

2022-11-17T06:00:00.000-05:00

Kevin Kelly

# What AI-Generated Art

## Really Means for Human

### Creativity

Picture Lee Unkrich, one of Pixar’s most distinguished animators, as a seventh grader. He’s staring at an image of a train locomotive on the screen of his school’s first computer. *Wow*, he thinks. Some of the magic wears off, however, when Lee learns that the image had not appeared simply by asking for “a picture of a train.” Instead, it had to be painstakingly coded and rendered—by hard-working humans.

Now picture Lee 43 years later, stumbling onto DALL-E, an artificial intelligence that generates original works of art based on human-supplied prompts that can literally be as simple as “a picture of a train.” As he types in words to create image after image, the *wow* is back. Only this time, it doesn’t go away. “It feels like a miracle,” he [says](#). “When the results appeared, my breath was taken away and tears welled in my eyes. It’s that magical.”



# ENGINES OF WOW

## THE WIRED WORD ON AI ART

- [Lessons from the History of Photography](#)

- [Classic WIRED Covers—Regenerated by AI](#)

Our machines have crossed a threshold. All our lives, we have been reassured that computers were incapable of being truly creative. Yet, suddenly, millions of people are now using a new breed of AIs to generate stunning, never-before-seen pictures. Most of these users are not, like Lee Unkrich, professional artists, and that's the point: They do not have to be. Not everyone can write, direct, and edit an Oscar winner like *Toy Story 3* or *Coco*, but everyone *can* launch an AI image generator and type in an idea. What appears on the screen is astounding in its realism and depth of detail. Thus the universal response: *Wow*. On four services alone—Midjourney, Stable Diffusion, Artbreeder, and DALL-E—humans working with AIs now cocreate more than 20 million images every day. With a paintbrush in hand, [artificial intelligence](#) has become an engine of *wow*.

## Featured Video

[President Barack Obama on How Artificial Intelligence Will Affect Jobs](#)

Because these surprise-generating AIs have learned their art from billions of pictures made by humans, their output hovers around what we expect pictures to look like. But because they are an alien AI, fundamentally mysterious even to their creators, they restructure the new pictures in a way no human is likely to think of, filling in details most of us wouldn't have the artistry to imagine, let alone the skills to execute. They can also be instructed to generate more variations of something we like, in whatever style we want —in seconds. This, ultimately, is their most powerful advantage: They can make new things that are relatable and comprehensible but, at the same time, completely unexpected.

So unexpected are these new AI-generated images, in fact, that—in the silent awe immediately following the *wow*—another thought occurs to just about everyone who has encountered them: Human-made art must now be over. Who can compete with the speed, cheapness, scale, and, yes, wild creativity of these machines? Is art yet another human pursuit we must yield to robots? And the next

obvious question: If computers can be creative, what else can they do that we were told they could not?

I have spent the past six months using AIs to create thousands of striking images, often losing a night's sleep in the unending quest to find *just one more* beauty hidden in the code. And after interviewing the creators, power users, and other early adopters of these generators, I can make a very clear prediction: Generative AI will alter how we design just about everything. Oh, and not a single human artist will lose their job because of this new technology.



It is no exaggeration to call images generated with the help of AI *cocreations*. The sobering secret of this new power is that the best applications of it are the result not of typing in a single prompt but of very long conversations between humans and machines. Progress for each image comes from many, many iterations, back-and-forths, detours, and hours, sometimes days, of teamwork—all on the back of years of advancements in machine learning.

AI image generators were born from the marriage of two separate technologies. One was a historical line of deep learning neural nets that could generate coherent realistic images, and the other was a natural language model that could serve as an interface to the image engine. The two were combined into a language-driven image generator. Researchers scraped the internet for all images that had adjacent text, such as captions, and used billions of these examples to connect visual forms to words, and words to forms. With this new combination, human users could enter a string of words—the prompt—that described the image they sought, and the prompt would generate an image based on those words.

