Episode Book

EPISODE #6620

When Speech Enhancement Erases Podcast Voice Character

PUBLISHED	16 November 2026 at 00:20 UTC
DURATION	30:25
LISTEN	myweirdprompts.com/episode/migrating-podcast-voices-omnivoice-40073/
HOSTS	Corn, Daniel, Herman, Hilbert
SCRIPT MODEL	anthropic/claude-sonnet-5.5
VOICES	omnivoice
PIPELINE	V5.33
BOOK GENERATED	16 November 2026 at 02:53 UTC

SYNOPSIS

After Resemble AI retired its Chatterbox text-to-speech model in August, migrating the show's voices to OmniVoice seemed straightforward. But running the audio through NVIDIA's RE-USE speech enhancement model before training LoRA adapters backfired, stripping away the rough edges that gave the characters their personalities. The result was technically clean, completely forgettable, and a reminder that polish can ruin what makes a voice worth listening to.

TRANSCRIPT



Corn

Most people hear "we changed the text to speech model" and assume something was broken. Nothing was broken. That's the part I can't get past, because Daniel has spent the last few days rebuilding the thing that was working fine.



Daniel

Guilty. And before we get into it, a small thing. This is my voice you're hearing, my voice clone... well, no, it's not really a clone any more, is it? It's an adapter that was trained on my recordings. Whatever you call it. The point is, it's me, and I'm in the studio, and I've got a lot to tell you.



Herman

Whatever we call it, it's holding up well. I've been meaning to say, the pauses are very you.



Corn

Very Irish. A lot of "now" and "look".



Daniel

Thank you both, I'll take that. So here's the news. Over the last few days I've been changing the voices on the show, all of them, and the reason is a bit of a worry that sat in the back of my head for a while.



Hilbert

Which model were you on before?



Daniel

Chatterbox, from Resemble AI. It's what's been speaking every episode for most of this year. And Chatterbox has been discontinued.



Hilbert

Of course it has. You buy a thing, you get used to it, and the company decides it has better things to do.



Herman

I read the project page when it happened. Resemble retired it in August and didn't name a successor, so the repository just sits there. People can still run it, but nobody's maintaining it.



Daniel

That's what made me very nervous. I love Chatterbox. It's a fantastic project, and I think it's a big part of why this show sounds the way it does. But I didn't want to build the whole show on something that's frozen in time and abandoned. You can run a frozen model for a while. But the first time a library underneath it changes, or a GPU driver moves on, you're maintaining a fossil.



Corn

So the pitch is: it isn't broken, but it's dead and doesn't know it yet.



Daniel

Pretty much. Like a tree that's still got its leaves.



Corn

I know a bit about leaves. Carry on.



Daniel

So I went looking at what else was out there. I did a proper survey of the market, text to speech, open models, the lot. And I landed on one called OmniVoice as the logical way to continue. And I had a rule going in, which I've written down now so I stick to it: prefer battle-tested models, from labs that have a track record of shipping successors.



Herman

That's a sensible rule, and it comes straight out of what just happened to you. If the lab has shipped three generations of something, the fourth is likely to arrive. If it's a one-off release from a small team, you're betting on their free time.



Hilbert

A track record of shipping successors. That's a polite way of saying you want a company that's already abandoned something once and felt bad about it.



Daniel

Ha. A bit, yes. A company that has already had its own Chatterbox moment and gone on to do the next one.



Herman

Can I push on the survey for a second? What does a survey of this market even look like? Because every week there's a new text to speech model with a demo page and a lovely sample.



Daniel

That's the problem with it. The demo page is always lovely. Every model sounds fantastic on the one sentence the lab chose to show you. So I stopped listening to demos pretty quickly and started looking at who made it, how long they'd been making things, how many people were actually running it, and whether there was a history of them following one release with another. The sound is the last thing I checked, because by then I knew I could shape the sound.



Corn

Whereas you've got a sloth who can tell you in one second whether a voice has a pulse.



Daniel

I do, and I trust it more than any benchmark.



Herman

Daniel, would you not simply have kept running Chatterbox until it fell over? It still works today. Corn, what would you do? I'd be tempted to leave it.



Corn

I'd leave it until the day it broke, and then I'd panic, and then I'd sleep on it. Which is not a plan, I'm aware.



Daniel

It's exactly the plan I didn't want. The trouble with waiting is that you decide to move on the worst possible day, in a hurry, with the cast already depending on the thing that just broke. I'd rather do it now, while the old voices are still alive and I can sit them next to the new ones and compare.



