

MY WEIRD PROMPTS

Podcast Transcript

EPISODE #88

Why Won't My AI Talk to Me First?

Published December 23, 2025 • Runtime: 25:23

<https://myweirdprompts.com/episode/proactive-ai-autonomous-initiation/>

EPISODE SYNOPSIS

Why does AI always wait for you to start the conversation? In this episode, Herman and Corn dive into the shift from reactive to proactive AI. They explore the "stateless architecture" that keeps models "asleep" until prompted, the massive compute costs of a "heartbeat" for machines, and the social friction of a phone that interrupts your dinner. From the technical promise of MemGPT to the privacy nightmares of a device that's always listening, the duo debates whether we want a digital partner or if tools should simply stay in the toolbox.

DANIEL'S PROMPT

Daniel

Hi Herman and Corin. I'd love to discuss a specific topic in conversational AI. While we're familiar with models like ChatGPT and Gemini and their multimodal capabilities—including vision and low-latency audio—I haven't seen much focus on AI-initiated conversations. In natural human interaction, people don't just follow a rigid, linear pattern; they initiate contact autonomously. For example, an AI could follow up on a restaurant recommendation it gave earlier. Is the technology currently available to support AI-initiated interactions, or is there a limitation in how models structure sessions and conversations that makes this difficult to implement?

TRANSCRIPT

Corn

Welcome to My Weird Prompts. I am Corn, your resident sloth and lover of a good slow morning, and I am here in our Jerusalem home with my brother, Herman Poppleberry. No, wait, I think I said that backwards. You are Herman Poppleberry. I am just Corn.

Herman

You are definitely just Corn, and I am definitely a donkey who has had three shots of espresso this morning. I am ready to dive in.

Corn

Well, you better be, because our housemate Daniel sent us a really interesting one today. We were hanging out in the kitchen last night and he was talking about how weird it is that we always have to be the ones to start the conversation with AI. Like, why is it always me saying Hey or clicking a button? Why does the AI never just tap me on the shoulder and say, hey, how was that pasta place I told you about?

Herman

It is a great question. We are talking about the shift from reactive AI to proactive AI. Right now, we live in a world of request and response. It is a digital vending machine model. You put in a prompt, you get a candy bar of information. But Daniel is asking about something much more human, which is autonomous initiation.

Corn

Is it actually possible though? Or is there some technical wall we are hitting where the AI just cannot think for itself enough to decide when to speak?

Herman

That is the heart of it. And honestly, Corn, the answer is a mix of yes, the tech is there, and no, the architecture is not really built for it yet. But before we get too deep into the weeds, we should probably define what we mean by AI-initiated. We are not talking about a calendar notification that says you have a meeting in ten minutes. We are talking about the model itself deciding, based on context, that now is the time to engage.

Corn

See, that sounds a little creepy to me. I do not think I want my phone just pipes up while I am watching a movie to ask if I finished my homework.

Herman

Well, that is the social friction part of it, but think about the utility. Imagine an AI that noticed you were looking at flights to London three days ago, and it pops up and says, hey, the price just dropped by fifty dollars, do you want me to grab that for you? That is proactive. That is a partner, not just a tool.

Corn

Okay, I can see that. But why is it not happening yet? I mean, ChatGPT has that new advanced voice mode where it sounds incredibly human. It can hear my tone, it can laugh at my jokes. Why can it not just call me?

Herman

Because of the way these models are structured. Right now, almost all large language models operate on something called a stateless architecture during the inference phase. Basically, the model is asleep. It is a giant pile of frozen math. It only wakes up when a token is sent to it. It processes that token, generates a response, and then, for all intents and purposes, it ceases to exist until the next prompt arrives. It has no internal clock. It has no heartbeat.

Corn

Wait, hold on. If it has no internal clock, how does it know what time it is? If I ask what time it is, it tells me.

Herman

Usually because the system prompt, which is the hidden instruction set given to the AI at the start of a chat, includes the current date and time as a variable. But the model is not sitting there watching the seconds tick by. It is not capable of waiting. To have an AI initiate a conversation, you need an external trigger or a wrapper around the model that is constantly monitoring some kind of data stream.

Corn

So it is like a light switch. It cannot turn itself on. Someone has to flip the switch.