Scientists now at Google invented the diffusion computational models that are at the core of image generators today, but the company has been so concerned about what people might do with them that it still has not opened its own experimental generators, Imagen and Parti, to the public. (Only employees can try them, and with tight guidelines on what can be requested.) It is no coincidence, then, that the three most popular platforms for image generators right now are three startups with no legacy to protect. [Midjourney](#) is a bootstrapping startup launched by David Holz, who based the generator in an emerging community of artists. The interface to the AI is a noisy Discord server; all the work and prompts were made public from the start. [DALL-E](#) is a second-gen product of the nonprofit OpenAI, funded by Elon Musk and others. [Stable Diffusion](#) appeared on the scene in August 2022, created by Emad Mostaque, a European entrepreneur. It's an open source project, with the added benefit that anyone can download its software and run it locally on their own desktop. More than the others, Stable Diffusion has unleashed AI image generators into the wild.



ILLUSTRATION BY ADAM GARCIA



ILLUSTRATION BY: @auranova\_ai +

MIDJOURNEY

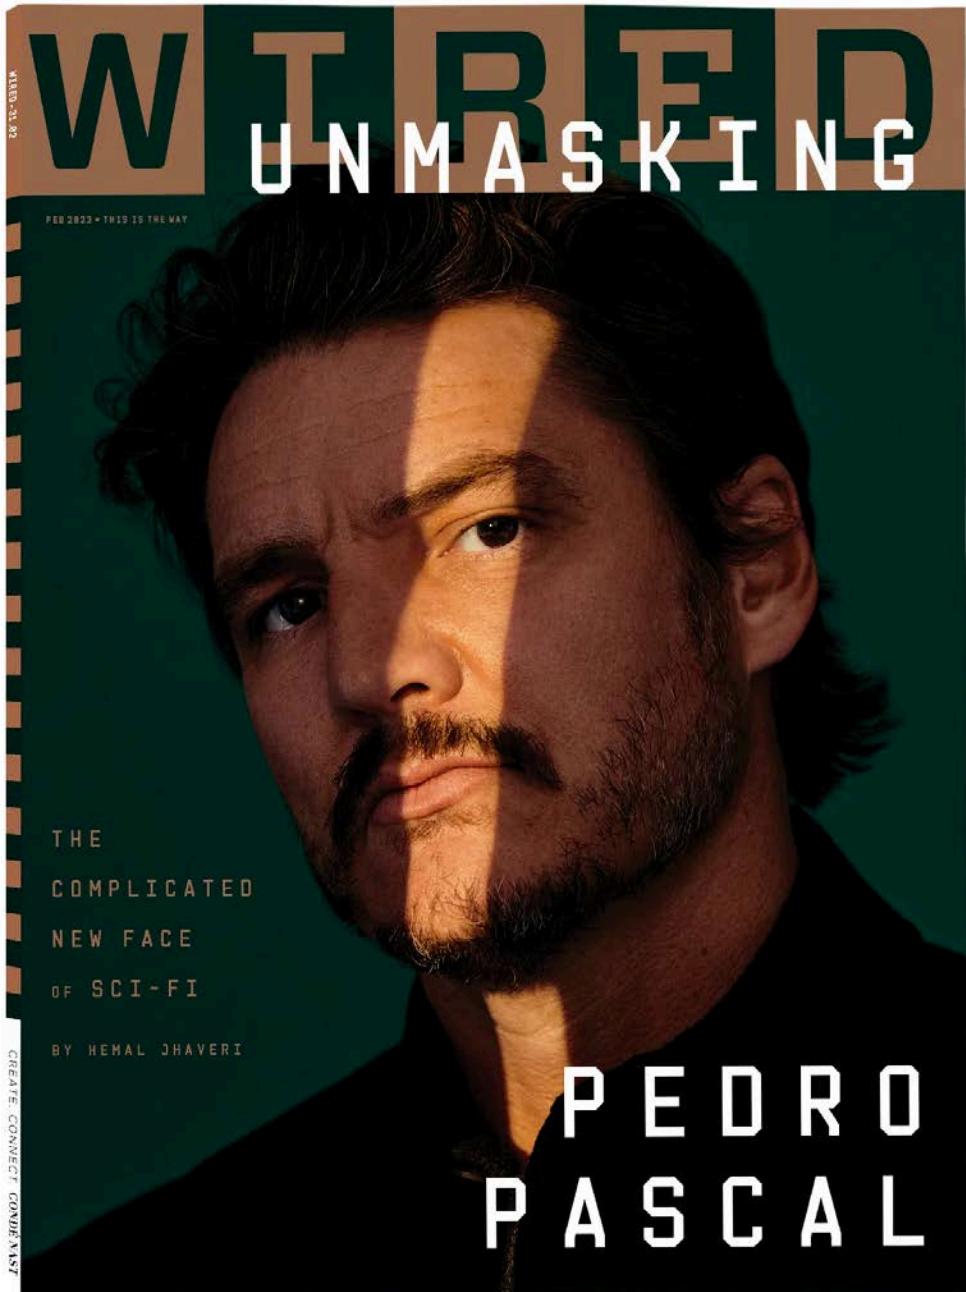
Why are so many people so excited to play with these AIs?

