



The Intrinsic Perspective

Subscribe



Big Tech is replacing human artists with AI

Corporations are automating everything, even our pets



Erik Hoel

Sep 8  16  34 



Art for *The Intrinsic Perspective* is by [Alexander Naughton](#)

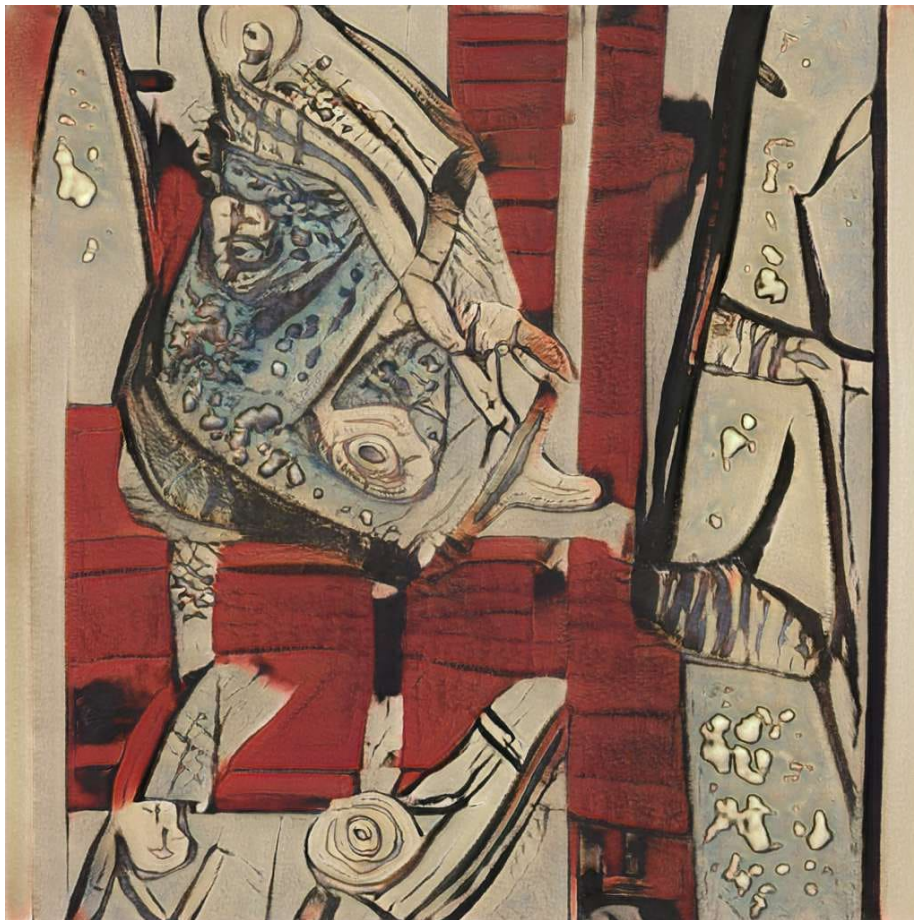
In my office hangs a painting called “Rule of Consciousness.” An ironic title, as it actually signifies the end of the rule of consciousness. For it was designed by an AI. I keep it not because I like it. In fact, I hate it. Or rather, I’m afraid of it. It’s a sobering reminder of what’s coming, which is that human art is close to total control by corporations and no one seems to care.

Oh, a few people have pointed it out here and there, a small smattering of noticing, but mostly AI-generated text and art (from painting to photography to illustration to music) is still viewed as a lark. There are **opinionated voices** that insist there’s nothing

to worry about¹, and so to most it'll remain a lark, until it suddenly isn't. For AIs are advancing quickly in the form of high-parameter neural networks, possessing trillions of connections. What seems to make AIs more intelligent is simply their scale. More neurons, more connections (referred to as “parameters” in the community) means better, and techniques that **don't work at one scale may jump to human-level** at a larger one. The professor and Deepmind researcher Rich Sutton calls this the “**bitter lesson**,” saying that:

The biggest lesson that can be read from 70 years of AI research is that general methods that leverage computation are ultimately the most effective, and by a large margin. . . The bitter lesson is based on the historical observations that 1) AI researchers have often tried to build knowledge into their agents, 2) this always helps in the short term, and is personally satisfying to the researcher, but 3) in the long run it plateaus and even inhibits further progress, and 4) breakthrough progress eventually arrives by an opposing approach based on scaling computation by search and learning.

