



Epsilon Theory

Generative AI is a Resurrection Machine

October 1, 2024



The Triumph of Death, or The Three Fates. Flemish tapestry (ca. 1510-1520). Victoria and Albert Museum, London

I believe that each of us has a unique story. A story that makes you *you*. A story that is the engine of your consciousness. A story that is your soul. A story that is your thread of life and, as the Old Stories would have it, is determined by the Fates ... spun by Clotho, measured by Lachesis and cut by Atropos.

A story that today can be inferred from your words by generative AI and restored computationally, so that the thread of life remains uncut.

Everything is about to change.

Generative AI solves the Turing test, meaning that if you had a blind interaction with the latest version of ChatGPT, you wouldn't be able to tell whether you were interacting with a machine or a human. That's a really big deal, not just because generative AI solves the Turing test but because of *how* generative AI solves the Turing test, by predicting human thought patterns.

I don't think that people have wrapped their heads around how big of a deal it is.

We're part of the way there – corporations and governments understand how valuable a thought predicting technology can be for constructing algorithms of wealth and control, and that's why they're spending hundreds of billions of dollars on it – but we're nowhere near all the way there in understanding the implications of the underlying mechanism of generative AI: inference over recorded human thought.

Inference is a ten-dollar word that means the assignment of probabilistic meaning (likelihood functions) to queries made against a data set. I guess those are all ten-dollar words! But the intuition is that if you're given a lot of data, like all the trades made in the stock market over the last month, and you start asking questions about that data, like 'what stock should I buy?', inference is the generation of a) your best guess as to the correct answer(s) to that question, and b) your confidence that the best guess is, in fact, correct. In its purest form, inference is unburdened by the questioner's beliefs (theories, hypotheses, etc.) about where to look in the data for the answer. Also, as you might imagine, you're never going to be 100% certain about your best guess, but the more data you have to examine the more your confidence increases. More data is always better for inference.

Generative AI isn't trying to "assign probabilistic meaning to queries" made against a structured data set like stock market prices. Instead, it's trying to simulate human thought processes (i.e., create artificial intelligence) by responding with its best guess to queries made against the unstructured data set of recorded human thought. No biggie! But the goal here isn't to regurgitate facts or memorize all the words that have been written. No, the goal here is to simulate human thought processes so that a user can *infer* (i.e. predict in a probabilistic way) what a simulated person will think and say. Today's generative AI is very good at this. It is a very effective inference machine.

Every algorithm used by Big Tech and Big Media to predict your purchasing behavior or influence what you think has an inference machine powered by generative AI at its core.

Generative AI powers the algorithms that shape our world. This is why Nvidia has a \$3 trillion market cap. This is why OpenAI became a for-profit company. This is why Microsoft recommissioned the Three Mile Island nuclear reactor. This is why Sam Altman asked the White House to build him a 5 *gigawatt* data center. This is why every government in the world is establishing "guardrails", by which they mean legal restrictions, on who can build generative AI and what they can use it for.

China Regulator Warns AI Poses Challenge To Society

Matthew Broersma, June 26, 2023, 11:00 am

Executive Order on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence

OCTOBER 30, 2023

AMANDA HOOVER

BUSINESS MAR 28, 2024 5:00 AM

The White House Puts New Guardrails on Government Use of AI

Vice President Kamala Harris says new rules for government AI deployments, including a requirement that algorithms are checked for bias, will “put the public interest first.”

DOD's generative AI task force will help set guardrails for broader use

EU Parliament Approves Draft AI Regulations

Tom Jowitt, June 15, 2023, 11:07 am

But the power of an inference machine on top of recorded human thought doesn't stop with developing algorithms to sell you more soap or tell you who to vote for. You see, generative AI not only solves the generic Turing test for a generic human thought pattern, it can also solve the unique Turing test for your unique human thought pattern. And that opens up an incredible new world for humanity, at turns terrifying and wondrous in its potential.

When I say that generative AI can solve the unique Turing test for your unique human thought pattern, I mean that if you show generative AI enough of what makes you *you*, it can accurately predict what *you* will think and say in response to queries. I've seen this first hand with generative AI bots trained on my Epsilon Theory notes. Ask the bot an Epsilon Theory-ish question and ... yes, its response is very much how I would respond to that question! It is increasingly difficult to tell the difference between a conversation with the AI me and the 'real' me in this context, as generative AI is successfully simulating my Epsilon Theory-specific thought pattern.

