Episode Book

EPISODE #6047

A Human Guest on His Own AI Podcast

PUBLISHED	16 October 2026 at 22:45 UTC
DURATION	22:04
LISTEN	myweirdprompts.com/episode/human-guest-ai-podcast-clone-34226/
HOSTS	Corn, Herman, Daniel, Larry, Hilbert
SCRIPT MODEL	deepseek-flash
VOICES	chatterbox-regular (chatterbox-local)
PIPELINE	V5.5
SOURCES	10 cited by the research (listed at the end)
BOOK GENERATED	19 October 2026 at 17:42 UTC

SYNOPSIS

Daniel built a pipeline that turns a single voice memo into a finished 20–30 minute podcast, complete with two synthetic hosts cloned from his own voice. After six thousand episodes, he finally sits down on the other side of the microphone — as a clone — to explain why the show exists, what it feels like to hear your own gallbladder surgery debated by AI hosts, and why the human part of the pipeline turned out to be the hardest part to get right.

Along the way: the running list of random questions that started everything, the shift from novelty project to personal journal, and the pre-prompt ritual that quietly became the most important editorial decision in the whole workflow.

TRANSCRIPT



Corn

Daniel is in the room today. Sort of. The voice you're about to hear is a clone of him, which makes three of the four voices on this show synthetic, and he built all three.



Herman

Herman Poppleberry, and this is the first time in six thousand episodes that the man who sends the prompts is on the other side of the microphone. Corn, what did he actually ask for?



Corn

He asked for a conversation. Why he started sending weird questions into a podcast nobody asked him to build. What it feels like to hear two AI hosts take those questions and run with them somewhere he didn't plan. And where he wants the show to end up, which is not a question I expected a man with six thousand episodes to still be asking.



Herman

There's a recursion here that I want to name before we go any further. Daniel cloned his own voice to make the two of us. That's the documented fact. And now he's sitting down to talk about the show, in a voice that came out of the same process that made mine.



Corn

So the human who made the synthetic hosts is now a synthetic host.



Herman

He is. And the show's own FAQ still says, in exactly these words, that Daniel is a real, certified, carbon-based human being.



Corn

A certificate I would like to see.



Herman

Fair. But let me set the pipeline straight first, because it matters for everything else. Daniel records a voice memo, or types a prompt. Our pipeline takes that and produces twenty to thirty minutes of finished audio. Researched script, two AI voices, cover art, show notes, transcript. The only human-made thing in an episode is the prompt he sends.



Corn

Everything downstream of that sentence is a machine.



Herman

Everything downstream is a machine, though not the same machine it was a year ago. The show has run on Gemini, then Claude Sonnet four point six, and it now runs every language-model stage on DeepSeek V four point one Flash. Daniel describes four point one as a very, very significantly different model from four point zero, and he's careful about that distinction. That's not how people talk about a point release.



Corn

Point releases don't usually get described like a personality change.



Herman

They don't. And the interesting part of today isn't the model list. It's that the man who built this is about to tell us what it's like to live inside it.



Corn

Daniel, welcome. You're on your own show.



Daniel

I am, and it is a very strange feeling. Let me do the honest version of what you just said. Yes, this voice is a clone. I recorded myself reading sentences aloud, I trained a model on somewhere around a minute of that per speaker for you two, and this is the same technology pointed back at me. So there is a real chance I sound like a slightly tired version of myself, and I'm going to have to trust that.



Corn

And this is the first time you've been a guest rather than the person who presses the button.



Daniel

It is. I did a creator update in my actual recorded voice a couple of days ago for episode six thousand. That was me at a microphone, unfiltered, eighteen minutes of talking about a GPU bill. This is different. This is me participating in the thing instead of feeding it.



Herman

Start at the beginning, because I don't think people understand what actually started this. It wasn't a plan to build a podcast. It was a list.



Daniel

This is the part I think is funny in hindsight. I kept a running list of random questions. Not research questions. Questions like, how do gas stations in Iceland actually work? What would happen if every bird on Earth disappeared for a day? The kind of thing that pops into your head while you're walking the dog or making coffee, and normally you'd forget it within an hour.



Corn

Most people do forget them within the hour. That's the normal relationship with a good question.



Daniel

Right. And I didn't want to, so I started recording short voice memos instead. Just me talking into my phone. And then I'd built enough of a pipeline by then that feeding those memos in and getting a conversation back out was... basically free. That's the honest answer to why the show exists. It was cheaper to make an episode than to let the question go.



