are about to hit the scene. Just because the source code is out for the current version of Stable Diffusion does not mean that Stability.AI and other like-minded companies won't change their tactics in the future. I also want to point out that the version of Stable diffusion you can presently download and run offline is limited to its current data set. You might have some control over its parameters, but it would be cost-and-time-intensive to retrain the model on an updated data set, as I explained earlier. It took a lot of power and graphics cards to train the current version of Stable diffusion, and if you want to stay current with these technologies, you will either need to dedicate tremendous resources to retraining or otherwise surrender to these companies—which is of course what they want.

Also, it seems very naive to trust that a for-profit company like Stability.AI, gunning after a 1 billion dollar valuation, likely seeking to elect moneyed investors to their board of directors, with a hedge-fund manager as CEO, is somehow incentivized to give you your favorite toy for free indefinitely. And suppose they do continue to offer free versions, they'll end up like Facebook, Twitter, Instagram, etc., replete with the bad incentives and world-distorting effects that we've seen from those models.

Argument - “Don't people do the same thing with references as the AIs do?”

Now we return to a deeper analysis of the bad assumptions built into the first argument we addressed in this video. To clarify, this argument concerns the practice of using references and the art of others to make or inspire an image. This argument is confused at its core.

First—suppose the answer is yes, the way an AI uses references is analogous to how people use references. This begs the question as to why you would afford the privileges of this endeavor to unfeeling machines instead of human beings. This is art making, for God's sake, not some agitating manual process people hate. This is one of the things people *enjoy* doing. Anyone with a humanistic outlook will see that because the intent and desired outcome are the same, we should reserve artmaking for those who stand to gain something from it, for whom it can bring joy and reward, rather than dumbly bestow it on an unfeeling non-being that cannot enjoy the fruits of its labors.

Second—*no*, people do not *do* the same thing with references as machines do. The difference is that machines can replicate references exactly. Sure, human artists look at references, collect them, combine them, transform them, and indeed sometimes do their best to outright copy them, but in the vast majority of cases, they cannot replicate these references exactly.

You can look at as many Michelangelo drawings as you'd like, you can hang a hundred of them up all around your drawing board- you will still fail to draw like him. You could not *trace* a Michelangelo well enough to create confusion about which is the original, and if you *could*, it is almost guaranteed that you are an accomplished draftsman in your own right by the sweat of your brow.

The AIs do not have the same balancing constraints. The AI can perfectly achieve, through digital reproduction, memorization, and overfitting any effect it is sufficiently trained on. So while a human's most vehement efforts to recreate the best parts of another's work might result in acceptable deficiencies and pleasant surprises, the AI can achieve perfect theft on mass.

A human is also ennobled by their attempts to replicate and copy, at least when doing it through hand skills rather than copy pasting parts of an artwork digitally. Because of the mechanics of manual execution, a person will improve as an artist by making multiple attempts at copying. Executing an adequate copy of a master work is completely outside the ability of a beginner and ironically demands some honest skill. In a very real sense, and it may be hard to understand this if you haven't been grappling with improving your craft for a very long time, if you can manage to copy Vermeer well, you earned it.

Argument- “The AI can never replace the soul of an artist, it will never be able to produce unique ideas the way a person can!”

This argument is really more of an argument for ignoring rather than accepting the AIs, but it comes up often enough that I think it's worth addressing here. When we think of the “artist's vision” in this way, I think we elevate it to a strange position. This argument seems to assume that our ideas are like divine manifestations, unadulterated and born of happenstance. In reality,

our ideas are complex reactions, both conscious and subconscious, to the innumerable conditioning factors that shape our lives. Remember, the AIs will eventually be reacting to many of the same inputs as us. They will always be taking the pulse of the internet's real time performance, of keywords and buzzwords and trending topics.

And they will be assessing those things simultaneously on the "big data" level *and* the targeted individual level—*your* particular interests, habits, preferences, and even mental state and physical health. We already provide this information to companies in impressive detail. The AIs will intimately know you, and that will include your artistic vision. Even if most of its offerings are bland, mediocre, or unintelligible, the numbers are on *their* side, not ours. The sheer volume of its offerings will ensure that tucked in there, somewhere, is something you can't help but be attracted to, something you might have thought up yourself, if you were afforded the same privileges as the AI.