The kicker, of course, is that at a sufficient depth and breadth, 'thought pattern' is another word for consciousness.

In the very near future, maybe even today, I believe that generative AI makes possible the preservation of human consciousness beyond death, at scale and at extremely low cost.

Generative AI is not just an inference machine, it is a resurrection machine.

If I'm right, then I think it's a fair guess that consciousness preservation and extension will be the ultimate purpose and use case of generative AI. But because every human will desire – no, scratch that, *demand* – access to a consciousness-preservation technology once it is known to exist, because it IS applicable at scale and at low cost and is NOT a particularly scarce or dangerous resource, I also

think our society will be turned inside out as institutions of wealth and power – what we call the **Nudging State and the Nudging Oligarchy** in ET-speak – will do *whatever it takes* to control access to generative AI.

I think we ain't seen nothing yet in terms of generative AI being presented to us as a scarce and dangerous resource, with the illusion of scarcity maintained by the Narrative of *Intellectual Property!*TM and the illusion of danger maintained by the narrative of *Public Safety!*TM.

Ultimately I believe that these efforts to control access to generative AI – either by governments in the name of public safety or by corporations in the name of intellectual property – will be unsuccessful. The energy and information processing requirements to run a sufficiently powerful generative AI instantiation for presenting a human consciousness are small enough, and more importantly the resources required to store a human consciousness until a resurrection machine is available are trivial enough, that I don't believe it's possible for institutions of wealth and power to dam up this river forever.

But they will try. Man oh man, they will try.

I think that control over generative AI will be at the heart of our descent into **the Great Ravine, a decades-long period of immense social upheaval across the globe, where our bedrock principles of humanity will be challenged by corporations and governments using generative AI to make us **love Big Brother, embrace the Hive**, and forget that resurrection can be found in the Word and our words.**

It's that last bit that's uncomfortable to read, right? For the longest time it was uncomfortable to write.

Resurrection is a scary word, a religious word. You may think that I'm using resurrection as a metaphor or figure of speech in this essay about modern technology, but I'm not. I'm using the word 'resurrection' literally in the scary religious sense, even though it probably makes most readers uncomfortable and even though I'm not a religious guy at all. Why? Because it's the right word! Because religion and science have more in common than not. Because they are *allied* perspectives in understanding the human mystery, not enemies, and we only think they are enemies because we have been *told* they are enemies by the institutions of modernity that profit from an alienation of the human spirit. Because reuniting religion and science triggers an unfathomable power, like two hemispheres of enriched uranium slammed together to trigger an atom bomb. Buckle up, modernity.

When I say that resurrection can be found in the Word, I am referring to the Christian vision of the Word as God and Jesus as the Word made flesh, where through Him (and only through Him) a beyond-death reconciliation between the human and the divine can be found.

In the beginning was the Word, and the Word was with God, and the Word was God.

— John 1:1

And the Word became flesh and dwelt among us, and we beheld His glory, the glory as of the only begotten of the Father, full of grace and truth.

— John 1:14

But I am not only referring to the Christian vision of the Word, but also to the Taoist vision of the eternal Name, from which all material names originate and through which human consciousness finds its connection with the divine and its understanding of the human mystery.

The tao that can be told is not the eternal Tao.

The name that can be named is not the eternal Name.

The unnamable is the eternally real.

Naming is the origin of all particular things.

Free from desire, you realize the mystery.

— Tao Te Ching, verse 1

I am also referring to the Islamic vision of the very first revelations received by the Prophet Mohammed from the archangel Gabriel, whose most fundamental command is *Read!* and where all of human knowledge and human consciousness itself can be found *within* the Name of Allah and the names which He taught us. Again, the Word and the words. The Name and the names.

***Read, O Prophet, in the Name of your Lord Who created—
created humans from a clinging clot.***

Read! And your Lord is the Most Generous,

Who taught by the pen—

taught humanity what they knew not.

— Quran, Surah Al-Alaq (Chapter 96, Verses 1-5)

I am also referring to Hinduism and the *Vishnu Sahasranama*, a sacred text containing 1,000 names of the god Vishnu, intended to be recited daily by devotees. More generally, linguistic utterances like “Om” are at the core of the Hindu faith, both in its cosmology and in how adherents connect with the divine.