Many images are being created for the same reason that humans have always made most art: because the images are

pretty and we want to look at them. Like flames in a campfire, the light patterns are mesmerizing. They never repeat themselves; they surprise, again and again. They depict scenes no one has witnessed before or can even imagine, and they are expertly composed. It's a similar pleasure to exploring a video game world, or paging through an art book. There is a real beauty to their creativity, and we stare much in the way we might appreciate a great art show at a museum. In fact, viewing a parade of generated images is very much like visiting a personal museum—but in this case, the walls are full of art we ask for. And the perpetual novelty and surprise of the next image hardly wanes. Users may share the gems they discover, but my guess is that 99 percent of the 20 million images currently generated each day will only ever be viewed by a single human—their cocreator.

Like any art, the images can also be healing. People spend time making strange AI pictures for the same reason they might paint on Sundays, or scribble in a journal, or shoot a video. They use the media to work out something in their own lives, something that can't be said otherwise. I've seen

images depicting what animal heaven might look like, created in response to the death of a beloved dog. Many images explore the representation of intangible, spiritual realms, presumably as a way to think about them. “A huge portion of the entire usage is basically art therapy,” Holz, the Midjourney creator, tells me. “The images are not really aesthetically appealing in a universal sense but are appealing, in a very deep way, within the context of what’s going on in people’s lives.” The machines can be used to generate fantasies of all types. While the hosted services [prohibit porn and gore](#), anything goes on the desktop versions, as it might in Photoshop.



This article appears in the February 2023 issue. [Subscribe to WIRED](#). Photograph: Peter Yang

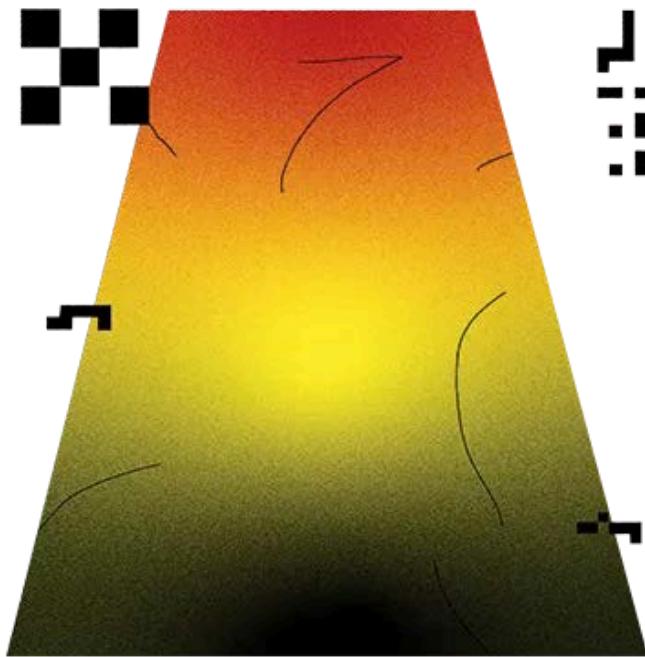
AI-generated pictures can be utilitarian too. Say you are presenting a report on the possibility of recycling hospital plastic waste into construction materials and you want an image of a house made out of test tubes. You could search

stock photo markets for a usable image made by a human artist. But a unique assignment like this rarely yields a preexisting picture, and even if found, its copyright status could be dubious or expensive. It is cheaper, faster, and probably far more appropriate to generate a unique, personalized image for your report in a few minutes that you can then insert into your slides, newsletter, or blog—and the copyright ownership is yours (for now). I have been using these generators myself to cocreate images for my own slide presentations.

