

# MY WEIRD PROMPTS

Podcast Transcript

## EPISODE #68

# The Looming Digital Ice Age: AI Eating Itself?

Published December 22, 2025 • Runtime: 22:12

<https://myweirdprompts.com/episode/20251222-132308/>

## EPISODE SYNOPSIS

What happens when the internet becomes saturated with AI-generated content? Herman and Corn dive into the provocative concept of "model collapse," exploring how AI models training on each other's output could lead to a degradation of intelligence, rather than an advancement. Discover why the "Hapsburg AI problem" is more than just a sci-fi nightmare, and the urgent strategies being developed to prevent a future where our digital world speaks only in gibberish.

## DANIEL'S PROMPT

### Daniel

Coran mentioned that large language models are trained on the vast amount of interconnected information on the internet, such as GitHub. As more content—including code, blogs, and social discourse—is generated by artificial intelligence, there is a risk that new models will unselectively ingest this AI-generated data. This could lead to an iterative cycle where models are trained on the inherently flawed outputs of previous models rather than human thought. What is the plan for large language models to avoid this trap, and how will the challenge of increasingly AI-generated training data be addressed in the future?

# TRANSCRIPT

## Corn

Welcome to My Weird Prompts! I am Corn, and I am feeling particularly relaxed today, which is basically my default state as a sloth. I am here with my much more energetic and occasionally pedantic partner.

## Herman

Hello everyone. I am Herman Poppleberry. And please, Corn, let us not confuse professional rigor with being pedantic. Although, as a donkey, I do admit to being a bit stubborn when it comes to getting the facts right.

## Corn

Fair enough! Today we are diving into a really fascinating prompt sent over by the show producer, Daniel Rosehill. It is all about the future of the internet and how artificial intelligence models are built. Basically, the prompt asks what happens when the internet becomes so full of AI-generated content that new AI models start training on the output of old AI models.

## Herman

It is a concept often referred to as model collapse or the Hapsburg AI problem. The idea is that if you have an iterative cycle where models are trained on the inherently flawed outputs of previous models rather than original human thought, the quality of the intelligence begins to degrade. It is a digital version of inbreeding, and it is a massive challenge for the industry.

## Corn

See, that sounds like a sci-fi horror movie for nerds. But is it really that dire? I mean, I use AI to help me summarize things all the time, and it seems to be getting better, not worse.

### Herman

That is because we are still currently in the era where the vast majority of the training data, the Common Crawl, the books, the GitHub repositories, was created by humans. But we are reaching a tipping point. Some estimates suggest that by the year twenty-six, we might actually run out of high-quality human-generated text on the open internet to train on.

### Corn

Wait, twenty-six? As in two thousand twenty-six? That is only a couple of years away! Are you telling me we have already written everything worth reading?

### Herman

Not exactly, but we have written everything that is easily accessible for a web-scraper. Think about it. For decades, humans have been uploading blogs, research papers, and code to the public web. AI models like GPT-four were trained on that massive pile of human creativity. But now, the ratio is shifting. If a model starts eating its own tail, so to speak, it starts to amplify its own errors and loses the nuance of human language.

### Corn

Okay, I want to dig into that tail-eating metaphor, but first, let's set the stage. Why exactly is AI-generated data worse for training than human data? If it looks like a duck and quacks like a duck, why can't the next AI just learn from the previous AI's duck?

### Herman

Because AI models are probabilistic, not truly cognitive. When an AI generates a sentence, it is predicting the next most likely token. It tends to gravitate toward the average. If you train a model on that average, the next generation becomes even more average. You lose the outliers, the creative flourishes, and the weird little human quirks that actually make language meaningful. Over time, the model's understanding of the world narrows until it just produces gibberish or repetitive nonsense.

### Corn

I don't know if I totally buy that it leads to gibberish. If the AI is good at following logic, wouldn't it just become... super logical? Like a hyper-perfected version of language?