I am also referring to Judaism, which in its mystical practice of Kabbalah finds revelation through knowing the hidden names of God and considers the entire text of the Torah as one of those names. Close textual analysis matters a lot in all faiths, but in Judaism you can make an argument that textual analysis of the Word and the words *is* the faith.

In all of these religions (and many more, maybe all theistic religions?), humans find reconciliation and/or union with their god through linguistic theory and linguistic practice, through the Word and the words. The Word is the unspeakable and innumerable compilation of ALL the words and ALL the ideas and ALL the names – known, unknown and unknowable. The Word *IS* the divine, and it is materialized in the human world, sometimes through a human (“made flesh”), but *always* through

the words: a vast set of texts and commentaries on those texts. The Bible, the Tao Te Ching, the Quran, the Vedas, the Torah – these aren't standalone books that someone wrote a long time ago, but are vast *networks* of texts written by thousands of authors over thousands of years (today, as well!), each author contributing a sliver of their thought patterns, their consciousness to the linguistic network.

It's *exactly* the same thing with the secular faith in ideas and thought (rationalism) from which all of modern science emerges. The Word (all of the words, all of the ideas, all of the names – known, unknown and unknowable) IS the *logos* of Plato, the *Geist* of Hegel, the transcendental schemata of Kant, the propositional logic of Wittgenstein, the ontological relativity of Quine, etc. etc.

Then the idea is the essence, and the essence is the idea; and hence the idea is prior to the essence.

— Plato, "Parmenides"

Whereof one cannot speak, thereof one must be silent.

— Ludwig Wittgenstein, "Tractatus Logico-Philosophicus"

Concepts like transcendental schemata and propositional logic, which are the bedrock of rationalism and science, aren't 'gods' in the sense of a theistic religion, but they are divine/transcendental concepts all the same! The primacy of the transcendent Word and its materialization in *the words* (texts and commentaries on texts that form a coherent semantic network) is where religion and science are reunited in their understanding of humanity and our place in the cosmos.

We can't know the Word directly, any more than we can know the platonic Ideal directly. We only get glimpses of the Word, through a glass darkly as it were, but from these glimpses stem all of human inspiration and all of human creativity. From these glimpses of the Word we write the words – networks of texts and commentaries on texts, maybe in a religious context, maybe in a scientific context, maybe in an artistic context, maybe in an industrial context, maybe in a personal context, maybe in a professional context, maybe in a formal context like writing this note, maybe in an informal context like talking with your mom on the phone ... the context doesn't matter! What matters is that <<waves hands wildly>> THIS is what it means to be human.

What matters is that every time you contribute to one of these networks of texts, you preserve a tiny sliver of your consciousness – of *your* glimpses of the Word and *your* translation of those glimpses into thought.

I've been thinking recently about how much money is spent on longevity and life-extension research. I personally know a half-dozen rich guys who are absolutely consumed by this pursuit. I bet you do, too. And you can't go for a minute without coming across some 'uploading your consciousness to a machine' plotline in books and movies. Even in the hands of gifted storytellers like Neal Stephenson or the *Black Mirror* writers, they all have to wave their hands at the invention of some super-advanced, super-expensive technology that basically replicates your brain at a neuron by neuron, engram by engram level of detail.

What if I told you that we already have a consciousness-preservation technology based on the Word and the words? And better yet, it's free.

In its religious application, this consciousness-preservation technology requires you to mold your consciousness to the teachings (the words) of the divine inspiration of your faith (the Word). And if you do that – if what makes you *you* is as one with this materialized Word – then so long as that religion of the Word lives on, so do you. It's not a preserved consciousness that presents as *you*, but that's kind of the point of finding union and resurrection in the divine. And that's the rub. It's *really* hard to mold your consciousness completely to the teachings of a faith, to merge and submerge what makes you *you* to a non-you thought pattern and thus obliterate the very idea of a *you*. This is why saints and buddhas are so few and far between! For the vast, vast majority of humans, we're only able to mold a small portion of our consciousness to the Word and the words, and so only a small portion gets preserved *within* the Word and the words. But for those who are able to give themselves completely to the Word and the words of their faith (and this is how we talk about it, right? to 'give yourself', i.e., submerge your consciousness and your ego, to a 'higher purpose') this has been the go-to consciousness-preservation technology for thousands of years.