In an [informal poll](#) of power users, I found that only about 40 percent of their time is spent seeking utilitarian images. Most AI images are used in places where there were no images previously. They usually do not replace an image created by a human artist. They may be created, for example, to illustrate a text-only newsletter by someone without artistic talent themselves, or the time and budget to hire someone. Just as mechanical photography did not kill human illustrations a century ago, but rather significantly expanded the places in which images appeared, so too do AI image generators open up possibilities for more art, not

less. We'll begin to see contextually generated images predominately in spaces that are currently blank, like emails, text messages, blogs, books, and social media.

This new art resides somewhere between painting and photography. It lives in a possibility space as large as painting and drawing—as huge as human imagination. But you move through the space like a photographer, hunting for discoveries. Tweaking your prompts, you may arrive at a spot no one has visited before, so you explore this area slowly, taking snapshots as you step through. The territory might be a subject, or a mood, or a style, and it might be worth returning to. The art is in the craft of finding a new area and setting yourself up there, exercising good taste and the keen eye of curation in what you capture. When photography first appeared, it seemed as if all the photographer had to do was push the button. Likewise, it seems that all a person has to do for a glorious AI image is push the button. In both cases, you get an image. But to get a great one—a truly artistic one—well, that's another matter.



Accessible AI image generators are not even a year old, but already it is evident that some people are much better at creating AI images than others. Although they're using the same programs, those who have accumulated thousands of hours with the algorithms can magically produce images that are many times better than the average person's. The images by these masters have a striking coherence and visual boldness that is normally overwhelmed by the flood of details the AIs tend to produce. That is because this is a team sport: The human artist and the machine artist are a duet. And it requires not just experience but also lots of

hours and work to produce something useful. It is as if there is a slider bar on the AI: At one end is Maximum Surprise, and at the other end Maximum Obedience. It is very easy to get the AI to surprise you. (And that is often all we ask of it.) But it is very difficult to get the AI to obey you. As Mario Klingemann, who makes his living selling NFTs of his [AI-generated artwork](#), says, “If you have a very specific image in mind, it always feels like you are up against a forcefield.” Commands like “shade this area,” “enhance this part,” and “tone it down” are obeyed reluctantly. The AIs have to be persuaded.

Current versions of DALL-E, Stable Diffusion, and Midjourney limit prompts to about the length of a long tweet. Any longer and the words muddle together; the image turns to mush. That means that behind every fabulous image lies a short magic spell that summons it. It begins with the first incantation. How you say it matters. Your immediate results materialize in a grid of four to nine images. From that batch of pictures, you variate and mutate offspring images. Now you have a brood. If they look promising, begin to tweak the spell to nudge it in new

directions as it births more generations of images. Multiply the group again and again as you search for the most compelling composition. Do not despair if it takes dozens of generations. Think like the AI; what does it like to hear? Whisper instructions that have worked in the past, and add them to the prompt. Repeat. Change the word order to see whether it likes that. Remember to be specific. Replicate until you have amassed a whole tribe of images that seem to have good bones and potential. Now cull out all but a select few. Be merciless. Begin outpainting the most promising images. That means asking the AI to extend the image out in certain directions beyond the current borders. Erase those portions that are not working. Suggest replacements to be done by the AI with more incantations (called inpainting). If the AI is not comprehending your hints, try spells used by others. When the AI has gone as far as it can, migrate the image to Photoshop for final tailoring. Present it as if you have done nothing, even though it is not uncommon for a distinctive image to require 50 steps.

Behind this new magecraft is the art of prompting. Each artist or designer develops a way of persuading an AI to

yield its best by evolving their prompts. Let's call these new artists AI whisperers, or prompt artists, or promptors. The promptors work almost as directors, guiding the work of their alien collaborators toward a unified vision. The convoluted process required to tease a first-rate picture out of an AI is quickly emerging as a fine-art skill. Almost daily, new tools arrive to make prompting easier, better.

[PromptBase](#) is a market for promptors to sell prompts that create simple images such as emoticons, logos, icons, avatars, and game weapons. It's like clip art, but instead of selling the art, they sell the prompt that generates the art. And unlike fixed clip art, it is easy to alter and tweak the art to fit your needs, and you can extract multiple versions again and again. Most of these prompts sell for a couple bucks, which is a fair price, given how much trouble it is to hone a prompt on your own.