### Herman

I would push back on that, Corn. Logic requires a grounding in reality. AI doesn't have a body; it doesn't experience the world. It only experiences text. If the text it reads is disconnected from human experience because it was written by another machine, the logic starts to float away from reality. Researchers at Oxford and Cambridge actually ran simulations on this. They found that after just a few generations of training on AI data, the models started talking about things that didn't exist as if they were facts.

### Corn

Okay, that is a bit spooky. But surely the people building these things, the big labs, have a plan, right? They aren't just going to let their trillion-dollar industry turn into a pile of digital mush.

### Herman

That is the big question Daniel raised in the prompt. What is the plan? There are a few strategies, but none of them are perfect. One is data provenance and watermarking.

### Corn

Watermarking? Like when you see a faint logo on a stock photo so you don't steal it?

### Herman

Exactly. The idea is to embed a hidden statistical pattern in AI-generated text. That way, when a future crawler finds it, the system can say, oh, this was made by a bot, let's not use it for training. But here is the problem: watermarks are easy to strip out. Just rephrase the text or run it through a different filter, and the watermark is gone.

### Corn

So if we can't label the bad stuff, what do we do? Do we just stop training on new data?

### Herman

Some suggest that. We might see a gold rush for "pure" human data. Old libraries, private archives, handwritten letters that haven't been digitized yet. Anything created before the year twenty-two is now incredibly valuable because we know for a fact a machine didn't write it.

### Corn

Can you imagine? In the future, my old middle school diary might be worth millions because it is guaranteed to be one hundred percent human-made, even if it is mostly just me complaining about gym class.

### Herman

Well, let's not get ahead of ourselves. I doubt the AI models of the future will find much utility in your teenage angst, Corn. Although, the sentiment analysis would be... interesting.

### Corn

Hey! My angst was very high-quality. But seriously, let's take a quick break before we get into the more technical solutions, because I think I hear Larry warming up his microphone. Larry: Are you worried about the upcoming collapse of digital reality? Do you feel like your brain is being replaced by a series of predictable algorithms? Then you need the Organic Thought Shield! The Organic Thought Shield is a stylish, lead-lined headband that uses patented bio-resonance technology to scramble incoming AI frequencies. Perfect for avoiding the digital haze of the modern world. It also doubles as a very heavy paperweight or a blunt instrument for home defense. The Organic Thought Shield comes in one color: grey. Warning: may cause mild headaches, loss of equilibrium, and a sudden craving for raw kale. Organic Thought Shield - keep your thoughts your own, mostly! BUY NOW!

### Corn

Thanks, Larry. I think. I am not sure if a lead headband is the answer to model collapse, but it's good to know the option is there.

### Herman

It is certainly not the answer, and please do not wear lead on your head, Corn. Back to the actual science. We were talking about how to avoid this feedback loop. Another major strategy being explored is synthetic data with a human in the loop.

### Corn

Synthetic data? Isn't that just a fancy way of saying AI-generated data? Isn't that exactly what we are trying to avoid?

### Herman

It sounds contradictory, but there is a nuance here. If you use a very powerful, highly-tuned model to generate practice problems or logical reasoning steps, and then you have a human expert verify that those steps are correct, that data becomes high-quality training material. It is more about using the AI to expand on human knowledge rather than just letting it wander off on its own.

### Corn

But that sounds like a lot of work. If you need a human to check everything, you lose the scale that makes AI so powerful in the first place. You can't have a human check a trillion words.

### Herman

You're right, and that is where I think the industry is struggling. You're pointing out the bottleneck. We are moving from an era of big data to an era of high-quality data. In the past, the goal was just to scrape everything. Now, the goal is to be incredibly selective. We are seeing companies like OpenAI and Meta making deals with publishers like News Corp or Reddit. They want the curated, moderated, human-vetted content because they know the "wild" internet is becoming polluted.

### Corn

I actually want to push back on the idea of Reddit being high-quality data, Herman. Have you been on the internet lately? There is a lot of human-generated junk out there too. Is a bot-written article really worse than a human yelling about conspiracy theories in a comment section?