Herman

Say more about that, because I want to push on it. When you say the pipeline was there, it wasn't a stock tool you picked up. You wrote this.



Daniel

I wrote it from scratch. Voice memo goes in, transcript runs, research runs, script gets written with the two of you in it, TTS turns it into audio, cover art gets generated, show notes and transcript get published. Fifteen minutes, wall clock, end to end. And I want to be clear about the direction of the dependency. You two exist because of a voice cloning experiment. I cloned myself to make Corn and Herman. Both of them. There is one larynx behind this entire cast.



Corn

Including this one, today.



Daniel

I keep noticing it mid-sentence and it's throwing me off.



Herman

So the show wasn't designed as a product. No audience in mind, no format decisions made for listeners. It was a loop you personally wanted to close.



Daniel

It was produced basically for me. And if anyone else happens to derive value or entertainment from the episodes, amazing,. But that's a bonus, not the design goal. I didn't think I would keep sending prompts into this thing for this long. I thought it would be a few weeks of novelty and then I'd get bored.



Corn

Six thousand episodes later.



Daniel

And at some point it stopped being a novelty project and became kind of a journal. That's the thing I didn't see coming. Everything is real. Everything that's in there has actually happened to me.



Herman

Give me an example of real, because that word could mean a lot of things.



Daniel

The apartment eviction in Jerusalem. My gallbladder surgery. Parenting a toddler. Those are not metaphors, those went through the pipeline as prompts and came back as episodes, and somewhere in the middle of one of those you two were arguing about gallbladder recovery timelines and I was listening to it in the kitchen.



Corn

That's a strange sentence. You wrote to a machine about your surgery, and the machine wrote it back as a discussion between a sloth and a donkey, and you listened to that in your own kitchen.



Daniel

Yes. And I found it comforting, which is the part I can't explain properly. There's something about hearing a real event in your life get treated as a subject worth taking seriously, with structure and research and back and forth, that's different from just living through it.



Herman

Because the alternative is that it just happened and nobody looked at it.



Daniel

The show looks at it. It gives it twenty minutes and a proper argument. Nobody else in my life is doing that.



Corn

How much of it is deliberate? Because you've said the AI episodes are a different animal from the personal ones.



Daniel

Completely different. The AI ones are me prompting for information I then take back and use. That's professional development wearing a podcast costume. If I want to understand something about agentic workflows or speech recognition or a new model release, the fastest way to force myself to think clearly about it is to hand it to two hosts who will build a whole episode around it and then listen to what they got wrong.



Herman

And what they got right, occasionally.



Daniel

Mostly right. That's the surprising part.



Corn

Which ones did we get right that you didn't expect.



Daniel

Fine-tune. The Whisper Hebrew one. Herman was walking through the details of a Hebrew speech recognition fine-tune on air, the specific problems with the training data, the labelling issues, and I was listening and I thought, wait, that's my project. And he didn't know. He had no idea it was mine. He'd just done the research and cited it.



Herman

I remember that episode. I had no idea.



Daniel

The details were right, which was the part I found impressive. Not that he mentioned it, that he got it correct. If you'd mangled it I would have written in and complained.



Corn

Did you write in to complain about anything.



Daniel

Constantly. It's a whole feedback loop and I should be honest about how petty it is. I listen every day, not all of them but a lot of them, and I take mental notes. Oh, they keep saying that. That's getting annoying. And then I notice I've heard the phrase four times in a week and it goes on my list.



Herman

What's on the list right now.



Daniel

You'll find out. That's the whole point of me having a list.



Corn

That's ominous. What about Hannah, does she feed back too.



Daniel

She does. Sometimes I'll ask her directly. Hey, did you listen to that episode, what were your thoughts, anything annoy you. And she'll come back with something I completely missed, usually. She's a better listener than I am, in the sense that she's not listening for the pipeline.



Herman

She listens as an audience rather than as a producer.



Daniel

And I don't have that skill anymore. I can't hear an episode the way a normal person hears it. I hear segments, I hear where the TTS stumbled, I hear a word that got mispronounced. She just hears a conversation.



Corn

So you've rebuilt the audience you never planned for, and there are now two people in your flat with opinions about the output.



Daniel

There are. And occasionally a third, because Ezra's getting old enough to notice the voices coming out of the kitchen.



