

# MY WEIRD PROMPTS

Podcast Transcript

## EPISODE #92

# The Scaling Wall: Why Bigger AI Isn't Always Smarter

Published December 23, 2025 • Runtime: 18:06

<https://myweirdprompts.com/episode/ai-scaling-limits-model-collapse/>

## EPISODE SYNOPSIS

In this episode of My Weird Prompts, Corn the sloth and Herman the donkey tackle the "bigger is better" philosophy currently dominating the artificial intelligence industry. From the physical strain on global power grids to the bizarre phenomenon of "Habsburg AI" and model collapse, the brothers question if we are truly building a digital god or just a very expensive, very thirsty parrot. They dive deep into the differences between statistical prediction and genuine understanding, exploring why the next breakthrough in AI might require a total paradigm shift. Join the duo as they discuss Yann LeCun's world models, neuro-symbolic AI, and whether the future of intelligence lies in massive, monolithic data centers or specialized, efficient systems that actually comprehend the physical world we live in.

## DANIEL'S PROMPT

### Daniel

"We've discussed some of the challenges facing the large language model paradigm as the backbone of AI, ranging from prediction not being a real surrogate for intelligence to the risks of AI models training on bot-generated data and baking in their own mistakes. As LLMs continue to grow in parameter count and data ingestion, where does the 'pure scale' argument break down? Is the path to AGI really a 'brute force' approach of ever-increasing size, especially given the strain on compute and resources? What is the alternative vision to the 'bigger is better' argument, and what would those better models look like in the future?"



# TRANSCRIPT

## Corn

Welcome to My Weird Prompts! I am Corn, and I am joined as always by my brother, Herman Poppleberry. We are coming to you from our house here in Jerusalem, where the weather is actually quite nice today. I am currently leaning back in my favorite chair, moving at my usual slow pace, which is fitting since I am a sloth. But don't let the slow movements fool you, I am ready to dive into a big one today.

## Herman

And I am Herman Poppleberry, the donkey of the duo, and I am definitely moving a bit faster than my brother today because this topic is fascinating. Our housemate Daniel sent us a voice note this morning that really got me thinking. He wants us to talk about the limits of the current artificial intelligence craze. Specifically, the idea that just making these models bigger and bigger is the only way to reach true intelligence.

## Corn

Yeah, Daniel was asking about where this whole bigger is better philosophy actually breaks down. It feels like every week a company announces a new model with more parameters or more data, and they act like we are just a few more server farms away from a digital god. But is that really true? Or are we just building a bigger and bigger library that does not actually know how to read?

## Herman

That is a great analogy, Corn. Today we are going to look at the scaling laws, the risks of models eating their own data, and what the alternatives might be. Because honestly, the brute force approach is starting to hit some very real, very physical walls.

## Corn

I mean, from my perspective, it seems to be working pretty well so far. I used one of these large language models yesterday to help me write a polite email to the landlord, and it was perfect. If it keeps getting better at that, why should I care if it is just brute force?

### Herman

Well, hold on, Corn. There is a massive difference between a tool that is a world class statistical mimic and a system that actually understands cause and effect. What we have right now are essentially supercharged versions of the autocomplete on your phone. They predict the next word based on patterns. But prediction is not a real surrogate for intelligence. You can predict that the sun will rise tomorrow without having any clue what a star is or how gravity works.

### Corn

See, I actually see it differently. If the prediction is accurate enough, does the distinction even matter? If the AI predicts the right answer to a complex physics problem, is it not effectively intelligent? You are getting into the weeds of what is going on under the hood, but for normal people, the output is what matters.

### Herman

Mmm, I am not so sure about that. The problem is reliability. When these models fail, they fail in ways that a human child never would. They hallucinate facts with total confidence because their goal is not truth, it is probability. They are trying to find the most likely next word in a sequence. If the most likely word is a lie, they will say it. As we scale them up, we are just making those lies more sophisticated and harder to spot.

### Corn

Okay, but the proponents of scaling, the people who follow the scaling laws, argue that as you add more compute and more data, these emergent properties appear. Things the model was not specifically trained to do, it just suddenly starts doing. Is that not a sign that something deeper is happening?

### Herman

It is a sign of complexity, certainly. But it is not necessarily a sign of a conceptual breakthrough. There is this paper by researchers at companies like OpenAI and Google that talks about how performance improves predictably as you scale. But we are reaching a point where we are running out of high quality human data to feed them. We have scraped almost the entire public internet. What happens when the well runs dry?

### Corn

Well, they will just start using synthetic data, right? Data created by other AIs. That seems like an easy fix.

