

What algorithmic art can teach us about artificial intelligence

Ceci n'est pas un algorithme

By James Vincent

The Verge Nov 6, 2018

(The essay contains several pictures.)

We live in a world that's increasingly controlled by what might be called "the algorithmic gaze." As we cede more decision-making power to machines in domains like health care, transportation, and security, the world as seen by computers becomes the dominant reality. If a facial recognition system doesn't recognize the color of your skin, for example, it won't acknowledge your existence. If a self-driving car can't see you walk across the road, it'll drive right through you. That's the algorithmic gaze in action.

This sort of slow-burning structural change can be difficult to comprehend. But as is so often the case with societal shifts, artists are leaping headfirst into the epistemological fray. One of the best of these is Tom White, a lecturer in computational design at the University of Wellington in New Zealand whose art depicts the world, not as humans see it, but as algorithms do.

White started making this kind of artwork in late 2017 with a series of prints called "The Treachery of ImageNet." The name combines the title of René Magritte's famous painting of a pipe that isn't a pipe, and ImageNet, a database of pictures that's used across the industry to train and test machine vision algorithms. "It seemed like a natural parallel for me," White tells The Verge. "Plus, I can't resist a pun."

To humans, the pictures look like haphazard arrangements of lines and blobs that lack any obvious immediate structure. But to algorithms trained to see the world on our behalf, they leap off the page as specific objects: electric fans, sewing machines, and lawnmowers. The prints are optical illusions, but only computers can see the hidden image. (In the video above you can see how White develops these images, plus how Verge staffers did at guessing what they show.)

White's work has attracted a lot of attention in the machine learning community, and it's getting its first major gallery show this month as part of an exhibition of AI artwork in India at Delhi's Nature Morte gallery. White says he designs his prints to "see the world through the eyes of a machine" and make "a voice for the machine to speak in."

That "voice" is actually a series of algorithms that White has dubbed his "Perception Engines." They take the data that machine vision algorithms are trained on — databases of thousands of pictures of objects — and distill it into abstract shapes. These shapes are then fed back into the same algorithms to see if they're recognized. If not, the image is tweaked and sent back, again and again, until it is. It's a trial and error process that essentially ends up reverse-engineering the algorithm's understanding of the world.

White compares the process to a "computational ouija board," where neural networks "simultaneously nudge and push a drawing toward the objective." He tells The Verge that this method gives him the control he wants out of the output, though it can take days to create a single image in this way, and he admits that the process is "kind of tedious."

Unlike some artists who work with machine learning, White doesn't pretend that his prints are the product of a some autonomous AI (a disingenuous narrative sometimes pushed by artists and promoters in order to create a feeling of technological mysticism). Instead, he's up front about his role: he sets a number of starting parameters for his perception engines, like the colors and thickness of lines, and winnows the output, rejecting prints that he doesn't find aesthetically pleasing. Although he is giving his algorithms a voice to speak in, he's also making sure the results are pleasant to hear. "I think I am trying to free the algorithm so it can express itself, so people can relate to what it's saying," he says.

And what is it saying? Well, as with any art, different people hear different things.

Some see the imagery made by White and his peers as a bad omen, another sign that artificial intelligence is not only getting smarter but beginning to think creatively and take on roles reserved for humans. Karthik Kalyanaraman, one half of the curation team responsible for the Nature Morte exhibition, tells *The Verge* by email that he arranged the show to draw attention to the “inevitable” questions we face about the future of humanity. “Once so much of our labor (manual, mental, emotional, artistic) is replaced by machines, what is left for us to do?” he asks. “How will we define ourselves?”

Kalyanaraman suggests that art made with AI demonstrates that computers may deserve credit as creative actors. The type of machine learning used by White and his peers works by sifting through large amounts of data and then replicating the patterns it finds. Kalyanaraman suggests that this is similar to the process by which humans learn art, but that our “mysticism” surrounding the notion of creativity stops us from seeing the parallels. “If a machine can make humanly surprising, stylistically new kinds of art, I think it is foolish to say well it’s not really creative because it doesn’t have consciousness,” he says.

Others frame the question in more ruthless economic terms. Writing for contemporary art magazine *frieze*, Mike Pepi suggests the promotion of AI creativity is essentially propaganda for corporate interests. Pepi says that despite “utopian prognostication,” the development of artificial intelligence is ultimately about replacing human labor, including white-collar jobs that need creative abilities. Says Pepi: “If machine intelligence can conquer this uniquely human realm, the march to artificial general intelligence must be nigh, and the profits unimaginable.”

White says his motivation is primarily to deconstruct what we think of as machine perception. In other words: to explain the algorithmic gaze. Take the example of the cello print in White’s series “The Treachery of ImageNet.” If you know what you’re looking for, you can see shapes that represent the instrument (a cluster of straight parallel lines bracketed by curves). But there’s also a confusing shape looming behind it. White says these shapes are there because the algorithms were trained using pictures of cellos with cellists holding them. Because the algorithm has no prior knowledge of the world — no understanding of what an instrument is or any concept of music or performance — it naturally grouped the two together. After all, that’s what it’s been asked to do: learn what’s in the picture.

This sort of mistake is common in machine learning, and it demonstrates a number of important lessons. It shows how critical training data is: give an AI system the wrong data to learn from, and it’ll learn the wrong thing. It also demonstrates that no matter how “clever” these systems seem, they possess a brittle intelligence that only understands a slice of the world — and even that, imperfectly. White’s latest prints for the Nature Morte gallery, for example, are abstract smears of color designed to be flagged as “inappropriate content” by Google’s algorithms. The same algorithms used to filter what humans see around the world.

Still, White says that he doesn’t see his artwork as a warning. “I’m just trying to present the algorithms as they are,” he says. “But I admit it’s sometimes alarming that these machines we’re relying on have such a different take on how objects in the world are grounded.”

And despite the error-prone nature of algorithmic gaze, it can also do very beneficial things. Machine vision could make the world a safer place by steering cars safely on roads or save lives by speeding up medical diagnoses. But if we really want to use this technology for good, we need to understand it better. Looking at the world through an algorithm’s eyes might be the first step.