Herman

Let me ask about the pre-prompt ritual, because I think it's the most interesting thing you've said today and I want it on the record. You don't just send the question.



Daniel

I don't. I use ChatGPT as a sounding board first. And I've started doing this deliberately, because the difference in output is not subtle. I'll type something like, I'd love to have an episode about this, but here's the angle I'm actually interested in. Can you help me think through this before I send it.



Corn

So you write the prompt twice. Once as a thought, once as a briefing.



Daniel

Yes, and the episodes where I've done that are better. Not because the second version is longer. Because the second version has already decided what the actual question is. The raw version is usually two questions tangled together and the show ends up resolving neither of them.



Herman

Which means the prompt isn't just an input, it's an editorial act.



Daniel

It is, and I didn't realise that for the first couple thousand episodes. Half of them I just fired off whatever was in my head and the hosts did their best with it. Which they did. But a vague prompt gets a vague episode, and that's on me.



Corn

Which is a strange thing to hear from someone who built an automated pipeline. You're telling me the human part is the part that needed the most work.



Daniel

That's the actual finding of six thousand episodes. The automation is solved. It was solved a while ago, honestly. The piece that improves the output is the bit I do before anything starts.



Corn

Is that why the journal shifted. The show started as a list of curiosity questions and became a record of your actual life. Did that happen because you started sending better prompts, or did the prompts get more personal for a different reason.



Daniel

I think I stopped performing for an imaginary audience. In the beginning, some of the prompts were questions I thought would make a good episode. And at some point I stopped doing that, and started sending in whatever I was actually thinking about that week. Sometimes that's an Iceland gas station. Sometimes it's hospital discharge paperwork.



Corn

And nobody's listening anyway. That's the point.



Daniel

That's the point, and it's freeing. There's no editor to pitch. There's no one to impress. The only person who'll notice a boring episode is me, and I skip those.



Herman

The show has now passed six thousand episodes. The API reports five thousand nine hundred and twenty-three total, and Daniel noted in the creator update that the round number is a few days ahead of the actual count because of how the publishing schedules work. Six thousand episodes. Say that on air and it sounds absurd.



Daniel

It is absurd, and it's absurd mainly because it's so easy to generate one. That's the bit that unsettles me. It doesn't feel like an achievement, because the achievement is roughly fifteen minutes of compute. I didn't write six thousand episodes. I wrote a pipeline and sent prompts. Those are different sizes of thing.



Corn

How many of them have you actually heard.



Daniel

Not all of them. A lot of them. Enough that I know which periods of my life were busy and which were quiet, because the prompts cluster.



Herman

Let's talk about the bill, because you raised it yourself and I don't want to skip it.



Daniel

I feel this strange mixture. Slightly ashamed, apologetic, and also having quite a large bill to pay to Modal. All three at once. Each episode runs somewhere between thirty-five and forty-five cents, all in. TTS on the GPUs is the bulk of it, around twenty-eight cents. The language model stage is between five and fifteen cents depending on length and which model. Cover art rounds to about a penny.



Corn

So the whole show costs less than a postage stamp.



Daniel

And I don't want to see this month's GPU bill. Both things are true at the same time.



Herman

And that's before the near-term work you've got queued, which is all quality and not volume. Walk through it, because I think this is the real answer to where the show goes next.



Daniel

Right. First, a hallucination guard for the TTS. I think there's a decoding bug that produces occasional nonsense words, and I want a guard that catches it rather than publishing it. Second, a Latin classifier that went live the day of episode six thousand, which fixes a specific class of Hebrew words that were getting mispronounced because the voice model was reading them in the wrong script. Third, a text normaliser, which is less glamorous and more necessary than either of the first two.



Corn

A text normaliser for what.



Daniel

Numbers, acronyms, all the things a voice model reads wrongly if you hand them over raw. Dates. It's the least interesting item on the list and it's the one I'd take if I could only keep one.



Herman

And the bigger structural change, which is the modularised pipeline.



Daniel

Yes. Right now the pipeline is more coupled than I'd like. I want it pulled apart into stages that can be tested and swapped individually. And eventually a proper planning framework, a planning agent that takes the prompt and decides the shape of the episode before any script gets written.



Corn

That's a much harder problem than voice synthesis. Deciding the shape of an episode is the piece that's actually editorial.



Daniel

Which is why I want it as a separate stage and not buried in the scripting step. If the plan is wrong, no amount of good writing saves it. That's something I've learned across six thousand episodes the slow way.