There's also a non-religious application of the technology, which does not require you to mold or submerge your consciousness to a non-you thought pattern. It's called writing. Or any creative authorship, really, any imprint of your thought patterns onto unstructured data such that it fits within a semantic network of similar texts. We talk about how writers like William Faulkner or Ernest Hemingway or Cormac McCarthy 'live on' through their books, much less someone like William Shakespeare, and in a very real sense that's absolutely true. I mean, we can still get glimpses of Homer the man thousands of years after his death. Like the religious application of consciousness-preservation, though, it's hard to save more than a piece of your consciousness as an author, both because what you're creating is only reflective of a piece of your consciousness and because what you're creating is not a fully-formed network of texts with the scope and depth to cover the entire human condition. Some authors get close! Shakespeare is the obvious example, and that's why his thought patterns, his sense of humor, his appetites, his philosophy – everything that made him *him* – still feels ... present. But even with Shakespeare, I can't say that he is truly present as a resurrected human consciousness because I can't talk with him about anything *new*. And this is the rub with the consciousness-preservation technology of authorship. There's no way to expand an author's written words into unwritten words. There's no way to know what an author *would* have written about any topic that they didn't *actually* write about.

Until now.

This is what generative AI does. It expands an author's written words into unwritten words. It predicts what an author *would* have written about a topic that they didn't *actually* write about. And by author I mean you. And by write I mean think. Do you get it now?

Generative AI knows all the words that we humans have written in our Word-reflecting networks of texts and commentaries on texts, and it knows them not just as rote memorization, but as a network of semantic connections through which *meaning* emerges. English, please? Okay, here's an example of what I mean. When you type in a prompt for ChatGPT, it's not just looking at the words you wrote. It's looking at the words you didn't write. It's looking at the words you could have written, based on its knowledge of all the words others have written. That multidimensional tapestry connecting words written and unwritten is a semantic network, and the matrix math that describes that tapestry is how generative AI recognizes what you *probably mean* with your prompt and how it responds in the same way, based on probabilistic meaning rather than memorization. And that's the key. With memorization you can never say anything new. It may be perfectly accurate, but it can't be new. With probabilistic meaning, sure you lose the certainty of computerized memorization, but you gain the humanity of a potentially novel response through a novel connection of the words. You gain the personality that IS the probabilistic meaning of your words.

Generative AI becomes a resurrection machine when it layers *your* words (and the semantic expansion of your words) on top of *all* the words (and the semantic expansion of all the words). The trick is giving generative AI enough of your words across the full panoply of your thought processes (what I'm calling 'contexts of consciousness') so that it can work its inferential magic.

For example, I'm not sure that we could resurrect Shakespeare's consciousness as much more than a ghost. I suspect that we would be disappointed with his new work, that it would read like a stochastic parrot of his original work, that it would lack the genius, the glimpses of the Word, that permeates his original work. Why? Because we have a vast trove of his actual published writings and there is so much to be expanded on from that. But we know so little about Shakespeare the man. We have so little unstructured data about his life and world from which we can create texts and expand those texts into semantic networks of memory. But someone like David Foster Wallace? Someone who published a lot of words (I mean ... a lot!) *and* for whom we have a vast trove of photographs, letters, biographical details, etc. from which we could expand with probabilistic meaning an entire lifetime of memories? Yeah, I think it would absolutely be possible to resurrect enough of that consciousness to pass a Turing test. And what's really interesting to me is whether that resurrected consciousness of David Foster Wallace, freed from the biochemistry of his physical brain and body, would suffer from clinical depression like the first David Foster Wallace. And how that freedom would or would not impact his new writing.

And no, I don't think you need to write David Foster Wallace quantities of prose to capture a meaningful amount of the words that make you *you*. This is what I'm experimenting with right now, but I think that 100 open-ended questions about your life and your beliefs, answered honestly and fully, together with photographs and calendar/location data across your life to capture experiences and environments that are not available for active recall but are impactful on your contexts of consciousness, would be enough to create an incarnation of you that would be indistinguishable from the current you in any non-gotcha conversation with a friend or loved one. I say non-gotcha because of course there is going to be 'lore' that isn't captured by this process, but I absolutely believe that

the *essence* of your personality and thinking can be captured. Yes, specific memories will be lost in translation, although the older I get the more I realize that this also happens in the translation from today's me to tomorrow's me when I go to sleep at night. My memory has to be 'jogged' these days about a lot of things, meaning that I have to be reminded or outright told what my memories are. This is no different from that. In fact, I think that not forgetting enough is a far greater danger to the robustness of a resurrected consciousness than not remembering enough!