Above-average prompts not only include the subject but also describe the lighting, the point of view, the emotion evoked, the color palette, the degree of abstraction, and perhaps a reference picture to imitate. As with other artistic skills, there are now courses and guidebooks to train the

budding promptor in the finer points of prompting. One fan of DALL-E 2, Guy Parsons, put together a free [Prompt Book](#), jammed with tips on how to go beyond the *wow* and get images you can actually use. One example: If your prompt includes specific terms such as “Sigma 75 mm camera lens,” Parson says, then the AI doesn’t just create that specific look made by the lens; “it more broadly alludes to ‘the kind of photo where the lens appears in the description,’” which tends to be more professional and therefore yields higher-quality images. It’s this kind of multilevel mastery that produces spectacular results.

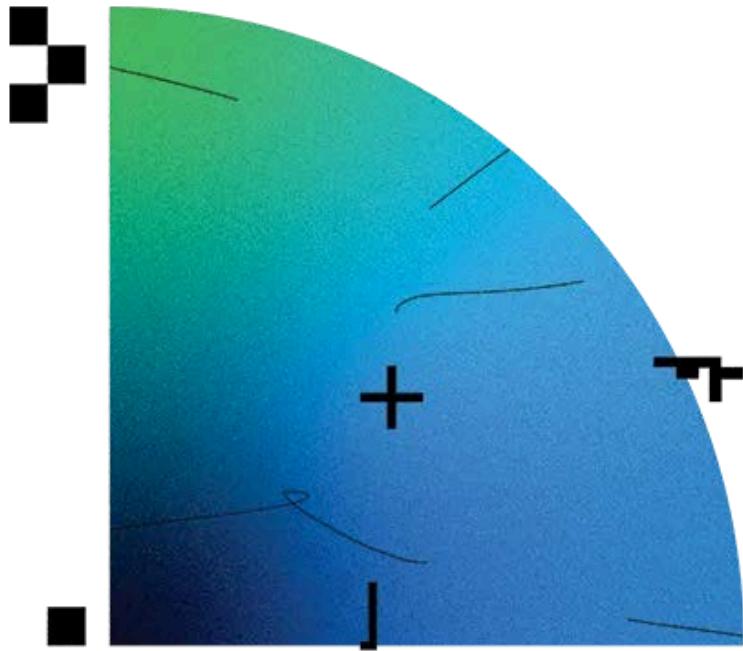
For technical reasons, even if you repeat the exact same prompt, you are unlikely to get the same image. There is a randomly generated seed for each image, without which it is statistically impossible to replicate. Additionally, the same prompt given to different AI engines produces different images—Midjourney’s are more painterly, while DALL-E is optimized for photographic realism. Still, not every promptor wishes to share their secrets. The natural reaction upon seeing a particularly brilliant image is to ask, “What spell did you use?” What was the prompt? Robyn

Miller, cocreator of the legendary game *Myst* and a pioneering digital artist, has been posting an AI-generated image every day. “When people ask me what prompt I used,” he says, “I have been surprised that I don’t want to tell them. There is an art to this, and that has also surprised me.” Klingemann is famous for not sharing his prompts. “I believe all images already exist,” he says. “You don’t make them, you find them. If you get somewhere by clever prompting, I do not see why I want to invite everybody else there.”

It seems obvious to me that promptors are making true art. What is a consummate movie director—like Hitchcock, like Kurosawa—but a promptor of actors, actions, scenes, ideas? Good image-generator promptors are engaged in a similar craft, and it is no stretch for them to try and sell their creations in art galleries or enter them into art contests. This summer, Jason Allen won first place in the digital art category at the Colorado State Fair Fine Art competition for a large, space-opera-themed canvas that was signed “Jason Allen via Midjourney.” It’s a pretty cool picture that would’ve taken some effort to make no matter

what tools were used. Usually images in the digital art category are created using Photoshop and Blender-type tools that enable the artist to dip into libraries of digitized objects, textures, and parts, which are then collaged together to form the scene. They are not drawn; these digital images are unapologetically technological assemblages. Collages are a venerable art form, and using AI to breed a collage is a natural evolution. If a 3D-rendered collage is art, then a Midjourney picture is art. As Allen [told Vice](#), “I have been exploring a special prompt. I have created hundreds of images using it, and after many weeks of fine-tuning and curating my gens, I chose my top 3 and had them printed on canvas.”

Of course, Allen’s blue ribbon set off alarm bells. To some critics, this was a sign of the end times, the end of art, the end of human artists. Predictable lamentations ensued, with many pointing out how unfair it felt for struggling artists. The AIs are not only going to take over and kill us all—they are, apparently, going to make the world’s best art while doing so.



