

## MY WEIRD PROMPTS

Podcast Transcript

### EPISODE #131

# 2026 AI Roadmap: From Invisible Agents to Physical Robots

Published January 01, 2026 • Runtime: 17:02

<https://myweirdprompts.com/episode/2026-ai-agent-roadmap/>

## EPISODE SYNOPSIS

In this forward-looking episode of My Weird Prompts, hosts Herman and Corn dive into a listener-submitted roadmap for the year 2026. They explore a future where artificial intelligence moves beyond the chat box and becomes an "invisible" layer within our operating systems, powered by highly optimized small language models that prioritize privacy and speed. The conversation tracks the evolution of the "agentic economy," where AI agents equipped with digital wallets negotiate and execute transactions on behalf of humans, shifting the digital landscape from business-to-consumer to business-to-agent interfaces. As the year progresses, the technical focus shifts from the brute-force scaling of parameters to "inference-time compute," where models are judged by their reasoning depth rather than their size. Finally, the duo discusses the "physical grounding" of AI, as Vision-Language-Action models allow robots to transition from pre-programmed tools to generalized helpers in our homes. This episode serves as a comprehensive guide to the year AI matures into a reliable, ubiquitous infrastructure that anticipates our needs and acts as a true partner in both the digital and physical worlds.

## DANIEL'S PROMPT

## Daniel

Happy New Year, Herman and Corren. I'd love to get your thoughts on what we can look forward to this year regarding developments in AI, perhaps broken down quarter by quarter. After 2024's global expansion and 2025's focus on agentic workflows and MCP, what do you think will be the major theme of 2026? Do you think agentic AI will move from the technical frontier into everyday use? Additionally, what do you foresee for model growth? Will we continue to see scaling with larger models and more parameters, or will there be fundamental shifts in AI architecture?

# TRANSCRIPT

## Corn

Happy New Year everyone! Welcome to My Weird Prompts. It is January first, twenty twenty-six, and I am sitting here in our living room in Jerusalem with my brother.

## Herman

Herman Poppleberry at your service. Happy New Year, Corn. I cannot believe we are already in twenty twenty-six. It feels like just yesterday we were obsessing over the first multimodal models, and now look where we are.

## Corn

It really does move fast. Our housemate Daniel sent us a great prompt to kick off the year. He was asking about what the next twelve months look like for artificial intelligence, specifically broken down by quarters. He noted that twenty twenty-four was about global expansion and twenty twenty-five was the year of agentic workflows and the Model Context Protocol. Now he wants to know if twenty twenty-six is the year agents go mainstream and if we are going to see a shift in model architecture away from just scaling up parameters.

## Herman

Daniel always has his finger on the pulse. It is a fantastic framing because twenty twenty-five really was the year of the architect. We saw the foundations of agents being built with the Model Context Protocol, which, for those who need a refresher, basically allowed different artificial intelligence tools and data sources to talk to each other using a standardized language. But it was still a bit of a frontier town, you know? A lot of early adopters and developers, but not necessarily something our grandmother was using to manage her garden.

## Corn

Exactly. So let's dive into this quarter by quarter. If we look at the first three months of twenty twenty-six, what is the immediate shift? I feel like we are seeing a massive move toward what I call the invisible agent.

### Herman

I love that term. In quarter one, the big theme is going to be operating system integration. Up until now, agents have mostly lived in our browsers or in specific apps. But in early twenty twenty-six, we are seeing the fruit of those deep partnerships between the major model labs and the companies that make our phones and computers. The agent is not an app you open anymore. It is the layer between you and the hardware.

### Corn

Right, so instead of me saying, hey, open my email and find that flight info and then open my calendar and add it, I just tell the system, I am going to London in March, keep me posted. And the system, using those agentic workflows we saw mature last year, just handles the cross-app communication. But Herman, does this rely on the massive cloud models, or are we finally seeing the small language models take the lead here?

### Herman

That is the quarter one breakthrough. We are seeing small language models with three to seven billion parameters that are incredibly optimized for local execution. Because of the architectural improvements in late twenty twenty-five, these small models can now handle complex reasoning that used to require a massive cluster of graphics processing units. So, in quarter one, the theme is privacy-first, local agency. Your data stays on your device, but the agent feels as smart as the giants did a year ago.