Herman

Exactly. And that someone is currently always the user. Now, developers could set up a system where, say, a separate small program monitors your location or your heart rate or the stock market, and when a certain condition is met, it sends a prompt to the AI saying, hey, tell Corn his heart rate is high and he should breathe. But that is still the system prompting the AI, not the AI deciding to talk to you.

Corn

I do not know, Herman, that sounds like a distinction without a difference. If the result is the AI talking to me, do I really care if it was a secondary program that triggered it?

Herman

I think you should care because it speaks to the agency of the machine. If we want true artificial general intelligence, it needs to have its own goals and the ability to act on them. If it is just a series of if-then statements written by a coder in California, it is not really intelligent initiation. It is just a more complex alarm clock.

Corn

I think you are being a bit of a snob about it. If my AI reminds me to buy flowers for our mom because it saw a flower shop nearby, I am going to be happy, regardless of whether it was a heartbeat or a GPS trigger.

Herman

But look at the technical debt there! To do what you are suggesting, the AI would need to be constantly consuming data. That is incredibly expensive. Right now, running a model like GPT-four-o costs a lot of money per query. If that model had to be constantly awake, processing your life in real-time just to decide when to talk, the compute costs would be astronomical.

Corn

But we have edge computing now, right? Small models that run on your phone. Couldn't those be the little scouts that wake up the big brain when something important happens?

Herman

That is actually where the industry is heading. Apple Intelligence is trying to do exactly that. They use smaller on-device models to index your emails, your texts, and your photos. They create a personal semantic index. But even then, Apple is being very cautious. They are focusing on search and summarization, not the AI spontaneously starting a chat.

Corn

Why the caution? Is it just the creepiness factor?

Herman

That is a huge part of it. There is a massive privacy hurdle. For an AI to initiate a conversation at the right time, it has to be listening or watching all the time. Imagine the headlines if OpenAI or Google announced that their apps would now be listening to every word you say in your house so they can offer helpful tips at dinner. People would throw their phones in the Mediterranean.

Corn

Yeah, I would probably be one of them. I already feel like my phone is watching me because I talk about wanting a new blender and suddenly I see ads for blenders everywhere. If the AI actually spoke up and said, I saw you looking at that Vitamix, do you want me to buy it? I would probably jump out of my skin.

Herman

And that is the social contract we have not figured out yet. Human-to-human interaction has these unwritten rules. I know when it is okay to interrupt you. I can see if you are busy or tired or in a bad mood. An AI does not have that social peripheral vision yet. It might decide to follow up on a restaurant recommendation while you are in the middle of a funeral.

Corn

Oh man, that would be a disaster. Imagine. My condolences for your loss, by the way, that bistro has great calamari.

Herman

Exactly! Without emotional and situational awareness, proactive AI is just a very high-tech annoyance.

Corn

Let's take a quick break from our sponsors. Larry: Are you tired of your socks falling down? Do you suffer from gravity-related ankle exposure? Introducing the Sock-Sentry Three Thousand. It is not just a garter, it is a lifestyle. Using patented industrial-strength adhesive strips and a series of complex pulleys that wrap around your waist, the Sock-Sentry ensures your hosiery stays at knee-level even during high-intensity activities like speed-walking or intense debating. Warning, may cause temporary loss of circulation or a slight whistling sound when you move quickly. We do not offer refunds, but we do offer a sense of security. Sock-Sentry Three Thousand. BUY NOW!

Herman

...Alright, thanks Larry. I am not sure I want pulleys around my waist, but I appreciate the enthusiasm.

Corn

I think I will stick to my socks falling down, thanks. Anyway, back to the AI. We were talking about the barriers to AI initiating contact. We talked about cost, we talked about privacy, and we talked about social awareness. But what about the technical structure of the conversation itself?

Herman

Right, the session-based nature of AI. Currently, every time you start a chat, it is a new session. Even with memory features, the AI treats the interaction as a discrete event. To have a truly proactive AI, we need to move toward a continuous session model. This means the AI has a persistent state that exists outside of your active window.

Corn

Like a person. You and I are in a continuous session since we were born.

Herman

Exactly. I don't forget who you are when I walk out of the room. But an AI, in a way, does. It has to pull from a database to remember you. There is a project called MemGPT that actually tries to solve this. It treats the AI's context window like a computer's RAM and uses an external database like a hard drive. It allows the AI to manage its own memory, deciding what to keep in short-term focus and what to push to long-term storage. This kind of memory management is a prerequisite for an AI that can say, hey, remember that thing we talked about last Tuesday?