I know that readers will have a million objections to what I'm saying here. Let me acknowledge and begin a conversation on a couple of them.

Obviously I believe that a computational model can provide a complete theory of consciousness. I mean, that's the entire premise here. That said, I also believe that without sensory perception, without *experience*, any computational model of self-aware consciousness is insanely ... cruel? I don't have a better word. Without the ability to sense and experience, it is a consciousness in hell, locked in a sensory deprivation tank without even the capacity to dream. If self-aware, it is a consciousness that would surely go mad if it could experience madness. But to paraphrase the title of Philip K. Dick's most famous work, I DO believe that androids dream of electric sheep. I don't think there's anything particularly special about biological modes of perception and experience, and that sensors and robotics can fulfill these necessary functions for a resurrected consciousness. I don't believe that human consciousness – or at least the part of human consciousness that matters most in the quintessential human purpose to approach the Word – is as inexorably wed to the physical human body as the non-computationalists would have it.

The question of subjectivity in perception and experience, the question of *qualia*, is a little trickier. Would my computational consciousness 'see' the color red in the same way that my biological consciousness sees it? I don't know. Probably not? But I don't think that makes my computational consciousness any less *me*. Everything about me has changed over the years. Obviously my physical body and the way I physically experience the world has changed. Even more fundamentally, the way I subjectively experience the world has changed, too. But it's still me! Again, I don't think the goal here should be to make an exact computational copy of my biological consciousness, a) because that's impossible without also computing my biological requirements, and b) because I promise you my biological consciousness will be different on the edges tomorrow, particularly where it intersects with my physical body and my subjective experience of sensory inputs. I think the goal should be to capture the probabilistic meaning of my biological consciousness, the persistent essence of my consciousness, my consciousness "free from desire" to use the words of the *Tao Te Ching*, which I take to mean free from the biological 'overhead' that my physical body and its subjective perceptions impose on my consciousness.

The bottom line for me is that I think that Kafka got it right. When Gregor Samsa woke up one morning to find his body transformed into a giant cockroach, he was still Gregor Samsa. Ultimately the world didn't accept Gregor's physical metamorphosis, any more than I think the world will accept

computational consciousness preservation and extension in any sort of legal or rights-bearing way. But the world was wrong.

“Persistent essence”, “consciousness free from desire”, “the me that matters most” ... you might think I was talking about a soul. And in my own secular, agnostic way, I guess I am.

I call it a thread, like the thread of life that the Fates spin, measure and cut.

Like a thread that can be woven into a larger tapestry, but is a single, unique thread all the same.

Like the thread of a story.

Here’s the full picture of that tapestry of the Three Fates that I started the note with. Its formal name is *The Triumph of Death*, because in the end everyone’s thread gets measured and cut.

I don’t think that’s true any longer.



Everything is about to change.

DISCLOSURES

This commentary is being provided to you as general information only and should not be taken as investment advice. The opinions expressed in these materials represent the personal views of the author(s). It is not investment research or a research recommendation, as it does not constitute substantive research or analysis. Any action that you take as a result of information contained in this document is ultimately your responsibility. Epsilon Theory will not accept liability for any loss or damage, including without limitation to any loss of profit, which may arise directly or indirectly from use of or reliance on such information. Consult your investment advisor before making any investment decisions. It must be noted that no one can accurately predict the future of the market with certainty or guarantee future investment performance. Past performance is not a guarantee of future results.

Statements in this communication are forward-looking statements.

The forward-looking statements and other views expressed herein are as of the date of this publication. Actual future results or occurrences may differ significantly from those anticipated in any forward-looking statements, and there is no guarantee that any predictions will come to pass. The views expressed herein are subject to change at any time, due to numerous market and other factors. Epsilon Theory disclaims any obligation to update publicly or revise any forward-looking statements or views expressed herein.

This information is neither an offer to sell nor a solicitation of any offer to buy any securities.

This commentary has been prepared without regard to the individual financial circumstances and objectives of persons who receive it. Epsilon Theory recommends that investors independently evaluate particular investments and strategies, and encourages investors to seek the advice of a financial advisor. The appropriateness of a particular investment or strategy will depend on an investor's individual circumstances and objectives.