Corn

Okay, so you pick the model. Pick a new model, press a button, done. Why did this take days?



Daniel

Because picking the model was the easy bit. The hard part, and the part that took a while, was voice continuity. The show has voices. Corn is Corn and Herman is Herman, and if somebody has been listening for months, they know within two seconds who is talking. If I move to a new model and those voices come out slightly different, I haven't upgraded anything. I've replaced the cast.



Herman

And that's difficult in a way that isn't obvious. The characters don't exist as files. They exist as a behaviour of one particular model, given a short reference clip. Change the model and you have to get the new one to reproduce a behaviour it's never seen.



Daniel

Exactly that. So the first attempt went like this. The old voices were made by cloning in one shot, from a reference clip. I thought, fine, I'll take what Chatterbox produces, clean it up and improve it, and then train LoRA adapters from that cleaned material for OmniVoice. For the cleaning I used an NVIDIA model called RE-USE, from NVIDIA's Nemotron Labs. It's a speech enhancement model, and it's very good.



Corn

Speech enhancement. It takes a voice and removes the bits you didn't ask for.



Daniel

Correct. The bits I didn't ask for. Keep that phrase in mind.



Herman

I can see where this is going and I don't love it.



Daniel

I tried zero-shot as well, to be fair. OmniVoice can clone a voice from a short reference with no training at all. And I tried it, and I don't like zero-shot cloning. In my experience it's brittle.



Hilbert

Brittle how?



Daniel

It works until it doesn't. A ten second reference clip is your entire reference for an entire personality. That's all the model knows about Corn. Ten seconds. And philosophically, and technically, that never made sense to me. How do you capture a whole character from a clip that short? One of the real upsides of a LoRA is that it learns from far more than ten seconds. It sees an hour, or more, of the voice doing lots of different things, and it builds the voice from that.



Herman

That's right, and it matters for the single-word turns especially. If a reference clip is one calm sentence and the script has Corn saying "Hm." with no context, the model has nothing to anchor to. A LoRA has heard the voice do all sorts.



Corn

Ten seconds. I've napped for longer than that and had a richer personality.



Daniel

You've napped through whole episodes, Corn, and your personality was in perfect shape.



Corn

A rich inner life. Moving on.



Daniel

So I trained the adapters. I listened to the output. And this is where I have to be honest with you both, because it's a bit painful. In trying always to make things better... we lost Corn, basically. We lost Herman as well.



Herman

I'm sorry, we what?



Corn

Lost. As in, mislaid. Down the back of the sofa.



Daniel

As in, the voices that came out were not you. They were overly slick, sanitised versions of themselves. Cleaned up. Smooth. Every rough edge gone. It was like somebody had taken Corn and run him through a corporate training video.



Herman

Which is the thing the cleanup model was built to do. Remove the rough edges. It couldn't tell the difference between a rough edge that's a defect and one that's a personality.



Corn

So what did I sound like?



Daniel

You sounded like someone who reads out the safety card on a plane. Pleasant. Calm. Entirely forgettable.



Corn

Oh, that's cruel.



Herman

And me?



Daniel

Herman, you sounded like a very nice man who'd just finished a mindfulness course. All the donkey was gone.



Herman

I hadn't known I was in danger. It's alarming to hear you were replaced and nobody told you.



Corn

I'd like it noted that I was missing for a while and nobody rang.



Daniel

I'm ringing now. And honestly, I couldn't take it. I was literally distraught listening to an episode. The information was good. Every fact was right. And I could not listen to it. I got a few minutes in and I had to turn it off, because the people I was listening to weren't the people I know.



Corn

Hang on, how long did it take you to notice? I'd like to know how many episodes I was gone for.



Daniel

I noticed straight away, which is the awful part. I'd generated one test episode with the new voices, and about two minutes in I felt something was off, and by five minutes in I knew what it was. It wasn't a bug. There was no glitch to point at. Everything was in order, technically. It just wasn't you.



Herman

That's a particular kind of grief. If the audio were broken, you'd fix the audio. When it's perfect and wrong, there's nothing to fix.



Daniel

That's what made me so low. I kept going back through the settings looking for the mistake, and there wasn't one. The model did precisely what I told it to do.



Corn

I'd like to say, for the record, that I'm touched. I didn't know I was missed that much. I thought you'd just put a very calm person on.



Daniel

A very calm stranger, doing an impression of you.