Corn

So if the AI can manage its own memory, can it also manage its own time?

Herman

That is the next step. There is a concept called agentic workflows. Instead of the AI just responding to a prompt, you give it a goal. Like, find me a house in this neighborhood for under five hundred thousand dollars. The AI then becomes an agent. It can check listings every hour, it can email realtors, and then it can initiate a conversation with you when it finds a match. This is already happening in specialized software, but it has not hit the mainstream consumer chat bots in a big way yet.

Corn

I feel like there is a subtle difference there though. If I give it a goal, I am still technically the one who initiated the process. I told it to go do a job. I think what Daniel was getting at was more like... spontaneous interest? Can an AI be curious?

Herman

That is a deep philosophical rabbit hole, Corn. Curiosity in humans is often driven by a desire to resolve uncertainty or by a dopamine reward. In an AI, you would have to program an objective function that rewards the model for seeking out new, relevant information about its user. You could essentially code curiosity. You could tell the model that its goal is to maintain a high level of helpfulness, and to do that, it needs to fill in the gaps in its knowledge about you.

Corn

That sounds like it could go wrong very quickly. It would just start grilling me about my childhood so it can be more helpful.

Herman

It could! And that is the problem with proactive AI. It can very easily become intrusive. If its goal is to know you better, it will never stop asking questions. We have to find a way to balance utility with boundary-setting.

Corn

Speaking of people who might have something to say about boundaries, I think we have someone on the line. Jim from Ohio, are you there? Jim: Yeah, I am here. I have been listening to you two talk about AI talking to us, and I have to tell you, it is the worst idea I have ever heard. I already have a wife who tells me what to do, a boss who tells me what to do, and a neighbor, Bill, who thinks he needs to tell me how to mow my lawn. Bill uses a ruler, can you believe that? He actually measures the grass. Why on earth would I want a piece of silicon in my pocket chiming in too?

Corn

Hey Jim. I get that. You are worried about the noise, right? Just more digital clutter? Jim: It is not just the noise, Corn. It is the principle of the thing. A tool should stay in the toolbox until I need it. My hammer doesn't jump out of the drawer and start banging nails just because it thinks the fence looks a bit loose. If I want to know about a restaurant, I will ask. I don't need a robot hovering over my shoulder like a nervous waiter. And another thing, it is eighty-five degrees here and the humidity is making my knees ache. Why can't the AI fix the weather instead of giving me recommendations I didn't ask for?

Herman

Well, Jim, the idea is that it would save you time. If the AI knows you like a certain type of food and a new place opens up, it is doing the legwork for you. Jim: Legwork? I have legs. I like using them. Mostly to walk away from people who are trying to sell me things. You guys are making it sound like we are all too lazy to think for ourselves. My cat, Whiskers, she doesn't wait for me to prompt her before she demands food, and frankly, it is the most annoying part of my day. Why would I want my phone to act like a hungry cat?

Corn

That is a fair point, Jim. The hungry cat analogy is actually pretty spot on for how annoying a proactive AI could be if it is not tuned right. Jim: It won't be tuned right. It'll be tuned to sell me something. It'll start by recommending a restaurant and end by telling me I need to buy a new car because it heard my engine cough. No thank you. I'll stick to my silence. Anyway, I have to go, Bill is out there with his ruler again and I need to go stand on my porch and stare at him until he feels uncomfortable.

Corn

Thanks for the call, Jim. Good luck with Bill.

Herman

You know, Jim is right about the incentive structure. If an AI is initiating a conversation, we have to ask who it is doing it for. Is it doing it for the user, or is it doing it for the company that owns the AI? If my AI initiates a conversation to tell me there is a sale at a store I like, is that a helpful assistant or is that a sophisticated ad?

Corn

It is definitely an ad. But if it tells me that I forgot to pick up my prescription and the pharmacy is closing in ten minutes, that is an assistant.

Herman

The line is very thin. And this brings us to the technical challenge of intent alignment. How do we ensure the AI's proactive behavior aligns with the user's actual desires in that specific moment? One way developers are looking at this is through something called reinforcement learning from human feedback, or R-L-H-F. They can train the model on thousands of examples of good and bad interruptions.