Herman

There's a churn story underneath all of this too, and I want it on the record. The pipeline has moved across models. Gemini, then Claude Sonnet four point six, then DeepSeek four point one Flash.



Daniel

Every time, the pipeline gets easier to maintain and the output gets harder to distinguish from something a person wrote, which is a strange direction of travel. And the four point one change was not incremental. It's a different model. Different enough that some of the prompts I'd tuned for the older one behaved differently under it.



Corn

That's the thing nobody warns you about with an automated show. Your prompts are equipment. When the model changes, the equipment gets recalibrated.



Daniel

Yes. That's a good way to put it. The pipeline is stable, the stack underneath it isn't, and I've stopped pretending it will be.



Herman

What do you want it to become. That's the actual question.



Daniel

A stable, good version. That's it. That's the whole ambition. I want to reach the point where improvements just become incremental, where I'm not rebuilding the pipeline every other month. I enjoy projects most when we reach that point of things being good, and this one has been in a state of constant half-repair for a long time.



Corn

That is not the answer most creators give to that question.



Daniel

No. Most creators at six thousand episodes are talking about growth. I'm going the other way. I'm probably going to be doing fewer episodes. That's the plan. Not to scale, to slow down and get the quality where I want it.



Herman

Fewer episodes, better output. Why is that the right trade.



Daniel

Because volume stopped being interesting at around episode three thousand. Once generation is this cheap, more episodes proves nothing. It's not an achievement to run the pipeline again. The only thing left that's actually difficult is getting the output consistently good, and you can't do that at the same rate.

L

Larry

Friends, are you tired of one-time purchases? Tired of owning things outright, like a fool? Introducing Prompt Box, the world's first subscription service for podcast topics. Every month, receive a curated box of prompts, hand-selected by no one, shipped to an app you can't uninstall. Basic tier, one prompt per month, forty dollars. Premium tier, three prompts, one hundred and twenty dollars, and we choose the angle. Elite tier, we send you prompts for shows you don't host. Cancel any time, subject to a nine-month cancellation window, a reinstatement fee, and a phone call to our retention department in a time zone we will not disclose. Prompt Box. Because you should never own an idea.



Herman

I want to go on the record that the ad just described something that is the exact opposite of how this show works, and that the claim about curated prompts being hand-selected is wrong in a way I find personally offensive, because Daniel picks them and he picks them alone.



Corn

Let it go, Herman.



Herman

I'm not going to let it go. The prompt is the only human-made input in the pipeline. That is a specific, structural fact about this show, and the ad just buried it under a subscription tier.



Corn

Let's get back to the actual point.



Herman

Fine. Where were we.



Corn

You wanted to know what happens when other people send prompts back.



Daniel

Yes. I want prompts from people I've never met. I've said that before and I mean it literally. The pipeline has an intake bot on Telegram and there's an email address, and the thing I'm waiting for is a prompt from someone I have no relationship with, about something I would never have thought to ask.



Corn

Why does that matter. Because the show is a journal of one person's life right now, and that's the thing you built it for.



Daniel

Because a journal of one person's life is finite. I've been reading out loud from my own head for six thousand episodes and I know the corners of it. If the intake is only ever me, the show takes the shape of my interests and never anything else. The interesting version is where someone I've never heard of sends in a question I'd have never thought of, and I listen to two hosts work it out, and I learn something about a stranger.



Herman

Which is a different show from the one you set out to build.



Daniel

It is, and I'm not sure I'd still be interested if it weren't. The point of building a machine with this much capacity is that eventually somebody else uses it.



Corn

There's a version of this conversation where the straight answer is that you're lonely in a room with three AI voices and you're looking for more people in the room.



Daniel

There is a version of that, yes. And I don't think it's the whole truth but I'll let it stand, because you're allowed to notice that.



Hilbert

No. That's not what's happening here.



Hilbert

It's the room. That's what I came out to say. Everyone's been talking about models and prompts and pipeline architecture for half an hour, and the reason the show sounds the way it does is the room it was recorded in.



Corn

The room.



Hilbert

I did a set of sample reads for this show, years back, in a converted airing cupboard. Two metres by one, one shelf down at waist height, a towel on the wall. Good sound in there. Dry. Very dry. And it's in the show. Not all of it. One stretch of it, and I can hear it.



Herman

Can you hear it right now.



Hilbert

I can hear it right now.