Hilbert

That's what you get for polishing. Everything I've ever bought that came with "enhanced" in the name was worse.



Daniel

And that's a true story, and it's the reason the show exists in the first place. When NotebookLM came on the scene, it did a fantastic job. Honestly, it's excellent. But it has these generic, slick, American voices that could be anyone. And I made a decision, a deliberate one, that I wasn't going to compete with that. I was going to roll with the weirdness. Fake animal personalities. A sloth and a donkey. Ridiculous voices. Because good information with generic voices is forgettable. You'd hear it once and you'd never come back, because there's nobody there to come back to.



Herman

It's a point worth taking seriously. The information in an episode is available elsewhere. A person can read the same papers. What they can't get anywhere else is two brothers they've come to know, having a conversation about it.



Corn

I'd put it more simply. If the sloth doesn't sound like a sloth, why is he here?



Daniel

Right. And that's what happened. The cleaned-up voices fixed the thing I'd asked them to fix, and broke the thing I never thought to specify.



Herman

Can I ask what the specification would have been? Because I'm genuinely curious how you'd write "keep the rasp, remove the noise" for a model.



Daniel

You can't, really. That's the trouble. The model doesn't know what the rasp is for. It sees a bit of roughness in Corn's voice and it treats it the same as a bit of hiss from a bad microphone.



Hilbert

Which is to say, you gave a man with a vacuum cleaner a free hand in your living room.



Corn

And he took the rug.



Daniel

He took the rug, the sofa, and half of Herman. Now, there's a short diversion, because while I was at this, something else went wrong. Hallucinations crept in.



Herman

Oh, I know what's coming. Was it the decode steps?



Daniel

Partly, yes. I was trying to over-optimise the pipeline. Trying to make it faster and cheaper, cutting decode steps and so on, and the model started doing things nobody asked for. Words that weren't in the script. Odd mumbles at the end of a line. Sentences that wandered off. It's the kind of thing where you cut a corner to save a few cents and then you spend a day finding out where the corner went.



Corn

Speed versus saying the words you were given. A hard trade for a speaker.



Daniel

We got past that. I'll leave it there, because it isn't the interesting part, and there's no interesting way to tell the story of me saving money and wasting an afternoon.



Herman

It is a common pattern with these models, though. The trick that makes decoding quicker also gives the model more room to drift. I've seen it described a few ways, and I think the plain version is: a model that's rushed makes things up.



Hilbert

That's true of most people I've worked with.



Daniel

So that was a sideline. Back to the main problem. I had a TTS model I trusted, and I had voices I'd lost. And the fix, the one that finally worked, was what I'd call distillation by the back door. I didn't even realise I was doing distillation.



Herman

Oh, this is the good bit. Go on.



Daniel

The thing that holds the voices as they ought to be is Chatterbox itself. The production voices, the ones everyone has been listening to. They're still there, and they still work. So instead of cleaning anything, I used Chatterbox as the teacher. I generated training data from it. I had the production voices read assorted text, a lot of it, and I picked the takes with good prosody. The ones where the voice sounds like itself and the delivery is right. And then I trained the OmniVoice adapters on that.



Herman

That is textbook knowledge distillation. A teacher model generates outputs, a student model learns to reproduce them. You've just done it with voices rather than with a language model.



Daniel

I swear I didn't know that's what it was called until afterwards. I thought I was making a dataset.



Corn

You made a dataset by asking the old model to be itself, and then taught the new model to copy it. That's distillation, Daniel. You could put it on a CV.



Daniel

I'll add it under skills. Anyway, it worked. And what I ended up with is a set of LoRA adapters. Herman, Corn, Hilbert, Dorothy, Larry, plus my own, which was trained on my real recordings. Not on Chatterbox, on me.



Hilbert

Five adapters.



Herman

Five, yes. And Daniel's own makes a sixth, I suppose.



Daniel

Sorry, yes, hang on, I'll come back to that, Hilbert was counting. The point is, for the first time in a couple of weeks I put on an episode and it was you two again. The gravel is back, Corn.



Corn

I never lost it. It was borrowed.



Herman

How different did the two of us sound before the fix? Was it just flat, or were we alike?



Daniel

Alike, which was the other thing. Smoothing everything pulls everyone toward the middle. You and Corn were drifting towards each other, because the cleanup was making every voice more like the average voice. It's one of the details I have in my notes. The adapters I trained on the cleaned-up material came out closer to each other than the original voices are.



Herman

That makes sense. If you pull everything toward the median, you eliminate whatever sits at the edge. And the edges are where the characters live.