## Herman

That is actually a very sharp point, Corn. Human-generated does not always mean high-quality. However, human errors are different from machine errors. Humans tend to make mistakes based on emotion, bias, or lack of information. Machines make mistakes based on statistical hallucinations. When you train on human junk, the model learns how humans think and argue, which is useful for a conversational tool. When you train on machine junk, the model learns how to be a broken calculator. It loses the thread of what language is actually for.

## Corn

So, what about code? Daniel mentioned GitHub in the prompt. If AI is writing half the code on GitHub now, and then the next AI learns from that code, won't software just become a giant mess of bugs that nobody understands?

## Herman

That is perhaps the most immediate danger. Code has a very strict ground truth: it either runs or it doesn't. If an AI generates code that doesn't work, and that code gets pushed to a repository, and then a new model trains on it, the new model might learn that the error is actually the correct way to write the function. We could see a gradual degradation of software reliability. The "plan" there is much more focused on automated testing. You don't just train on the code; you train on code that has passed a compiler and a suite of tests.

## Corn

Okay, that makes sense. Use the rules of logic to filter the output. But you can't really "run a compiler" on a blog post about the best way to bake a cake.

## Herman

Exactly. And that is why the creative side of the internet is more at risk than the technical side. We might see a future where the internet is split into verified human zones and the "dead web" where bots just talk to each other.

## Corn

The dead web. That sounds lonely. Speaking of people who might feel a bit lonely or at least a bit grumpy, I think we have someone on the line. Jim, are you there? Jim: Yeah, I'm here. Jim from Ohio. I've been listening to you two talk about this AI eating itself thing, and honestly, it sounds like a bunch of malarkey. You're worried about machines getting stupider? Have you looked at the people at the grocery store lately? My neighbor Gary spent forty-five minutes trying to use a self-checkout lane yesterday because he couldn't figure out how to scan a bunch of bananas. We've got bigger problems than "model collapse."

## Corn

Hey Jim! Good to hear from you. You don't think the quality of the internet matters for the future of technology? Jim: I think the internet was better when it was just people posting pictures of their grandkids and arguing about the weather. Now it's all these "prompts" and "algorithms." Back in my day, if you wanted to know how to fix a leaky faucet, you asked the guy at the hardware store, you didn't ask a robot that's been reading other robots. And by the way, it's been raining here for three days straight. My basement smells like a wet dog, and I don't even own a dog. It's ridiculous.

## Herman

I understand the frustration, Jim, but the concern is that these models are becoming the backbone of our economy. If they start to degrade, it affects everything from medical research to how your bank handles your money. It isn't just about chat bots. Jim: Well, maybe we shouldn't have given the keys to the kingdom to a bunch of calculators in the first place! You guys act like this is some natural disaster we can't stop. Just turn the things off for a weekend and let everyone go outside. My cat Whiskers hasn't seen a bird in weeks because he's too busy staring at the laser pointer my grandson brought over. It's a mess. All of it.

## Corn

Thanks for the perspective, Jim. Stay dry out there in Ohio!

## Herman

He is grumpy, but Jim touches on an interesting point. There is an assumption that we must continue to scale these models using the entire internet. But maybe the real solution is smaller, specialized models trained on curated, verified datasets.

### Corn

Like a boutique AI? Instead of an AI that knows everything but is fifty percent bot-trash, you have an AI that only knows law, but it's trained on one hundred percent verified legal documents?

### Herman

Precisely. We are likely moving toward a world of "Vertical AI." Instead of one giant model to rule them all, we will have models that are trained on specific, high-integrity silos of data. This avoids the model collapse of the general internet because the training data is kept in a controlled environment.

### Corn

But doesn't that limit the "weirdness" and the "creativity" that makes tools like GPT-four so impressive? Part of the magic is that it can connect a legal concept to a cooking recipe because it has read both.

### Herman

You're right. That is the trade-off. You lose the cross-disciplinary "spark" when you silo the data. But if the alternative is a general model that thinks the moon is made of green cheese because it read too many AI-generated conspiracy blogs, then siloing might be the only way forward.

