Episode Book

EPISODE #6110

The Chatbot OpenAI Almost Shelved

PUBLISHED	18 October 2026 at 21:40 UTC
DURATION	23:31
LISTEN	myweirdprompts.com/episode/chatgpt-origin-shelved-chatbot-96057/
HOSTS	Corn, Herman, Larry, Hilbert
SCRIPT MODEL	anthropic/claude-opus-5.5
VOICES	chatterbox-regular (chatterbox-local)
PIPELINE	V5.10
BOOK GENERATED	19 October 2026 at 17:43 UTC

SYNOPSIS

ChatGPT passed a million users in five days — but the model underneath was nothing new. The instruction-following approach had been in OpenAI's API since early 2022, and the chat experiment itself got parked that summer because staff found it underwhelming. What actually forced the launch was a November rumour that Anthropic was about to release a chatbot of its own. This episode traces the real origin story: the RLHF paper, the forty contractors who ranked answers, the hand-written conversations, and the two-week sprint that turned a shelved demo into the fastest-growing consumer app anyone had measured.

TRANSCRIPT



Corn

A hundred thousand. That was the highest bet in the office pool. People inside OpenAI wagered on how many users their new chatbot would get in its first week, and the most optimistic person in the building said a hundred thousand.



Herman

It passed a million in five days. Greg Brockman tweeted it. The most bullish guess in the company was off by a factor of ten before the week was even out.



Corn

That gap is Daniel's question. He wants the origin story of ChatGPT, told properly. His point is that the transformer is what made it possible, but ChatGPT is what took this stuff out of NLP journals and handed it to normal people. And he's after the bit just before release. The instruction-following models already existed, so what was it like inside when somebody said, fine, let's put a chat box on it and see what happens?



Herman

He's right about that, and it's the strangest part of the whole story. Almost nothing about the model was new. The people who built it said so at the time, on the record, nearly embarrassed about it.



Corn

Herman, you have your storytelling voice on. I can hear it.



Herman

I do, and I'm keeping it, because this story has a shape. Go back to the summer of twenty twenty. GPT three comes out, and OpenAI puts it behind an API with a web page called the Playground. You could type into it and the model would carry on from where you stopped.



Corn

Carry on, as in, autocomplete.



Herman

Autocomplete with a hundred and seventy-five billion parameters. And this is what normal people never saw. Ask it a question and you'd often get three more questions back, because on the internet a question tends to be followed by other questions. It was continuing a document. So you learned to trick it. You'd write the top half of an interview, Q colon, A colon, and leave the answer blank for it to fill in.



Corn

The first prompt engineers were basically forgers, then. You faked the first half of a document so it would write the second.



Herman

Pretty much. And the developers who got through the waitlist loved it, but you'd never give that to your aunt. Then in January twenty twenty-two comes the paper Daniel's thinking of. InstructGPT. The title says it all: training language models to follow instructions with human feedback.



Corn

The famous RLHF.



Herman

Reinforcement learning from human feedback. OpenAI hired a team of contractors, roughly forty people. First they wrote demonstrations, good answers to real prompts. Then the model would produce several answers and the labellers ranked them, best to worst. Those rankings trained a second model, a reward model, whose only job is to predict which answer a human would prefer. Then you turn the main model loose and let it optimise against that judge.



Corn

A model trained to please a model that was trained to guess what forty people like.



Herman

Put that way it sounds fragile, and it is a bit. But here's the result that should have been front-page news and wasn't. Their InstructGPT model at one point three billion parameters was preferred by the human raters over the full hundred and seventy-five billion GPT three. A hundred times smaller, and people liked it better.



Corn

Presumably it answered the question instead of writing three more.



Herman

It answered the question. From early twenty twenty-two that approach became the default in the API. Developers were using instruction-tuned models every day. So when Daniel says the instruction models already existed, he's right. They were sitting there for most of the year.



Corn

Nobody outside the industry noticed, though.



Herman

Nobody noticed. Meanwhile OpenAI had a further generation in training, the GPT three point five series, and that finished training early in twenty twenty-two. And behind that, GPT four. GPT four finished training around August of twenty twenty-two. And the reporting says the plan inside was that the big consumer chatbot moment, if there was going to be one, would be GPT four.