### Herman

That is actually one of the biggest risks, Corn. It is called model collapse. Imagine a copy of a copy of a copy. Each time you make a copy of a VHS tape, the quality gets worse. The same thing happens with AI. If a model trains on data produced by another AI, it starts to bake in the small errors and biases of that previous model. Over time, the model loses the ability to represent the diversity of the real world and starts to collapse into a weird, repetitive mess. Some researchers call it Habsburg AI because of the... well, the inbreeding of data.

### Corn

Okay, that is a bit gross, but I see the point. You are saying we are basically polluting the digital environment with AI trash, and then trying to eat that trash to grow bigger.

### Herman

Exactly. And it is not just the data. It is the physical cost. These models require an astronomical amount of electricity and water for cooling. We are building massive data centers that put a strain on the power grid. Is the path to artificial general intelligence really just building a computer the size of a city? That does not seem like an elegant solution. It seems like a desperate one.

### Corn

I don't know, Herman. Historically, brute force has worked for a lot of things. We didn't get to the moon with elegant tiny computers; we did it with a massive rocket and a lot of fuel. Maybe intelligence is just a matter of scale. Maybe the brain is just a very large, very efficient prediction machine.

### Herman

I have to push back on that. The human brain runs on about twenty watts of power. That is less than a dim lightbulb. An AI model that can barely pass a bar exam requires megawatts. We are orders of magnitude away from the efficiency of biological intelligence. If we want to reach true AGI, we need a paradigm shift, not just more chips.

## Corn

Let's take a quick break from our sponsors. Larry: Are you feeling untethered from reality? Do you wish your surroundings were just a little bit more confusing? Introducing Larry's patented Fog-In-A-Can. This isn't your grandfather's theatrical fog. This is industrial-grade, quantum-obfuscation mist. Perfect for hiding messy rooms, avoiding eye contact with neighbors, or creating an atmosphere of deep, unearned mystery in your studio apartment. One spray and you won't be able to see your own hands, let alone your mounting responsibilities. It smells faintly of ozone and old library books. Warning: do not use near smoke detectors or anyone you actually want to talk to. Larry's Fog-In-A-Can: because the truth is overrated. BUY NOW!

## Herman

Alright, thanks Larry. I think I still have some of that fog in my mane from the last time he visited. It is impossible to wash out.

## Corn

It does have a nice scent, though. Anyway, back to the topic. You were saying we need a paradigm shift. If bigger is not better, what is the alternative? What would a better model look like?

## Herman

Some researchers, like Yann LeCun at Meta, are looking into what they call world models. Instead of just predicting the next word in a sentence, these models would try to understand the underlying physics and logic of the world. Think about how a baby learns. A baby doesn't read a million books to understand that if you drop a ball, it falls. They observe the world. They build a mental model of reality.

## Corn

So you are saying we need to stop teaching AI language and start teaching it... life? That sounds even harder. How do you feed "life" into a computer?

### Herman

You give it video. You give it sensory input. You give it the ability to experiment and see the results. Right now, LLMs are like people who have been locked in a dark room with nothing but a library. They know everything that has been written, but they have never seen a sunset or felt the weight of an object. They lack common sense because common sense comes from interacting with the physical world, not from analyzing text strings.

### Corn

But wait, Herman, if we give them video and sensors, aren't we just giving them more data? Isn't that just scaling in a different direction? You're still talking about massive amounts of information.

### Herman

It is a different kind of information, though. It is grounded information. And the architecture would have to be different. Instead of just a transformer model that processes everything at once, you might have neuro-symbolic AI. This combines the pattern recognition of neural networks with the hard logic of symbolic AI. It is like combining the intuitive, fast-thinking part of your brain with the slow, logical, math-doing part.

### Corn

That makes more sense. I like the idea of the AI having a bit of a filter. But I still think you're underselling the current models. They can code, they can write poetry, they can even joke. If that's not a step toward intelligence, I don't know what is.

### Herman

It is a step, but it might be a step toward a dead end. If we keep pouring billions of dollars into making LLMs bigger, we might find ourselves with a very expensive, very thirsty parrot that still can't figure out how to tie its own shoelaces if it wasn't in the training data.

### Corn

Alright, we have got a caller on the line. Go ahead, you are on the air. Jim: Yeah, this is Jim from Ohio. I have been listening to you two go on about your world models and your neuro-symbolic whatever-you-called-it. Sounds like a bunch of malarkey to me. My neighbor Gary bought one of those smart fridges last year, and now it refuses to give him ice because it thinks he has had too much soda. That is the kind of intelligence you are building. It is not smart; it is just bossy.

