

WHAT WE TALK ABOUT WHEN WE TALK ABOUT AI (PART ONE)

A Normal Person's Explainer on What Generative AI is and Does

Part 1 – In the Beginning was the Chatbot

(Go to Part Two)

*"Are you comfortably seated? Yes, well, let's begin." *Clears throat theatrically**

"Our experience, in natural theology, can never furnish a true and demonstrated science, because, like the discipline of practical reason, it can not take account of problematic principles. I assert that, so far as regards pure logic, the transcendental unity of apperception is what first gives rise to the never-ending regress in the series of empirical conditions. In this case it remains a mystery why the employment of the architectonic of human reason is just as necessary as the intelligible objects in space and time, as is proven in the ontological manuals. By means of analysis, it must not be supposed that the transcendental unity of apperception stands in need of our sense perceptions. Metaphysics, for example, occupies part of the sphere of the transcendental aesthetic concerning the existence of the phenomena in general..."

It was 1995, and several of us who worked in my community college's Macintosh lab were hunting around the net for weird software to try out, back when weird software felt fun, not dangerous. Someone found a program on the nascent web that would almost instantly generate pages of thick and unlovely prose that wasn't actually Kant, but looked like it. It was, to our definitionally untrained eyes, nearly indistinguishable from the Immanuel Kant used to torture undergrad college students.



The logo of the Kant
Generator Pro

We'd found the Kant Generator Pro, a program from a somewhat legendary 90s programmer known for building programming tools. And being cheeky. It was great. (recent remake [here](#)) We read Faux Kant to each other for a while, breaking down in giggles while trying to get our mouths around Kant's daunting vocabulary. The Kant Generator Pro was cheeky, but it was also doing something technically interesting.

The generator was based on a Markov chain: a mathematical way of picking some next thing, in this case, a word. This generator chose each next word using a random walk through all Kantian vocabulary. But in order to make coherent text rather than just random Kant words, it had to be weighted: unrandomized to some extent. The words had to be weighted *enough* to make it form human-readable Kantian sentences.

A text generator finds those weights using whatever text you tell the computer to train itself on. This one looked at Kant's writing and

built an index of how often words and symbol appeared together. Introducing this “unfairness” in the random word picking gives a higher chance for some words coming next based on the word that came before it. For instance, there is a high likelihood of starting a sentence with “The,” or “I,” or “Metaphysics,” rather than “Wizard” or “Oz.” Hence, in the Kant Generator Pro “The” could likely be followed by “categorical,” and when it is the next word will almost certainly be “imperative,” since Kant went on about that so damn much.

The Kant Generator Pro was a simple ancestor of ChatGPT, like the small and fuzzy ancestors of humans that spent so much time hiding from dinosaurs. All it knew, for whatever the value of “knowing” is in a case like this, was the the words that occurred in the works of Kant.

Systems like ChatGPT, Microsoft Copilot, and even the upstart Deepseek use all the information they can find on the net to relate not just one word to the next, like Kant Generator Pro did. They look back *many* words, and how likely they are to appear together over the span of full sentences. Sometimes a large language model takes a chunk as is, and appears to “memorize” text and feed it back to you, like a plagiarizing high schooler.

But it’s not clear when regurgitating a text verbatim is a machine copying and pasting, versus recording a statistical map of that given text and just running away with the math. It’s still copying, but not copying in a normal human way. Given the odds, it’s closer to winning a few rounds of Bingo in a row.

These chatbots index and preserve the statistical relationships words and phrases have to each other in any given language. They start by ingesting all the digital material their creators can find for them, words, and their relationships. This is the training people talk about, and it’s a massive amount of data. Not good or bad data, not meaningful or meaningless, just everything, everywhere people have built

sentences and left them where bots could find them. This is why after cheeky Reddit users mentioned that you could keep toppings on pizza by using glue, and that ended up becoming a chatbot suggestion.

Because people kept talking about using glue on pizza, especially after the story of that hilarious AI mistake broke, AI kept suggesting it. Not because it thought it was a good idea. AI doesn't think in a way familiar to people, but because the words kept occurring together where the training part of the AI could see them together. The AI isn't right here, we all know that, but it's also not wrong. Because the task of the AI isn't to make pizza, the task is to find a next likely word. And then the next, and next after that.

Despite no real knowing or memorizing happening, this vast preponderance of data lets these large language models usually predict what is likely to come next in any given sentence or conversation with a user. This is based on the prompt a user gives it, and how the user continues to interact with it. The AI looks back on the millions of linguistic things it has seen and built statistical models for. It is generally very good at picking a likely next word. Chatbots even to feel like a human talking most of the time, because they trained on humans talking to each other.

So, a modern chatbot, in contrast to the Kant Generator Pro, has most of the published conversations in modern history to look back on to pick a next good word. *I put leash on the*, blimp? Highly unlikely, the weighting will be very low. *Véranda*? Still statistically unlikely, though perhaps higher. *British politician*? Probably higher than you'd want to think, but still low. *Table*? That could be quite likely. But how about *dog*? That's probably the most common word. Without a mention of blimps or parliamentarians or tables in the recent text, the statistics of all the words it knows means the chatbot will probably go with dog. A chatbot

doesn't know what a dog is, but it will "know" dog is associated with leash. How associated depends on the words that have come before the words "dog," or "leash."

It's very expensive and difficult to build this data, but not very hard to run once you have built it. This is why chatbots seem so quick and smart, despite at their cores being neither. Not that they are slow and dumb – they are doing something wholly different than I am when I write this, or you as you read it.

Ultimately, we must remember that chatbots are next-word-predictors based on a great deal of statistics and vector math. Image generators use a different architecture, but still not a more human one. The text prompt part is still an AI chatbot, but one that replies with an image.

AI isn't really a new thing in our lives. Text suggestions on our phones exists somewhere between the Kant Generator Pro and ChatGPT, and customize themselves to our particular habits over time. Your suggestions can even become a kind of statistical fingerprint for your writing, given enough time writing on a phone or either any other next word predictor.

We make a couple bad mistakes when we interact with these giant piles of vector math and statistics, running on servers all over the world. The first is assuming that they think like us, when they have no human-like thought, no internal world, just mapping between words and/or pixels.

The other is assuming that because they put out such human-like output, we must be like them. But we are not. We are terribly far from understanding our own minds completely. But we do know enough to know biological minds are shimmering and busy things faster and more robust than anything technologists have ever yet built. Still, it is tempting, especially for technologists, to have some affinity for this thing that seems so close to, but not exactly, us. It feels like it's our first time getting to

talk to an alien, without realizing it's more like to talking to a database.

Humans are different. Despite some borrowing of nomenclature from biology, neural nets used in training AI have no human-style neurons. The difference shows. We learn to talk and read and write with a minuscule dataset, and that process involves mimicry, emotion, cognition, and love. It might also have statistical weighting, but if it does, we've never really found that mechanism in our minds or brains. It seems unlikely that it would be there in a similar form, since these AIs have to use so much information and processing power to do what a college freshman can with a bit of motivation. Motivation is our problem, but it's never a problem for AIs. They just go until their instructions reach an end point, and then they cease. AIs are unliving at the start, unliving in the process, and unliving at the end.

We are different. So different we can't help tripping ourselves up when we look at AI, and accidentally see ourselves, because we want to see ourselves. Because we are full of emotions and curiosity about the universe and wanting to understand our place in it. AI does not want.

It executes commands, and exits.

(Go to Part Two)