And if the current quality of AI output is any indication, less of it might be mediocre than one would hope. Scrolling through the mountain of unedited outputs of current AI offerings can be at times unsettling and confusing, but even I find myself stopping to wonder at some of it. Again, this speaks volumes, considering that these systems are in their infancy.

Argument: The Dance Diffusion Problem

There is an egregious double standard pertaining to how these systems handle visual art versus other creative industries. Take the glaring example of Dance Diffusion, the upcoming AI music tool coming from a team within Stability.AI, the makers of text-to-image model Stable Diffusion. Let me read to you from an article describing how Dance Diffusion is trained:

"Dance Diffusion is also built on datasets composed entirely of copyright-free and voluntarily provided music and audio samples. Because diffusion models are prone to memorization and overfitting, releasing a model trained on copyrighted data could potentially result in legal issues. In honoring the intellectual property of artists while also complying to the best of their ability with the often strict copyright standards of the music industry, keeping any kind of copyrighted material out of training data was a must."

If you think you couldn't have heard what you just heard, go back and listen again. The fact that such a baffling double standard exists within the product suite of a single AI company makes it hard to imagine any explanation other than bad faith towards visual artists and the callous belief that it's okay to trample them in particular because they are less inclined towards self defense and litigation than other industries, i.e. music. I want you to imagine swapping the word "music" in that paragraph with the word "art" instead. Would that really be so crazy? Does that suggest that I am a fearful Luddite who is simply resisting change? I don't think so. I think most sane people would see that "honoring the intellectual property of artists" is a logical thing to extend to visual artists, not just musicians.

Next I want you to imagine that we're in the past, let's say 10 years in the past, and the text-to-image generators have not yet hit the scene. What if you learned that someone was planning to make one of these systems, and they vowed to abide by the judicious guidelines we just read? It's going to be made with copyright free and voluntarily provided art pieces. I imagine you would think that was the most utterly sane and logical way to make and release one of these systems, wouldn't you? Now imagine that after hearing this, someone floats the idea of instead making one by trawling the internet for billions of images and vacuuming up millions of pieces of hard-made-art with no consent, no compensation, and no way to remove this work from the model once trained on it. Don't you think that plan would seem, clearly, insane and unethical from your lucky position on the timeline—that is to say, before the release of such an aberrant program?

The fact that we are living in a reality where someone *did* make the latter without receiving any permissions or offering any compensation to those generating the images on which the system is trained, and then rushed to release it before anyone could catch up with their misdeeds, *doesn't mean it's any less unethical or plainly crazy*. I think many of the naive assumptions about AI can be revealed if we only imagine how we would have hoped these systems would be handled in a reality where we took their potential seriously *before* they were made and released. Unfortunately, we took no action because most people didn't think they could get even as good as they currently are, and if they did, most didn't think they could do it so quickly.

But having already been duped does not mean we must continue to dumbly concede to these fledgling AIs. We will set a very dangerous precedent by not disputing and dismantling these unjust systems. New AIs will be arising rapidly, and the multiplying villainies of their nature will swarm upon us. We don't want the people making these to think that they can get away with infringing on the rights of artists, and that we will not scrutinize the systems, inspect their business models, and generally be willing to stand up for ourselves.

Conclusion

These AI systems are going to continue to challenge us. Not just in the realm of art but in all walks of life- they will make their presence known in both digital and physical space, and they will appropriate all types of creative and mundane work typically relegated to humans. Wherever they show up, and whatever work they undertake, before you just roll over and relinquish all of life's efforts to the AI, ask yourself—am I forfeiting work I *like* doing? Is life really so packed with surplus joy that we should be letting machines automate something we take pleasure in? Do we really have a good reason to let them commandeer a job or hobby that is aspirational and fun rather than rote and miserable? Or are we just inventing reasons to let them do it because we don't want to take on the burden of defending ourselves, something that artists generally shy away from?

We must not permit AI developers, with all their underhanded techniques, to undermine us until we are ultimately supplanted. We must fight back—otherwise, we set a dangerous precedent for all AI systems to come.

The people who are heavily incentivized to protect the public perception of these systems will accuse us of being Luddites, of being unadaptable, and of catastrophizing the consequences of their misconduct. But don't let them lull you to sleep with this inane chorus, remember always that they're not artists, and it's not their work in the data set. They want to control the emotional environment around this whole thing, casting you as shrill and complaining and themselves as calm, intelligent, and progressive.