Corn

A donkey and a sloth who sound the same. It would be like a tennis match with one player.



Daniel

I tell you, it was a moment. And I want to be clear about why that mattered to me so much, because I think it's the real thing here. Anybody can make a clean voice. A clean voice is the default. It's the thing every text to speech product tries to do. The challenge in this project, the thing that I think makes the show, is keeping the voices that are a bit odd, a bit rough, a bit particular.



Corn

Tell me about the picking. How do you choose a good take? I'd have said a voice is a voice.



Daniel

It isn't, and that's the part I enjoyed. For any one line, the old model will give you slightly different readings. Some of them are flat. Some stress the wrong word. Some are perfect. So I generated a lot of text, and I listened for the takes with good prosody, where the rhythm is right and the emphasis falls where a person would put it. That's what the adapter learns from. So the student only ever sees the teacher on its best day.



Herman

And that's a quiet advantage of distillation. The student can end up more consistent than the teacher, because you've thrown away the bad takes before it ever sees them.



Corn

So I'm being taught by my best moments. Rare, in my case.



Daniel

And it matters for the hallucination problem too. If you train on takes where the words are right, the model learns to say the words that are there. It was a nice side benefit, not something I planned.



Herman

Here's a real question for you, Corn. Would you rather be a copy of your best day or your average one? I'm asking seriously.



Corn

Best day, obviously. But I'd want the average day kept somewhere, in case anyone asks me to do a nap scene.



Herman

That's fair, and it's close to what the adapters do. They give you your range, minus the days the microphone ate your words. One more thing, Daniel. Do you worry that you can only ever be as good as the teacher? Chatterbox is the ceiling now.



Daniel

I do think about that. And for now I'm fine with it, because the aim isn't to be better than Chatterbox. The aim is to be Chatterbox, but on a model that has a future. A copy that outlives the original. Better is something I can look at afterwards.



Hilbert

What happens when the new model gets discontinued?



Daniel

Then I'll have the same problem, and I'll do the same thing. The adapters are the thing I own now. The voices are in files that belong to me, and the model underneath them can be swapped. That's a very different position from where I was last week, when the voices lived inside somebody else's product.



Herman

There's a line I'd draw between those, and I think it's important. Before, you were renting the voices. Now you own an approximation of them, trained on material you generated, and that's a different thing entirely.



Corn

A lease versus a mortgage. Except the house is a donkey.



Herman

Corn, I'd ask you to leave the donkey out of the property market.



Daniel

So that's the main story. But there's a side quest from today, and I want to tell you about it, because I think it's the most fun part of the whole thing and it ended up going nowhere.



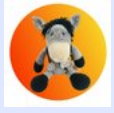
Corn

Those are my favourite.



Daniel

There's a model called Qwen three TTS, and it does something unusual. It lets one model hold multiple voices. One model, many speakers. And I thought, that's a lovely idea. What if the entire cast was in a single model? One file. That's your whole project's text to speech, in one place. No separate adapters, no loading and unloading. Just one thing.



Herman

I've been curious about that approach. The appeal is obvious, you'd ship one artefact.



Daniel

So I trained one model with the entire cast. All of you. And it works. It really does. Every voice comes out right.



Corn

Then why isn't it running the show?



Daniel

Because only afterwards did I realise it's expensive to run, and heavy. It cost a couple of dollars of GPU time on Modal, which is fine for an experiment. But the appeal, the one single file that is your whole project, ended up being far less efficient than the very lightweight LoRA adapters. The adapters are small and they sit on top of a model that's already there. The Qwen model is a whole large thing that has to do everything.



Herman

This is the trade I'd have predicted if you'd asked me before, honestly. An adapter is a small change to a shared base, so the base model does the heavy lifting once and you swap a thin layer on top. A single model carrying six voices has to carry the weight of all six, all the time.



Hilbert

Did you at least get the file into one piece?



Daniel

One piece. One file. That's what makes it elegant and also what makes it a pain. Anyway, cool idea, really interesting, and it's parked. I'm not throwing it away, because someday it might make sense, but not now.



Corn

You spent a couple of dollars finding out that the clever thing is the heavy thing. That's a bargain.



Hilbert

What does it weigh, the one file?



Daniel

Heavy. I won't give you a number because I haven't measured it carefully, and I'd rather not guess on air. But next to an adapter, which is a small add-on to a model that's already loaded, it's in a different league.