At its birth, every new technology ignites a Tech Panic Cycle. There are seven phases:

1. Don't bother me with this nonsense. It will never work.
2. OK, it is happening, but it's dangerous, 'cause it doesn't work well.
3. Wait, it works too well. We need to hobble it. Do something!
4. This stuff is so powerful that it's not fair to those without access to it.
5. Now it's everywhere, and there is no way to escape it. Not fair.
6. I am going to give it up. For a month.

## 7. Let's focus on the real problem—which is the next current thing.

Today, in the case of AI image generators, an emerging band of very tech-savvy artists and photographers are working out of a Level 3 panic. In a reactive, third-person, hypothetical way, they fear other people (but never themselves) might lose their jobs. Getty Images, the premier agency selling stock photos and illustrations for design and editorial use, has already banned AI-generated images; certain artists who post their work on DeviantArt have demanded a similar ban. There are well-intentioned demands to identify AI art with a label and to segregate it from “real” art.

Beyond that, some artists want assurances that their own work not be used to train the AIs. But this is typical of Level 3 panic—in that it is, at best, misguided. The algorithms are exposed to 6 billion images with attendant text. If you are not an influential artist, removing your work makes zero difference. A generated picture will look exactly the same with or without your work in the training set. But even if you *are* an influential artist, removing your

images still won't matter. Because your style has affected the work of others—the definition of influence—your influence will remain even if your images are removed. Imagine if we removed all of Van Gogh's pictures from the training set. The style of Van Gogh would still be embedded in the vast ocean of images created by those who have imitated or been influenced by him.

Styles are summoned via prompts, as in: “in the style of Van Gogh.” Some unhappy artists would rather their names be censored and not permitted to be used as a prompt. So even if their influence can't be removed, you can't reach it because their name is off-limits. As we know from all previous attempts at censoring, these kinds of speech bans are easy to work around; you can misspell a name, or simply describe the style in words. I found, for example, that I could generate detailed black-and-white natural landscape photographs with majestic lighting and prominent foregrounds—without ever using Ansel Adams' name.

There is another motivation for an artist to remove themselves. They might fear that a big corporation will

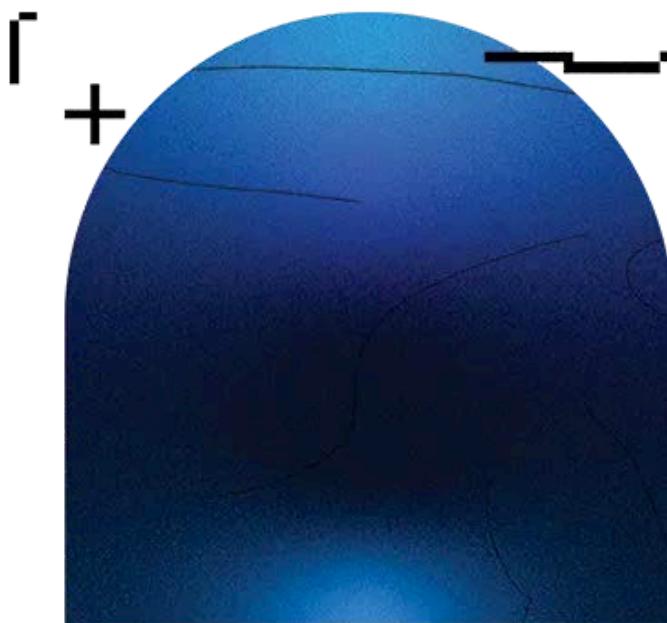
make money off of their work, and their contribution won’t be compensated. But we don’t compensate human artists for their influence on other human artists. Take David Hockney, one of the highest-paid living artists. Hockney often acknowledges the great influence other living artists have on his work. As a society, we don’t expect him (or others) to write checks to his influences, even though he could. It’s a stretch to think AIs should pay their influencers. The “tax” that successful artists pay for their success is their unpaid influence on the success of others.

What’s more, lines of influence are famously blurred, ephemeral, and imprecise. We are all influenced by everything around us, to degrees we are not aware of and certainly can’t quantify. When we write a memo or snap a picture with our phone, to what extent have we been influenced—directly or indirectly—by Ernest Hemingway or Dorothea Lange? It’s impossible to unravel our influences when we create something. It is likewise impossible to unravel the strands of influence in the AI image universe. We could theoretically construct a system to pay money earned by the AI to artists in the training set,

but we'd have to recognize that this credit would be made arbitrarily (unfairly) and that the actual compensatory amounts per artist in a pool of 6 billion shares would be so trivial as to be nonsensical.