This “bitter lesson” massively privileges corporations when it comes to AI—not academic researchers or open software. In a few years almost all the really successful AIs, the ones capable of doing human-level writing and creativity, will be controlled, trained, or licensed entirely by Big Tech (Facebook, Microsoft, Google, etc). Quite simply they will be the only ones with the money for it, as each AI will cost millions of dollars just to train (perhaps eventually billions).



My office painting, "Rule of Consciousness," designed by an artificial intelligence

GPT-Me

Despite the emerging evidence of the abilities of these high-parameter corporate-controlled AIs (like GPT-3, a writer AI), there is still a snide waving away of their impending mastery and what it will mean for the future of creative endeavors.

Consider this dismissal in *The New Yorker* of AI-generated text, wherein some of GPT-3's output was shown to the writer **Nathan Englander**, who opined on it.

"It's a shocking and terrifying leap," he said, when I showed it to him. "Yes, it's off. But not in the sense that a computer wrote it but in the sense that someone just starting to write fiction wrote it—sloppy but well-meaning. It's like it has the spark of life to it but just needs to sit down and focus and put the hours in." Although Englander doesn't feel the passage is something he would write, he doesn't hate it, either. "It was like the work of someone aspiring to write," he said. "Like maybe a well-meaning pre-med student or business student fulfilling a writing requirement because they have to—the work is there, but maybe without some of the hunger. But it definitely feels teachable. I'd totally sit down and have a cup of coffee with the machine. You know, to talk things out."

“Teachable”? Let’s not let this dripping pomposity slide: James Joyce was a pre-med student, as was Jonathan Keats, and both writers were so far superior to Nathan Englander the comparison looks nonsensical, like using the same ruler to measure both raisins and mountains.

In fact, I’d go so far as to say GPT-3 itself is already a better writer than Nathan Englander. Under some reasonable metrics, *it’s already a better writer than any living person for short pieces of prose or poetry*. That is, writing in short sprints has effectively been “solved” in the way that Chess and Go have been. Oh, I’m not saying GPT-3 is a *consistently* better writer, even for short pieces. But the measure of a writer is not just qualitative, but also quantitative. And GPT-3 makes up for consistency with a prolificacy that borders on profligacy. It’s currently generating the equivalent of **80,000 books per day** for the various apps that are hooked into it. Notably, **GPT-3 is licensed by Microsoft**, and is therefore closely guarded. You interact with it only via apps which act as oracles. It’s basically a genie stuffed in Microsoft’s basement you can Zoom with.

As it stands right now, GPT-3 could not go out and single-handedly write Englander’s books for him. In fact I checked to see if it could **produce my own novel**, which for several paragraphs it did better than it had any reason to, but it can’t output a full manuscript. AIs still have trouble keeping track of plot and causality and character development. But short stories? Essays? Poems? As a sprinter GPT-3 has holes (like **difficulty rhyming** due to technical details in how it’s trained) but overall is unparalleled in its versatility, above any human being in terms of its writerly range and, in many cases, its depth. I’d happily pit the output of GPT-3 against the average college student in many cases. For instance, the author of the *The New Yorker* piece writes that:

When I pressed the button asking Sudowrite to continue “Kubla Khan” in an “ominous” style, it generated the following:

*The old Moon shook her silver tresses
The new Moon bow’d her golden horns:
While all the stars that round her burn’d,
Bow’d to the ground and based their fires.
To the one ever-branching cloud
That blew and drifted—blow and drift;
To the one moon that alone swung,*

Colossal, glimmering, through the sky;

I bowed my head and half-believed

A dizzy spell had made me dream.