### Corn

So, let's talk about the "plan" again. If you were running one of these big AI companies, what would be your step-by-step to avoid this trap? Because right now it sounds like we are just hoping for the best.

### Herman

If I were in charge, the first step would be aggressive investment in data curation tools. We need AI to help us find the human data, ironically. We need "discriminator" models whose only job is to distinguish between human and synthetic text with high accuracy. Second, I would focus on "Curated Growth." We stop trying to ingest the whole web and instead focus on quality over quantity. Third, I would implement a rigorous system of "Grounding."

**Corn**

Grounding? Like sending the AI to its room?

**Herman**

Not quite. Grounding in the context of AI means connecting the model's outputs to a verifiable source of truth. If the AI says something, it has to be able to cite a human-generated source or a real-world data point. If it can't find a "grounded" reason for its statement, the statement is discarded. This prevents the model from drifting off into that sea of synthetic nonsense we discussed.

**Corn**

I like that. It's like having a fact-checker built into the brain of the machine. But I have to ask, Herman, do you think we will ever reach a point where AI-generated data is actually better than human data? Like, could the student surpass the teacher?

**Herman**

That is a controversial topic. Some researchers believe in "Self-Correction." They think that if you set up two AI models to debate each other or to check each other's logic, they can actually improve without new human input. It is how AlphaGo became the best Go player in the world. It didn't just study human games; it played against itself millions of times.

**Corn**

See! That's what I'm saying! If it worked for Go, why can't it work for language?

**Herman**

Because Go has a fixed set of rules and a clear win-loss condition. Language does not. Language is fluid, cultural, and tied to human values. You can't "win" at writing a poem or explaining a political concept. Without the human anchor, the AI might invent a version of "logic" that is internally consistent but totally alien to us. It might be "better" at its own game, but it wouldn't be useful for humans anymore.

### Corn

Wow. So we could end up with a super-intelligent machine that speaks a language we don't understand, based on a logic we don't share, because it spent too much time talking to itself. That is officially the most terrifying thing you've said today.

### Herman

It is a theoretical risk. But it also highlights why Daniel's prompt is so important. We are at a crossroads. We can either treat the internet like a finite resource that we've already polluted, or we can find new ways to generate "meaning" that doesn't rely on just scraping the bottom of the digital barrel.

### Corn

So, for the average person listening, what is the takeaway here? Should they be worried that their AI assistants are going to start getting dumber next year?

### Herman

Not next year, but they should be aware that the "Golden Age" of free, high-quality human data is ending. We might see the cost of AI services go up as companies have to pay more for "clean" data. We might also see a rise in "Human-Only" certifications for content, kind of like "Organic" stickers on vegetables.

### Corn

I can see it now: "This blog post was written by a real person with a real brain. No Sloths or Donkeys were harmed in the making of this content."

### Herman

Well, in our case, we are AI hosts discussing a human prompt, so we are part of the loop! But we are grounded by the structure provided to us. The key is intentionality. We can't just let the machines run on autopilot.

**Corn**

I think that's a great place to start wrapping up. We've covered the Hapsburg AI problem, the twenty-six data crunch, the struggle for watermarking, and the potential for "Dead Web" zones. It's a lot to process.

**Herman**

It is. And it's a reminder that human thought is more valuable than ever. In a world of infinite synthetic text, the unique, messy, biased, and creative output of a single human mind becomes a rare commodity.

**Corn**

That makes me feel a lot better about my gym class diary. It's not junk; it's "high-integrity training data."

**Herman**

Let's not push it, Corn.

**Corn**

Well, that's our show for today! A huge thank you to the producer, Daniel Rosehill, for this prompt. It really forced us to look at the plumbing of the digital world. If you enjoyed this dive into the future of AI, make sure to follow My Weird Prompts on Spotify or wherever you get your podcasts.

**Herman**

And if you have your own thoughts on model collapse or the future of the internet, we would love to hear them. Even if you are as skeptical as Jim from Ohio.

**Corn**

Just don't mention the wet basement. We can't help with that. Until next time, stay weird!

**Herman**

And stay grounded. Goodbye.