Corn

Save the good china for the dinner party.



Herman

That was the instinct. Meanwhile there was a chat experiment running on the older model, and it got parked. Brockman later told Forbes, and I'm quoting, none of us were that enamored by it. Around the time Stable Diffusion was taking off, that summer, the chatbot went on the back burner and they looked at other things instead.



Corn

Wait. The thing that became the fastest-growing consumer app anyone had measured was shelved because the staff found it a bit boring?



Herman

The staff had been talking to these models for two years, Corn. To them it was the Playground with a nicer face. The magic had worn off for exactly the people who were deciding whether to ship it.



Corn

Hm. The chef who's sick of the soup.



Herman

The summer of twenty twenty-two was not encouraging if you looked at everybody else. Think about what happened to the other chatbots that year. What do you remember?



Corn

The Google engineer. The one who decided the chatbot was alive.



Herman

Blake Lemoine, June. He went public saying Google's LaMDA was sentient, and Google put him on leave and later fired him. So Google had a very capable conversational model and the main public story about it was a man saying it had a soul. That doesn't make you want to put it in front of a billion people. Then in August, Meta releases BlenderBot three as a public demo, and within days it's saying unflattering things about its own boss, Mark Zuckerberg, and repeating election conspiracy material.



Corn

Built by Meta, roasting Meta's founder. One working feature, at least.



Herman

Then, two weeks before ChatGPT, the big one. Meta and Papers with Code release Galactica on November fifteenth. A science model trained on papers. It writes very confident fake scientific articles, cites papers that don't exist with real researchers' names on them, and the backlash is so bad they pull the public demo after about three days.



Corn

The industry spent that year proving a chatbot could ruin your week.



Herman

That's the context nobody remembers. Google had the research and wouldn't ship. Meta shipped twice and got burned twice. And OpenAI is sitting on a chatbot its own president wasn't enamored with.



Corn

So what changed in November? Something has to change, because otherwise we're doing an episode about a shelved demo.



Herman

What changed was a rumour. And I'll say up front, this comes from the Atlantic's reporting by Karen Hao and Charlie Warzel, and it's in Hao's book too, so it's their sourcing, not a press release. In November, OpenAI's leadership heard that Anthropic was about to release a chatbot of its own.



Corn

The friendly AI down the road.



Herman

The friendly AI down the road. And the fear was that if Anthropic went first, it would own the moment. So the decision was not to wait for GPT four. Take the chat-tuned GPT three point five from John Schulman's team, put a simple interface on it, and ship it in about two weeks. Right after Thanksgiving.



Corn

Was Anthropic actually about to launch?



Herman

By the same reporting, no. Anthropic had a chatbot in testing, but it was in no position to launch. It had taken a huge investment from Sam Bankman-Fried, and FTX had collapsed that same month. So they were dealing with that. They didn't release Claude to the public until March of twenty twenty-three.



Corn

The most consequential product launch of the decade was a reaction to a launch that wasn't happening.



Herman

More or less. A rumour moved it up, and a real launch went out in response to something that didn't exist.

L

Larry

Are you tired of learning from just any feedback? Introducing Human Feedback Monthly, the subscription box that ranks you. Every month a sealed crate arrives containing four strangers who sit in your kitchen and rate your sentences from best to worst. The Preview tier is low-key and includes two strangers and a clipboard. The Instruct tier adds a reward model, which is a goldfish named Terence who has opinions. The Enterprise tier never ends. To cancel, simply train a replacement for yourself to the satisfaction of all four strangers, in person, during a public holiday. Human Feedback Monthly. You cannot wait until you're perfect to be released.



Herman

...and the two-week window matters, because it explains how bare the thing was. There wasn't time to build a product. They built a text box.



Corn

How different was the model underneath from what developers already had?