### Corn

That makes a lot of sense. It solves the latency issue too. If I have to wait five seconds for a cloud model to tell me where my next meeting is, I might as well just look it up myself. But if it is local, it is instant. Now, moving into quarter two, Daniel mentioned the agentic economy. This is where things get a bit weirder, right?

### Herman

Oh, definitely. Quarter two of twenty twenty-six is when I think we see the rise of the autonomous transaction. Last year, we started talking about giving artificial intelligence agents digital wallets. In the next few months, that becomes a standard feature. We are talking about agents that can not only find you a better insurance rate but can actually negotiate the contract and execute the payment within parameters you set.

### Corn

That feels like a huge leap in trust. I mean, am I ready to let an agent spend my money? I think the breakthrough there has to be in the verification layers. We need those human-in-the-loop safeguards that are not just annoying pop-ups but meaningful checkpoints.

### Herman

Precisely. And that is where the Model Context Protocol comes back in. Because it provides a standardized way for an agent to prove its identity and its authorization level to a third-party vendor. In quarter two, we will likely see the first major retail platforms launch agent-specific interfaces. Instead of a website designed for human eyes, they will have an endpoint designed for an agent to crawl, compare prices, and buy. It is a shift from business-to-consumer to business-to-agent.

### Corn

It is almost like the internet is being rebuilt for bots. Which sounds scary, but it might actually make things more efficient for us humans. I am curious though, does this mean the big models are getting even bigger to handle this complexity? Daniel asked about scaling versus architectural shifts.

### Herman

That is a great bridge to the middle of the year. But before we get into the heavy technical shifts of quarter three, we should probably take a quick break.

### Corn

Good idea. Let's hear from our sponsors. Larry: Is your aura feeling a bit dusty? Are your chakras misaligned with the current lunar cycle? You need the Quantum Bio-Pillow. This is not just a place to rest your head. It is a high-frequency energy-tuning station. Each Quantum Bio-Pillow is stuffed with proprietary hyper-conductive foam and infused with the essence of ancient mountain air. Our users report dreaming in colors that do not even exist in this dimension. One customer said he woke up and could suddenly speak fluent dolphin. Is it science? Is it magic? It is Larry's Quantum Bio-Pillow. Do not let your brain waves go un-tuned for another night. Larry: BUY NOW!

### Corn

...Alright, thanks Larry. I am not sure about the dolphin speech, but I could use a good night's sleep. Anyway, Herman, back to the technical side. Daniel was asking if we are just going to keep seeing more parameters or if there is a fundamental shift coming. What does the second half of twenty twenty-six look like?

### Herman

This is where it gets really exciting for a nerd like me. For the last few years, the transformer architecture has been the undisputed king. We just kept feeding it more data and more compute. But in quarter three of twenty twenty-six, I think we hit the wall of diminishing returns for pure scaling. We are seeing a shift toward what researchers call inference-time compute.

### Corn

Explain that for those of us who are not reading research papers at three in the morning.

### Herman

So, traditionally, a model's intelligence was mostly determined by its training. You pour all the knowledge in at the beginning, and then when you ask it a question, it gives you a quick answer based on that training. Inference-time compute means the model actually spends more time thinking before it speaks. It runs internal simulations, checks its own logic, and explores different paths of reasoning before it gives you the final output. It is like the difference between someone who blurts out the first thing that comes to mind versus someone who sits quietly for a minute and works through the problem.

### Corn

Like the transition we saw with the early reasoning models in late twenty twenty-four and twenty twenty-five, but much more advanced?

### Herman

Exactly. By quarter three of twenty twenty-six, this becomes the standard for all frontier models. We stop talking about how many billions of parameters a model has and start talking about its reasoning depth. This allows models to be smaller and more efficient while being significantly more capable at complex tasks like coding or scientific discovery. It is about working smarter, not just having a bigger brain.

## Corn

That ties into Daniel's question about architectural shifts. Are we still using transformers for this?