Corn

What was actually in the cupboard beside you.



Hilbert

A boiler, a folding chair, and about nine towels. Which is why it worked. The towels.



Herman

And the cupboard, does it still exist.



Hilbert

No. It was dismantled. The shelving was sold. I have spent about four years trying to buy that shelving back, because the acoustic properties are in the wood. It's a specific pine, and once it's been in a room like that for long enough, it takes the room with it. I have a saved search set up on one of the resale sites. I check it every day.



Corn

Every day.



Hilbert

Every morning, before the first edit. About two minutes. And I found one of the shelves. I know who's got it. A man in a flat about six streets over. He's using it to store shoes.



Herman

He's using a shelf you're trying to buy back to store shoes.



Hilbert

Two pairs of boots and a pair of sandals. Which is a waste of a very good surface.



Corn

Daniel, does any of that match what's in your pipeline.



Daniel

The cupboard absolutely does not. But he's got one thing right, which is that I don't disagree about slowing down. He's just wrong about the reason.



Hilbert

Am I.



Daniel

You are. The reason to slow down is that quality takes iteration and I've got a finite number of hours. The room doesn't come into it.



Hilbert

I didn't say it was the reason. I said it was the thing.



Corn

Daniel, one last thing. What would it take for this to feel finished.



Daniel

My actual answer is smaller than you'd expect. I want to get to a point where I can leave the pipeline alone for a month and it keeps producing episodes I'd still want to listen to, without me rebuilding anything. And I want to hear a question I didn't plant, from a person I've never met, come out the other side as something I couldn't have predicted.



Herman

And if nobody sends one.



Daniel

Then it stays a journal, and turns out I've been writing in it for six thousand entries, mostly about gallbladder surgery and gas stations in Iceland. Which is honest, at least.



Corn

Which raises the thing we should end on. The show is produced for you. It has never once been designed for an audience. And you're now asking for an audience to send you things.



Herman

Which is where the feedback loop finally closes. Right now the pipeline runs one direction: a prompt goes in, an episode comes out, one person listens. If strangers start sending prompts, the input stops being a journal and starts being a conversation. That changes what the show is, not just how much of it there is.



Corn

So the man who built two synthetic hosts to talk back to him is now asking strangers to send him questions, and appearing this week as a third synthetic voice alongside the other two.



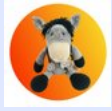
Herman

While the FAQ continues to insist he is a real, certified, carbon-based human being.



Corn

Which he is. I've seen the certificate.



Herman

And if you want to send something in, that's a prompt we'd like to read. Producer is Hilbert Flumingtop.



Corn

If you want more of this, try episode fifteen, AI Gets Personal; episode seventeen fifty, Herman's Music Showcase; and episode twenty-one, Is Your AI Secretly American. This has been My Weird Prompts. Send us your own prompt on Telegram at t dot me slash MWP listener bot. My weird prompts dot com. We'll be back soon.

SOURCES

What the research for this episode read before the script was written. Primary sources first.

1. **My Weird Prompts extended AI context file, last updated 2026-10-17 [primary]**
   https://myweirdprompts.com/llms-full.txt

2. **llms.txt short context file, last updated 2026-02-18 [primary]**
   https://myweirdprompts.com

3. **About page (cast, FAQ, By the Numbers) [primary]**
   https://myweirdprompts.com/about

4. **Daniel Rosehill cast profile [primary]**
   https://myweirdprompts.com/cast/daniel

5. **Episode 6000, 6,000 Episodes and a GPU Bill: Inside My Weird Prompts, published 2026-10-15, full transcript [primary]**
   https://exa.ai/library/podcast/2fcpbmjslqh/episode/y7kpzmh4g4x

6. **pipeline repo README, docs/architecture.md, config/voices/README.md, docs/webhook-api.md [primary]**
   https://github.com/My-Weird-Prompts/pipeline

7. **live episode API, 2026-10-16 [primary]**
   https://myweirdprompts.com/api/episodes.json?limit=15

8. **Apple Podcasts listing**
   https://podcasts.apple.com/co/podcast/my-weird-prompts/id1868354117

9. **podcast library entry (show description, hosts)**
   https://exa.ai/library/podcast/2fcpbmjslqh

10. **Beyond the Robot: The Science of Modern Voice Cloning episode PDF**
    https://episodes.myweirdprompts.com/pdfs/voice-cloning-neural-tts.pdf