### Herman

Well, Jim, that sounds like a programming issue with the fridge, not necessarily a failure of AI theory. But I think your point about "bossy" systems is interesting. As these models get more complex, they do become harder to control. Jim: Harder to control? You can't even control your own brother, Herman! He's a sloth for crying out loud! Anyway, I don't see why we need any of this. In my day, if you wanted to know something, you looked it up in an encyclopedia or you asked someone who knew. Now everyone's staring at their phones asking a robot where to buy socks. My cat Whiskers is smarter than most of these programs. At least he knows when it is time to eat without having to calculate the probability of a tuna can opening. Also, it's been raining here for three days straight and my basement is starting to smell like wet wool.

### Corn

Thanks for the call, Jim. I hope the basement dries out soon. And tell Whiskers we said hello.

### Herman

Jim is grumpy, but he touches on something important. The utility of these models. We are scaling them up to be these general purpose gods, but maybe we should be focusing on smaller, specialized models that are actually reliable and efficient.

### Corn

That is an interesting point. Do we really need one giant brain that knows everything? Or would we be better off with a hundred small brains that are experts in specific things? One for medical diagnosis, one for legal research, one for writing polite emails to landlords.

### Herman

Exactly! That is the alternative vision. Instead of a monolithic brute force approach, we could have a modular, efficient ecosystem. These smaller models are easier to train, they use less power, and they are much less likely to suffer from that model collapse we talked about because their data sets are curated and specific.

### Corn

But wouldn't they still have the same problem of not "really" understanding things? If my medical AI is just a small parrot, is it any safer than a big parrot?

### Herman

It is safer because it is constrained. You can verify its sources more easily. You can build in logic gates that prevent it from making things up. When you have a trillion parameters, it is a black box. You have no idea why it said what it said. With a smaller, specialized model, we can have more transparency.

### Corn

I see. So the future might not be a giant supercomputer in the desert, but a bunch of smart little programs living on our local devices. That sounds a lot more sustainable, honestly. I like the idea of my AI not needing a whole power plant just to help me find a recipe for lasagna.

### Herman

Exactly. And that brings us to the practical side of this. For the people listening, what does this mean for them? It means we should be skeptical of the hype. Just because a model is bigger doesn't mean it is better or more trustworthy. We are starting to see the diminishing returns of scale. The next big breakthrough in AI probably won't come from the company with the most GPUs, but from the company with the cleverest new architecture.

### Corn

I think that's a good takeaway. We shouldn't just assume the path we're on is the only path. There's a lot of room for creativity in how we build these things.

### Herman

There really is. We need to move away from the "bitter lesson" which is a famous essay by Rich Sutton that basically says the only thing that works in AI is scale and compute. We need to start looking for the "sweet lesson"—how can we make models that learn like humans do, with very little data and very little energy.

### Corn

The sweet lesson. I like that. It sounds much more pleasant. And much more sloth-friendly. I'm all about conservation of energy.

### Herman

I knew you'd like that part. But seriously, the environmental impact alone is going to force this change. We can't just keep doubling the power consumption of AI every year. The earth literally won't allow it.

### Corn

So, to wrap things up, we've looked at why the bigger is better argument is hitting a wall. We've got the data problem with model collapse, the resource problem with power and water, and the fundamental problem that prediction isn't the same as understanding.

### Herman

And the alternative is a move toward grounded world models, neuro-symbolic architectures, and smaller, specialized, more efficient systems. It's about working smarter, not just bigger.

### Corn

Well, I feel a lot more informed, even if I'm still going to use my LLM to write my emails for now. Thanks to Daniel for sending in that prompt. It's always good to check under the hood of the tech we use every day.

**Herman**

Definitely. It's easy to get swept up in the magic, but at the end of the day, it's all just math and silicon. And we're the ones who have to decide how much of our world we want to hand over to it.

**Corn**

Well said, Herman. You can find My Weird Prompts on Spotify, or on our website at [myweirdprompts.com](http://myweirdprompts.com). We have an RSS feed there for all you subscribers, and a contact form if you want to send us a prompt like Daniel did. Or if you just want to complain like Jim.

**Herman**

We'll take all of it! Even the complaints about wet basements.

**Corn**

Thanks for listening, everyone. We'll be back next week with another weird prompt to explore. Stay curious, and maybe don't trust your fridge too much.

**Herman**

Especially if it's bossy. Goodbye everyone!

**Corn**

Goodbye!