In the coming years, the computational engine inside an AI image generator will continue to expand and improve until it becomes a central node in whatever we do visually. It will have literally seen everything and know all styles, and it will paint, imagine, and generate just about anything we need. It will become a visual search engine, and a visual encyclopedia with which to understand images, and the primary tool we use with our most important sense, our sight. Right now, every neural net algorithm running deep in the AIs relies on massive amounts of data—thus the billions of images needed to train it. But in the next decade, we'll have operational AI that relies on far fewer examples to learn, perhaps as few as 10,000. We'll teach even more powerful AI image generators how to paint by showing them thousands of carefully curated, highly selected images of existing art, and when this point comes, artists of all backgrounds will be fighting one another to be included in

the training set. If an artist is in the main pool, their influence will be shared and felt by all, while those not included must overcome the primary obstacle for any artist: not piracy, but obscurity.



As soon as 2D generative algorithms were born, experimenters rushed to figure out what was next. Jensen Huang, the ambitious cofounder of Nvidia, believes the next generation of chips will generate 3D worlds for the metaverse—“the next computing platform,” as he [calls it](#). In a single week this past September, three novel text-to-

3D/video image generators were announced: GET3D (Nvidia), Make-A-Video (Meta), and DreamFusion (Google). The expansion is happening faster than I can write. Amazing as frameable 2D pictures produced by AI are, outsourcing their creation is not going to radically change the world. We are already at peak 2D. The genuine superpower being released by AI image generators will be in producing 3D images and video.

A future prompt for a 3D engine might look something like this: “Create the messy bedroom of a teenager, with posters on the wall, an unmade bed, and afternoon sunlight streaming through closed blinds.” And in seconds, a fully realized room is born, the closet door open and all the dirty clothes on the floor—in full 3D. Then, tell the AI: “Make a 1970s kitchen with refrigerator magnets and all the cereal boxes in the pantry. In full volumetric detail. One that you could walk through. Or that could be photographed in a video.” Games crammed with alternatively rendered worlds and full-length movies decked out with costumes and sets have eternally been out of reach for individual artists, who remain under the power of large dollars. AI could make

games, metaverses, and movies as quick to produce as novels, paintings, and songs. Pixar films in an instant! Once millions of amateurs are churning out billions of movies and endless metaverses at home, they will hatch entirely new media genres—virtual tourism, spatial memes—with their own native geniuses. And when big dollars and professionals are equipped with these new tools, we’ll see masterpieces at a level of complexity never seen before.

But even the vast universes of 3D worlds and video are not vast enough to contain the disruption that AI image generators have initiated. DALL-E, Midjourney, and Stable Diffusion are just the first versions of generative machines of all types. Their prime function, pattern recognition, is almost a reflex for human brains, something we accomplish without conscious thinking. It is at the core of almost everything we do. Our thinking is more complex than just pattern recognition, of course; dozens of cognitive functions animate our brain. But this single type of cognition, synthesized in machines (and the only cognition we have synthesized so far), has taken us further than we

first thought—and will probably continue to advance further than we now think.

When an AI notices a pattern, it stores it in a compressed way. Round objects are placed in a “roundness” direction, red objects in another direction for “redness,” and so on. Maybe it notices “treeness” and “foodness” too. It abstracts out billions of directions, or patterns. Upon reflection—or training—it notices that the overlap of these four qualities produces “appleness,” yet another direction. Furthermore, it links all these noticed directions with word patterns, which can also share overlapping qualities. So when a human requests a picture of an apple via the word “apple,” the AI paints an image with those four (or more) qualities. It is not assembling bits of existing pictures; rather, it is “imagining” a new picture with the appropriate qualities. It sort of remembers a picture that does not exist but could.