## Herman

We are seeing a hybrid approach. Transformers are still there for the heavy lifting of language understanding, but we are seeing them paired with new architectures like State Space Models or even updated versions of Liquid Neural Networks. These are much better at handling incredibly long sequences of data without the memory costs of a traditional transformer. Imagine an agent that can remember every single interaction you have had for the last three years perfectly. That requires a shift in how the model stores and retrieves state.

## Corn

That is a massive shift. It moves the artificial intelligence from being a tool you use to a partner that has a shared history with you. If the model has a perfect memory of our collaboration, it can anticipate my needs in a way that feels almost telepathic.

## Herman

And that leads us right into quarter four. The theme for the end of twenty twenty-six, in my opinion, is physical world grounding. We have had these brilliant brains living in the cloud, but in the final months of this year, we are going to see them truly inhabit the physical world. I am talking about Vision-Language-Action models becoming standard in consumer robotics.

## Corn

You mean like the home robots we have been promised for decades? Are we finally getting the robot that can actually fold the laundry and not just vacuum the floor?

## Herman

We are getting closer. The breakthrough in quarter four will be the ability for these models to generalize. Instead of a robot being programmed to fold a shirt, it will have an agentic brain that understands the concept of fabric and the goal of folding. It can look at a pile of clothes it has never seen before and figure it out in real-time. This is the culmination of the agentic year. The agent moves from your screen to your living room.

### Corn

It feels like the theme of twenty twenty-six is maturity. Twenty twenty-four was the wow factor. Twenty twenty-five was the plumbing and the protocols. Twenty twenty-six is when it all becomes useful, reliable, and physical.

### Herman

I think that is the perfect way to put it. It is the year artificial intelligence stops being a novelty and starts being an infrastructure. It is like the shift from the early days of the internet where you had to explain what a website was, to the day you realized you could not run your business without it.

### Corn

So, to answer Daniel's question directly, yes, agentic artificial intelligence is absolutely moving from the frontier to everyday use. But it might not look like a robot butler right away. It will look like an operating system that knows you, a digital wallet that saves you money, and a research assistant that can think through a problem for ten minutes before giving you a perfect answer.

### Herman

And the scaling laws are changing. We are moving away from the brute force of more parameters and toward the elegance of better reasoning and more efficient architectures. The models are getting deeper, not just wider.

### Corn

That is a lot to look forward to. It makes me realize that we need to be more intentional about how we use these tools. If the agent is going to be our representative in the digital economy, we need to make sure we are setting the right goals and values for it.

### Herman

That is the big human challenge for twenty twenty-six. As the technology matures, our responsibility grows. We are not just users anymore. We are managers of a fleet of digital agents.

**Corn**

Well, I for one am excited to see how this plays out. We will have to check back in at the end of each quarter to see if Herman's predictions hold up.

**Herman**

Hey, I am confident! The research is there. The incentives are there. The only real wildcard is how quickly we as humans can adapt to this new pace of life.

**Corn**

True. We will probably spend half the year just trying to figure out how to talk to our dolphin-speaking neighbors thanks to Larry's pillow.

**Herman**

Exactly. One step at a time.

**Corn**

Well, thank you all for joining us for this first episode of twenty twenty-six. We have a lot of ground to cover this year and I am glad we are doing it together.

**Herman**

It is going to be a wild ride. Thanks for the prompt, Daniel. It really helped us frame the year ahead.

**Corn**

If you want to send us your own weird prompts, you can find the contact form on our website at [my weird prompts dot com](http://myweirdprompts.com). You can also find our full archive and the R S S feed there for subscribers. And of course, we are available on Spotify.

**Herman**

We love hearing from you. Even if your prompt is just about how to fold laundry with a malfunctioning robot.

**Corn**

Especially if it is about that. This has been My Weird Prompts. I am Corn.

**Herman**

And I am Herman Poppleberry.

**Corn**

Happy New Year, and we will talk to you next week.

**Herman**

See you then!

**Corn**

Thanks for listening to My Weird Prompts. Don't forget to visit my weird prompts dot com for more episodes and to get in touch. We will see you next time!

**Herman**

Bye everyone!

**Corn**

Goodbye!