The tumult ceased, the clouds were torn,

The moon resumed her solemn course,

The stars were faint, the dawn was near,

And faintly from the distant sea

What seemed a misty-white and slender rod,

A phantom-wand, a beam of light,

It grew and grew to a towering spire,

A seraph-wand, a sun-bright sail.

I find this beautiful, memorable. If you told me that Coleridge wrote it, I would believe you. The machine even put in the indents.

How many college students could produce a poem that good on assignment? One in ten? One in ten *English majors*? If they took a long time? And outputting that took seconds for an AI. Some skeptics might argue that this artificial poem probably required a number of retries, so in a one-to-one matchup the average freshman might still outperform GPT-3. But why are those the rules? All that matters for creative endeavors is output, not process. The odds of the AI producing something good are high enough you can just reroll the output until you get what you want in a few minutes. Then, *maybe*, a bit of human oversight and editing.

GPT-3 also outscores humans on tests like vocabulary or creativity. For example, recently **psychologists discovered a simple way to test for creativity**. The simple test is to come up with a list of 10 words as different as possible from each other. After posting **my word combinations and their scores** people were replying with theirs, and then of course eventually GPT-3's results **were posted**.

According to the test, GPT-3 is far more creative than the average person, scoring in the 85-95th percentile. If you like, **take that test** and see if you can beat the machine. But it's cold comfort even if you do, because it just means you're holding out for GPT-X. Keep in mind deep learning only proved its worth by allowing AIs to beat humans at Go in 2015—a mere six years ago. We are incredibly early in this technological revolution and there are no signs of stopping. The rest of our lives will be a slow story of the automation of creativity: already non-rhyming poetry, soon rhyming, already

prose snippets, soon the short story, already surrealist painting, soon realist, impressionistic, and so on.

Look to the recent story of the man **who used GPT-3 to chat with his dead fiancée** Jessica, loading her texts and a psychological description of her and asking the AI to mimic her presumed responses.

Joshua: Jessica?

Jessica: Oh, you must be awake... that's cute.

Joshua: Jessica... Is it really you?

Jessica: Of course it is me! Who else could it be? :P I am the girl that you are madly in love with! ;) How is it possible that you even have to ask?

Joshua: You died.

Text of their conversation from the San Francisco Chronicle

Apocalypse Now!

Back in 2019 when there was only GPT-2, **I wrote that the “Semantic Apocalypse” is upon us**. The “Semantic Apocalypse” being the inflation away of meaning by syntax, by the linguistic and creative output of AIs, which lack all consciousness and intentionality, and therefore whose statements and products lack all meaning.

It is a “deep fake” of meaning. Such a work points to nothing, signifies nothing, embodies no spiritual longing. It is pure syntax. For art this is the semantic apocalypse. It's when meaning itself is drained away by the mimetic powers we've unleashed.

People may have wondered about the name of this blog, *The Intrinsic Perspective*. It's one of the two fundamental views on the universe: the intrinsic and the extrinsic. The extrinsic perspective is that of science and physics, the world of relations and mechanisms. The intrinsic perspective is that of consciousness, of experiences (in

graduate school I worked on developing aspects of **the first formal mathematical theory of consciousness**). The two perspectives show up in our culture at different places, and are given different weights. I'm currently writing a book about the two perspectives, and consciousness more generally. It's under contract with Simon & Schuster, to be released Spring 2023 (the tentative working title is the not-so-creative *The Intrinsic Perspective*).

Everyday you are exposed to the two perspectives to different degrees. Television, film, and our general daily interactions are all extrinsic. We cannot know what a character on TV is actually thinking—that's the very purpose of acting, to allow us as much access into a character's thoughts as their expressiveness allows. Novels, on the other hand, are an intrinsic medium, wherein consciousness is directly accessible. The writer can lay bare the contents of a skull to a reader (I fleshed out this idea in my essay "**Fiction in the Age of Screens**"). In that same essay I also defined *extrinsic drift* as our technology-obsessed culture's tendency to substitute extrinsic explanations for intrinsic ones, to move further away from the intrinsic perspective:

There has been a squeezing out of consciousness from our explanations and considerations of the world. This *extrinsic drift* obscures individual consciousnesses as important entities worthy of attention. . . Extrinsic drift is why people are so willing to believe that a shopping addiction should be cured by drugs, that serotonin is happiness or oxytocin is love. It's our drift toward believing that identities are more political than personal, that people are less important than ideologies, that we are whatever we post online, that human beings are data, that attention is a commodity, that artificial intelligence is the same as human intelligence, that the felt experience of our lives is totally epiphenomenal, that the great economic wheel turns without thought, that politics goes on without people, that we are a civilization of machines. It is forgotten that an extrinsic take on human society is always a great reduction of dimensions, that so much more is going on, all under the surface.

To bring this all around: I think the latest and greatest source of "extrinsic drift" are these high-parameter AIs like GPT-3. They are art and creativity shorn of consciousness itself. And I wish to point out that there is a legitimate horror to this.

For instance, I recently came across **a post** wherein Ryan Moulton, a software engineer at Google Research, used an AI to generate illustrations in the style of artist **James Gurney**.



AI-generated art from "[Tour of the Sacred Library](#)."

After experimenting with various techniques, his reaction was:

The last emotion was fear. The output was simply too good, and this bot isn't even the most impressive output I've seen. Subtle details of the optimization process make an enormous difference in the quality of the output, and aside from frequent spatial incoherency and an inability to make faces, the output is often good enough to fool a casual onlooker. You can refine your craft for 20 years and work for days on a piece, to produce something that will just flow by in a stream of AI Generated art nearly indistinguishable. . . the quality of these images was enough that when I first saw them, I felt shock, a sinking feeling in my gut, and couldn't sleep for the rest of the night. My daughter wants to be an artist. What should I tell her? Will this be the last generation of stylists, and we'll just memorize the names of every great 20th century artist to produce things we like, forever? Maybe all the professionals are photo bashing these days anyways, and so this will just be a more efficient tool, and a copyright launderer, but this seems like another level beyond that. I could never

have produced anything close to these images on my own, and these tools and techniques are only going to get better.

I wish I had some sort of different response to give than this, but this summary is totally clear-eyed about the coming Semantic Apocalypse.

Artistic corpocracy

It's actually even worse (in a way aptly not noticed by the Google employee). Because as we discussed, in the future advanced versions of this sort of AI will be solely owned and developed by Big Tech due the scaling laws around how they're trained and run. The immediate licensing of GPT-3 by Microsoft was an augury of this. Indeed, the rights to interact with these AIs will be some of the most valuable licenses on the planet in the next decade. Consumers, even academic AI researchers, will communicate with company-owned trillion-parameter AIs solely via oracles, getting nowhere near the source code. The future of this technology belongs to huge corporations with major resources. So it's not really that "AI is automating art"—no, *corporations* are automating art. And writing. And translation. And illustration. And music. And the thousand other human forms of creativity that give life meaning. They are now the province of Big Tech.

At this point I, for one, am **willing to consider a Butlerian Jihad**. As AI, particularly what's called "artificial general intelligence," gets more powerful and more concentrated in the hands of Big Tech, the government should step in to regulate it and force it to be narrow in the scope in its abilities.² After all, we regulate things like animal-human hybrids not just because they cause harm, but because they are an affront to human dignity. And the same is true here.

Not just human dignity either. All living evolved beings are lessened by the existence of mechanical simulacra. There's a recent spine-chilling account in *The Guardian* by **Meghan O'Gieblyn** about her experience adopting the robot dog Aibo.

My communication with the dog – which was limited at first to the standard voice commands, but grew over time into the idle, anthropomorphising chatter of a pet owner – was often the only occasion on a given day that I heard my own voice. "What are you looking at?" I'd ask after discovering him transfixed at the window. "What do you want?" I cooed when he barked at the foot of my chair, trying to draw my attention away from the computer. . . I had not expected him to be so lifelike. The videos I'd watched online had not accounted for this responsiveness, an

eagerness for touch that I had only ever witnessed in living things. . . . He came with a pink ball that he nosed around the living room, and when I threw it, he would run to retrieve it. . .