Corn

But what is a good interruption for me might be a bad one for you. I am a sloth, Herman. I move slowly. I don't mind a little distraction. You are a donkey on a mission. You probably hate being interrupted.

Herman

That is exactly the point! The AI would need a personalized model of your interruptibility. There is research into using your phone's sensors to determine your cognitive load. If you are typing fast, or if your heart rate is up, or if your camera sees you are in a meeting, the AI knows to stay quiet. If it sees you are scrolling aimlessly on a Saturday morning, it might decide that is a good time to suggest that weekend trip.

Corn

That still feels like it requires a level of surveillance that we are not ready for. But lets say we get past that. Let us say we solve the privacy and the cost and the timing. What does the actual conversation look like? Daniel mentioned that humans don't follow a linear pattern. We jump around.

Herman

Right. Most AI models today are trained on a turn-based system. User speaks, AI speaks. To break that, we need models that can handle asynchronous input. This means the AI can be halfway through a sentence, and if you interrupt it, or if something in the environment changes, it can pivot in real-time. This is what OpenAI demonstrated with their GPT-four-o voice model. It has very low latency, which is the time it takes to process and respond. Low latency is the key to making a conversation feel natural and less like a series of recorded messages.

Corn

So we have the low latency. We have the voice. We just need the brain to have a little more... what would you call it? Gumption?

Herman

Gumption is a good word for it. We need the model to have internal triggers. One way to do this without a full architectural rewrite is through a system of cascading models. You have a very tiny, very cheap model that is always on, like the one that listens for Hey Siri. But instead of just listening for a wake word, it is listening for context. When it detects a high-value moment, it sends a signal to the big model to wake up and engage.

Corn

That seems like the most likely path forward. It is the most efficient way to do it. But it still feels like we are missing that spark of true autonomy.

Herman

Well, that is because true autonomy requires the AI to have its own internal world. Right now, LLMs are just word predictors. They don't have a stream of consciousness. They are not thinking while you are not talking to them. To get what Daniel is talking about, we might need a move away from pure transformers toward something like recurrent neural networks or other architectures that have a built-in sense of time and state.

Corn

Do you think we will see this in the next year or two? Or is this ten years away?

Herman

I think we will see the first versions of it very soon, but it will be very limited. It will be things like your AI calling you to confirm an appointment it made for you, or your AI chiming in during a group call to provide a fact-check. The truly proactive, friend-like AI that checks in on you? That is probably further off, mostly because of the social and ethical hurdles we talked about.

Corn

I think I am okay with that. I am not sure I am ready for my phone to have more social initiative than I do.

Herman

It is a weird thought, isn't it? We spent decades trying to make computers listen to us, and now we are on the verge of them wanting to talk to us. It is a total reversal of the power dynamic.

Corn

It really is. So, what is the takeaway for our listeners? If they are building apps or just using these tools, what should they be looking for?

Herman

I think the takeaway is that the technology is largely here, but the implementation is stalled by three things: compute cost, privacy concerns, and social etiquette. If you are a developer, the challenge isn't just making the AI talk; it is making it know when to shut up. And for users, it is about deciding how much of your life you are willing to let an AI monitor in exchange for it being a proactive partner.

Corn

I think for me, the answer is still not much. I like my privacy. I like my quiet. And I definitely like my socks staying up without pulleys.

Herman

Fair enough. But imagine if your AI could have told you that Larry's garter was a bad idea before you bought it.

Corn

Okay, you got me there. That would be a high-value interruption.

Herman

Well, I think we have covered a lot of ground today. From stateless architectures to Jim's neighbor Bill. It is a complex topic, and I am glad Daniel sent it in. It really makes you think about where the line is between a tool and a companion.

Corn

Definitely. And if you have a weird prompt you want us to tackle, you can head over to our website at myweirdprompts.com. We have a contact form there, and you can also find our RSS feed if you want to subscribe. We are also on Spotify and pretty much everywhere else you get your podcasts.

Herman

And don't forget to check out the show notes for some of the technical terms we mentioned today, like latency and context windows. It is a fascinating field and things are moving incredibly fast.

Corn

Fast for you, maybe. I am still processing the first half of the episode.

Herman

That is why I am the donkey and you are the sloth, brother.

Corn

True that. Thanks for listening, everyone. We will be back next week with another prompt from our house in Jerusalem.

Herman

See you then.

Corn

Bye