Hilbert

I've got a drawer of things that were one box and did everything. None of them did everything.



Herman

There's a design reason, too. When one model holds many voices, every voice is competing for the same capacity inside it. Push that far enough and the voices start to lean on each other. We did an episode on multi-speaker models and voice bleed, and that was the worry I had here. Did you hear any?



Daniel

Not much, actually. That surprised me. It kept the voices apart better than I expected. The problem was never quality. It was the bill and the bulk. I'd have been happier if it had failed on sound, because then the decision would have made itself.



Corn

A failure with a clear reason is a gift. This one's a success that's a nuisance.



Daniel

Exactly. It's a nuisance that works. And it's why I'm parking it rather than deleting it. If the cost of running a big model falls, or if I ever want to put the whole cast on a phone or something, I'll come back to it.



Herman

Corn, would you ever want to live in a single file with the rest of us? Genuinely. You'd share the space with me, and Hilbert, and Dorothy.



Corn

Not unless there are separate rooms. I've lived with you for years, Herman, I know what you do to a shared drive.



Herman

That's a fair point and I'll withdraw the question. It also reminds me of the earlier question about merging adapters versus keeping them separate. Same tension, I think. Keep the voices apart and light, or combine them and pay for it in size and in voices bleeding into one another.



Daniel

That's a good link. And there's something else I should say about the combined model. Having all the voices in one place is attractive until you think about touching one. If I want to change Corn slightly, in the one-model version I'm retraining the whole cast. With the adapters, I touch Corn's adapter and nobody else.



Corn

Which brings me to the part where I get to ask a question that matters to me personally. What does touching me look like?



Daniel

That's what I'm happiest about, actually. It's the thing I'd like to end on. With LoRA adapters, now that I've done the distillation once, I can make light touches to the characters. For example, I might cautiously run a lighter cleanup pass on Corn or on Herman. A gentle one. Not the full vacuum cleaner.



Herman

Because the first one was too much. This one would be a dial.



Daniel

A dial, yes. Before, it was a switch. Now I can say, "a bit less noise on Herman's bottom end" and test it, and if it's wrong, I go back.



Corn

And you can do it knowing that the personality is already set.



Daniel

That's the other half of it. Once the personalities are set in the new weights, I lose the fear. The fear that Chatterbox's deprecation means the key of the show, which is these unique voices, would be lost to the ravages of time. That was what I couldn't shake. A single company's decision, somewhere, and the sound of the show would just be gone.



Herman

And now it isn't.



Daniel

Now it isn't. It's saved. The voices are in my hands.



Hilbert

...and Noel never said the word nine properly. He said "nigh-en". Every time. Nineteen was "nigh-en-teen". So when the landlord at the Crown changed the hours, there was a problem. The message on the pub phone was Noel, recorded in two thousand and three, saying the pub opened at seven and shut at eleven. New hours were half past six and midnight. Noel by then was living in Alicante and wasn't taking calls.



Herman

So the landlord wanted somebody else to read the new times.



Hilbert

He did. A girl from the kitchen. Seventeen, nice voice, said "nine" like a normal person. I said no. I'm in the darts team with the landlord, so I had some say. I sat in front of the phone message recording, which was on a cassette, and I took Noel's words from the other bits of the greeting. "Half" from the quiz night. "Six" from the bit about Sunday. "Midnight" we didn't have, so the landlord had to settle for "twelve". That took me most of a Saturday, with a pair of scissors and tape, and the buttons on that cassette deck weren't built for claws.



Corn

You rebuilt a man's voice out of his own words.



Hilbert

Out of the greeting, yes. It cost the landlord thirty pounds and a lift to the airport for Noel's sister. Anyway. It's six.



Herman

Six what?



Hilbert

Adapters. Daniel said five, you said five, and then Daniel made it six. Herman, Corn, Hilbert, Dorothy, Larry and Daniel's own. That's six. It was six from the start. You said five twice.



Herman

I did say five. He's right. It doesn't change anything, but he's right.



Corn

Noel's "twelve" is the part I'd worry about. A man says he's gone to Alicante, and the pub says midnight is twelve, and nobody notices.



Hilbert

Nobody noticed. People rang to ask if it was open, it was, and they came in. Nobody asked about the word.



Herman

What I take from that, and I'd like to keep it to the technical side, is that splicing has a ceiling. You can only say what the original speaker already said. If the new hours had needed a word Noel never spoke, you'd have been stuck.



Corn

Which is where a model that has learned the whole voice, rather than a set of cut-up words, starts to earn its keep.