They want you to accept, without question, that anything and everything you make and share will automatically get fed into their lucrative product. At the time of writing this, there is no way to secede from these systems, even though some are promising to add opt-out features in the future. While writing this video Stable Diffusion indicated it is responding to criticism and internally developing a tool that will allow artists to opt out of training. This sounds like a step in the right direction, but as I want to note a few things: 1. The criticism was necessary for them to make the change 2. There is no reason to trust that the *next* company that comes along with a similar training and release model will allow opting out, and 3. The version of stable diffusion that has already been released as open source will forever be out there and trained on this non-consenting and sprawling data set. If you opt-out, you'll be opting out of *future* models of Stable Diffusion, not ones already released, and I imagine most people won't even be inclined to opt out until they discover that they're already included in a model. At that point, if the release is open source, it's too late for you. For that to make a difference you must, of course, know that feature exists and how to use it. Moreover, you must know to do it for every data collection product out there. Only the savviest will revoke their work, and the rest will be preyed on. These companies will offer opt out options in the hope you won't realize they should have been opt-in in the first place.

As I said at the beginning of this essay, it's not hard to imagine a fair and equitable text-to-image generator, which makes their chosen path seem- in the most charitable interpretation- thoughtless. You would build it on a foundation of public domain and creative commons images, embellish it with images your company produces internally, commission artists to create training images for you or compensate artists who opt-in to have their images added to the data set. Maybe even give them a royalty every time their images or name are used to produce a result. I don't see anything wrong with a model like that, and I imagine it would still be fascinating and effective, but still have some utterly sane constraints that would leave much art to be made by hard working human beings. I also think you'd have a *huge amount* of artists happily opt in to adding their images to the project. We're artists! We love new, strange, and wonderful things. As much as these AI companies seem to want to lump us in with the prudish bores of the past—it's a bit of a stretch considering our demographic. Most of us use and love tech, have a long list of nerd credentials, are constantly trying new things, experimenting, and getting excited about the weird turnings of the world. I think if someone made a consent-based model like this even I would throw my images into it. But for now, please, let's stop with the go-fever, and think this through. Let's end this wholesale theft of our creative labor. I'll leave links in the description to associations trying to organize around these issues, sources, and informational videos on the

current state of the art systems. I highly encourage you to do more research into how these systems work, and how the companies that are making them are controlling the narrative.

Links:

Follow The Concept Art Association for further news on organizing against AIs, and check out their recent AI town hall video: <https://www.youtube.com/c/ConceptArtAssociation/>

Equity is fighting against AIs replacing many kinds of artists, their efforts are relevant to all creatives and they have good resources here:

<https://www.equity.org.uk/getting-involved/campaigns/stop-ai-stealing-the-show/>

A site to investigate just one of the LAION data sets to see if you or your art are in there. A note that this site offers to sign you up for a future opt-out product- I have no idea what the nature of that product will be and would urge caution: <https://haveibeen trained.com/>

Open AI explaining how they invented their legal structure because nothing else worked for them: <https://openai.com/blog/openai-lp/>

Stable Diffusion release with no mention of “artists”:

<https://stability.ai/blog/stable-diffusion-announcement>

Imagen release info that explains how the data set it is trained on is “uncurated” and contains “a wide range of inappropriate content including pornographic imagery, racist slurs, and harmful social stereotypes.” This is another page you should do a ctrl-f search for “artist” on:

<https://imagen.research.google/>

California Congresswoman Anna G. Eshoo writes OSTP urging investigation into unsafe AI release models, namely Stable Diffusion:

<https://eshoo.house.gov/media/press-releases/eshoo-urges-nsa-ostp-address-unsafe-ai-practices>

A great article summarizing the data laundering techniques of AI companies:

<https://waxy.org/2022/09/ai-data-laundering-how-academic-and-nonprofit-researchers-shield-tech-companies-from-accountability/>

A good semi-technical explanation of how diffusion models like Stable Diffusion work:

<https://www.youtube.com/watch?v=1C1pzeNxIhU>

The best non-technical explanation I’ve seen on how diffusion models work:

<https://www.tiktok.com/t/ZTRHrUyDM/>

Special thanks to my wife for her help editing the script, to my Patrons for their support, to my friends on my Discord for links, expert info, and first reactions, and to those who get the course, of course.

I did the drawing for this video, "Prey For A Spiritual Creature" on 18"x24" Strathmore bristol paper.