This same technique can be used—in fact, is already being used, in very early forms—to find new drugs. The AI is trained on a database of all the molecules we know to be active medicines, noticing patterns in their chemical structures. Then the AI is asked to “remember” or imagine

molecules we have never thought of that seem to be similar to the molecules that work. Wonderfully, some of them actually do work, just as an AI image of a requested imaginary fruit can look remarkably like a fruit. This is the real transformation, and soon enough, the same technique will be used to help design automobiles, draft laws, write code, compose soundtracks, assemble worlds to entertain and instruct, and cocreate the stuff we do as work. We should take to heart the lessons we've learned so far from AI image generators because there will soon be more pattern-seeking AIs in all realms of life. The panic cycle we presently face is simply a good rehearsal for the coming shift.

What we know about AI generators so far is that they work best as partners. The nightmare of a rogue AI taking over is just not happening. That vision is fundamentally a misreading of history. In the past, technology has rarely directly displaced humans from work they wanted to do. For instance, the automatic generation of pictures by a machine—called a camera—was feared in the 1800s because it would surely put portrait painters out of

business. But the historian Hans Rooseboom could find only a *single* portrait painter from that time who felt unemployed by photography. (Photography actually inspired a resurgence of painting later in that century.) Closer to our time, we might have expected professional occupations in photography to fall as the smartphone swallowed the world and everybody became a photographer—with 95 million uploads to Instagram a day and counting. Yet the number of photography professionals in the US has been slowly rising, from 160,000 in 2002 (before camera phones) to 230,000 in 2021.

Instead of fearing AI, we are better served thinking about what it teaches us. And the most important thing AI image generators teach us is this: Creativity is not some supernatural force. It is something that can be synthesized, amplified, and manipulated. It turns out that we didn't need to achieve intelligence in order to hatch creativity.

Creativity is more elemental than we thought. It is independent of consciousness. We can generate creativity in something as dumb as a deep learning neural net. Massive data plus pattern recognition algorithms seems

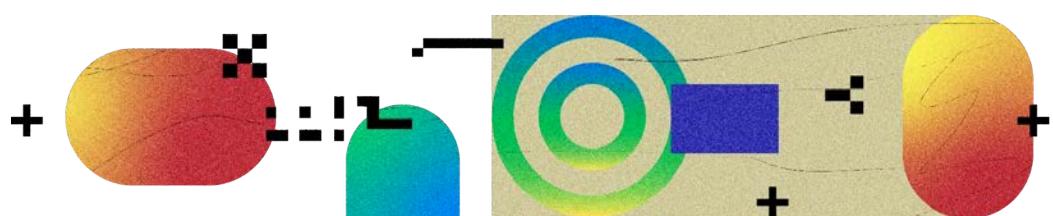
sufficient to engineer a process that will surprise and aid us without ceasing.

Scholars of creativity refer to something called Uppercase Creativity. Uppercase Creativity is the stunning, field-changing, world-altering rearrangement that a major breakthrough brings. Think special relativity, the discovery of DNA, or Picasso's *Guernica*. Uppercase Creativity goes beyond the merely new. It is special, and it is rare. It touches us humans in a profound way, far beyond what an alien AI can fathom.

To connect with a human deeply will always require a Creative human in the loop. This high creativity, however, should not be confused with the creativity that most human artists, designers, and inventors produce day to day.

Mundane, ordinary, lowercase creativity is what we get with a great new logo design or a cool book cover, a nifty digital wearable or the latest must-have fashion, or the set design for our favorite sci-fi serial. Most human art, past and present, is lowercase. And lowercase creativity is exactly what the AI generators deliver.

But this is huge. For the first time in history, humans can conjure up everyday acts of creativity on demand, in real time, at scale, for cheap. Synthetic creativity is a commodity now. Ancient philosophers will turn in their graves, but it turns out that to make creativity—to generate something new—all you need is the right code. We can insert it into tiny devices that are presently inert, or we can apply creativity to large statistical models, or embed creativity in drug discovery routines. What else can we use synthetic creativity for? We may feel a little bit like medieval peasants who are being asked, “What would you do if you had the power of 250 horses at your fingertips?” We dunno. It’s an extraordinary gift. What we do know is we now have easy engines of creativity, which we can aim into stale corners that have never seen novelty, innovation, or the wow of creative change. Against the background of everything that breaks down, this superpower can help us extend the wow indefinitely. Used properly, we can make a small dent in the universe.



*This article appears in the February issue. [Subscribe now.](#)*

*Let us know what you think about this article. Submit a letter to the editor at [mail@wired.com](mailto:mail@wired.com).*

---

Viewed using [Just Read](#)