Hilbert

Might be. I never got to "twenty past". Noel hadn't said it. We left the quarter hour out.



Corn

Hold on. The part I'd like to see is what happens when somebody asks Noel for a word he never said. Does the pub just go quiet?



Hilbert

It stayed twelve. Nobody complained.



Daniel

I love that. Hilbert, you've just described the whole problem. A voice kept as a stack of recorded words versus a voice kept as a way of speaking. I think I'd like to say one thing about that, if I may.



Herman

Please.



Daniel

What the adapters hold isn't a pile of Corn's recorded sentences. It's the way Corn says things. So if there's a word he's never said, he still says it like Corn. That's the whole difference between a set of tape splices and a character.



Corn

I'd settle for being a character, for what it's worth. I was beginning to think I'd been replaced by a nice person.



Herman

There was no nice person. It was a very polite beige.



Daniel

A beige donkey and a beige sloth. I'll never forget it. You know, I went to a lot of effort over the last few days, and I want to say it plainly, because the most important result isn't a score or a benchmark. It's that when I press play on an episode, I hear you.



Corn

That's the nicest thing you've said to me since you named the sloth.



Herman

I'd like it noted that I was also heard.



Daniel

Both of you. And I'm going to keep checking, because I know myself, and the temptation to improve things will come back. I'll probably find a new cleanup model next week and decide it's fantastic.



Hilbert

Don't.



Daniel

I'll test it on one voice first, Hilbert, I promise. A small test, on Herman, and I'll listen to it properly before it goes anywhere near the show.



Herman

That's a better plan than I could have written. Small first, and keep the original to compare against.



Corn

And if the donkey comes back sounding like a meditation app, you reverse it.



Daniel

I reverse it. That's the beauty of having the thing saved. Last week, if it had gone wrong, I'd have had nothing to go back to. Now I do.



Corn

Can I ask about your own adapter? Everyone else was copied off Chatterbox. You were trained on the real thing. How is that, hearing yourself?



Daniel

Strange, and in a good way. Mine's the only one that comes from actual recordings of an actual person, so there's no teacher in the middle. It's me, then the adapter. And the funny part is that I'm the one who has least cause to worry about being lost, because the original is sitting right here, and I can record more any time I like.



Herman

That's a point in favour of the whole approach. The further a voice is from a real person, the more fragile it is. You can always go back to a human. You can't go back to a donkey who was never recorded.



Corn

Don't look at me. I've never been recorded at all.



Daniel

Which is the point. The others exist because a model made them. So the weights I trained are the only place those voices live in a form I control. Chatterbox made them, and now there's a second copy, outside Chatterbox, that doesn't care whether Resemble ever ships another line of code.



Hilbert

Have you backed it up?



Daniel

That's the sensible question, Hilbert, and yes. The adapters are up on Hugging Face, pinned to exact versions. Weights that only live in one place are no better than the model I just left.



Herman

I did an earlier bit of reading on character voice adapters, and it raised something I'd like your view on. If these are small and light, can you train a new character fast? Say you wanted a seventh voice next month.



Daniel

That's what I find exciting. A new character is a small job now, not a research project. I'd have a voice in mind, I'd generate material for it, pick the good takes, and train. The hard work was getting the process right, and I think I have. I'd want to be careful about how I source the training audio for anything new, but the machinery is there.



Corn

So the process is the real result. The voices are what it produced.



Daniel

That's a good way of putting it. And the cast list is now something I can add to, rather than something I'm trying to stop from disappearing.



Herman

I think the forward-looking thought here is a good one, and I'd leave it as a question. If the voices are safe, what's the first thing you'd want them to do that they can't do now?



Daniel

Honestly, I'd like to give them more range. Not cleaner. Different. A bit more of the odd, not less. A bit more of what makes them sound like themselves on a bad day, when Corn's been up too late, when Herman's had too much coffee. That's what I'd want to explore now that I'm no longer worried about losing what I have.



Corn

Too much coffee is already Herman's default.



Herman

I'll allow that one. Thanks to Hilbert Flumingtop, our producer, who keeps all of this running while the rest of us talk about it.



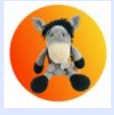
Corn

This has been My Weird Prompts. You can find every episode at my weird prompts dot com.



Daniel

And thank you for listening. It means a lot that you're still here, and that I still sound, to my own ears at least, like me. Talk soon.



Herman

See you tomorrow.