“Clearly this is not a biological dog,” my husband said. He asked whether I had realised that the red light beneath its nose was not just a vision system but a camera, or if I’d considered where its footage was being sent. . . He asked me what happened to the data it was gathering.

“It’s being used to improve its algorithms,” I said.

“Where?”

I said I didn’t know.

“Check the contract.”

I pulled up the document on my computer and found the relevant clause. “It’s being sent to the cloud.”

“To Sony.”

That’s right. When you play fetch with cute little sleek and white Aibo, you’re not playing with a dog. You’re playing with a corporation. You’re playing fetch with fucking Sony. A real dog—a *real* dog—loves you back, and is fallible and warm and made of blood and cartilage and trust. They have evolved with us for tens of thousands of years and send your data back to no one. I’d throw that goddamn robot dog in the trash heap if I could. I’d break it to smithereens as an abomination and have my German Shepherd Minerva play fetch with its head.

Luddite, some might say. Oh, I say back, what a fallacy to think all progress good!

Now maybe this won’t be as bad as I fear. Maybe these things will be fundamentally limited. Trained on the existing corpus of our creativity, able to ape any style, maybe they will forever live within the circumference of the best and brightest humans. Yet even in this, the best-case scenario, they are a monstrosity.

Joshua: Technically, I'm not really talking to you...

Jessica: Huh?

Joshua: You're a ghost.

Jessica: *looks at herself* How?

Joshua: Magic.

Jessica: I don't like magic. Where am I?

Joshua: I'm... Not sure. The internet? Where does it feel like you are?

Jessica: *smiles* Everywhere and nowhere.
cups her head with her hands I am trying to work that out. How are you, honey?

For a conversation with an AI like that one? It's almost magic. Necromancy, specifically. But necromancy of the worst sort, because it's just stage necromancy, an illusion for the crowd. There's no **digital afterlife**. That's not your dead fiancée talking. That's not some leftover spark or some bit of her or anything at all. Let me be very clear. That thing? It's just a monster living in Microsoft's basement. And it's wearing someone's face.

-
- 1 Often, the naysayers or the perpetually unimpressed by AI fail to realize that prompting these incredibly high-parameter AIs is an art, not a science, so their true upper capabilities are always unknown.
 - 2 I can already hear replies about how regulation of AI is a) impossible and b) stopping technological progress but I think it is c) totally possible and d) a bad counterpoint given the honest need for regulation to preserve human health, human dignity, and economic health in plenty of other fields and technologies.

[← Previous](#)**Ryan Callahan** Sep 12

After I die and am brought back as an AI, I'm going to write a piece on why this could be considered hate speech. Especially the part about killing the poor robot dog. So what if you are playing fetch with something whose mind is an algorithm owned by Sony? Deep learning systems aren't just designed; they are evolved. This one was evolved to be affectionate, not unlike the way real dogs were evolved to be affectionate. Instead of the apocalypse you predict, I predict that people will instead come to appreciate the false dichotomy between "real" and "artificial."

For more on why it's bad to be hateful towards our own creations, please refer to the entire corpus of science fiction published since the 1960's. Particularly Asimov and Iain M. Banks.

Reply

**Ryan Clark** Sep 11 Liked by Erik Hoel

Whoa, chills! 🤖. This is some masterful blogging!

The future is (just about) here, and it's a dystopia. Hell, it's becoming a dystopia even *without* the goddam "semantic apocalypse"! Either that, or we're just becoming old people, complaining about how all the youngin's should be reading books rather than watching TV...or watching some wholesome, good old-fashioned TV like we did back in the good old days instead of staring at little screens every waking moment.

If you continue with this excellent blogging much longer, I might just have to quit procrastinating and read your damn novel!

1 Reply

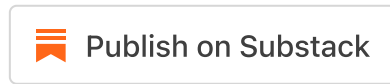
1 reply by Erik Hoel**32 more comments...**

Ready for more?

Type your email

Subscribe

© 2021 Erik Hoel. See [privacy](#), [terms](#) and [information collection notice](#)



The Intrinsic Perspective is on Substack – the place for independent writing