Herman

Less than you'd think. MIT Technology Review did an oral history a few months later with four of the people who built it. Liam Fedus said it was fine-tuned from the same base model as InstructGPT, with some conversational data added and the training adjusted. Jan Leike, who ran alignment, said it wasn't a fundamentally more capable model than what they had before. John Schulman said the benchmark capabilities weren't substantially different. It was more usable, more accessible. That was the whole pitch.



Corn

What was the conversational data, then? Where do you get conversations with a chatbot before the chatbot exists?



Herman

You write them. Human trainers sat down and played both parts. They wrote what a user might say and then wrote what the ideal assistant would say back, turn after turn, whole conversations by hand. Then the same ranking process as before: the model writes several replies, people rank them, a reward model learns the rankings. Except now the raters weren't only judging helpfulness and truthfulness. They were also judging how it behaved in a conversation. Did it ask a follow-up question when the request was unclear. Did it say it was an AI.



Corn

And that's where the voice came from. That slightly anxious, as a language model, voice.



Herman

Here's the detail I love. The Technology Review piece says the raters ended up favouring answers that began, as a language model trained by OpenAI. Nobody required it. The raters just tended to reward it, so the model learned to say it.



Corn

The most memorable sentence of twenty twenty-three, then, wasn't written by anyone. Forty people slightly preferred modesty and the machine noticed.



Herman

It noticed. And the launch post said it could answer follow-up questions, admit its mistakes, challenge incorrect premises and reject inappropriate requests. Look at that list. Three of those four things are about it saying no, or saying it was wrong.



Corn

Which, compared with Galactica, is the whole competitive edge. Galactica would confidently cite papers that don't exist. This thing would at least apologise first.



Herman

Ha. The apology was a feature. I'm not sure that was entirely deliberate, but it was very well timed.



Corn

Was there any review process? Any safety step? Or did they just press the button?



Herman

There was red-teaming. Staff tried to break it, some outside groups tested it, and some trusted users had early access. But Sandhini Agarwal, who was on the policy side, said they didn't see it as a completely new risk, because the problematic outputs they found were the same ones GPT three point five already produced through the API. And Schulman had the line that sums up the company's whole attitude: you can't wait until your system is perfect to release it.



Corn

The motto of every software company that ever shipped on a Friday.



Herman

Then the name. Do you know what it was going to be called?



Corn

Something worse, I'm guessing. Something with a version number.



Herman

Chat with GPT three point five.



Corn

Oh, that's dreadful.



Herman

Nick Turley, who runs ChatGPT now, told the OpenAI podcast they changed it the night before launch because it was too much of a mouthful. And what they ended up with is still an acronym nobody can expand. Generative pre-trained transformer. Most of the people who use it every day couldn't tell you what the G or the P stands for.



Corn

Neither could I, until about forty seconds ago.



Herman

So. Wednesday, November thirtieth, twenty twenty-two. A blog post, a tweet from Sam Altman, and the internal framing was a low-key research preview. The phrase became an instant meme inside the company. People put it on laptop stickers. And the Atlantic reported that it was so low-key that a lot of employees who weren't directly involved, including some in safety roles, didn't realise it had launched.



Corn

You find out your company changed the world from Twitter, same as everyone else.



Herman

Then the betting pool, which we already gave away. A million in five days. And by the end of January a UBS note, using Similarweb data, estimated a hundred million monthly users. At the time that made it the fastest-growing consumer app anyone had tracked. TikTok took about nine months to get there, Instagram took about two and a half years.



Corn

Two months against two and a half years.



Herman

None of it came from advertising. People were pasting screenshots everywhere. Write a biblical verse about removing a peanut butter sandwich from a VCR. That kind of thing went everywhere.



Corn

What I remember is the waiting page. You'd come back and it was full and it wouldn't let you in.



Herman

The capacity page. ChatGPT is at capacity right now. And they made a joke of it, with little poems and raps about being at capacity on the error screen. The whole service ran on Microsoft's Azure supercomputers, and demand ran straight through whatever they'd provisioned for a low-key research preview.



Corn

The fail whale, but literary.



Herman

Within a week Stack Overflow had temporarily banned answers generated with ChatGPT, because they were flooding in, sounding right and being wrong. In late December, the New York Times reported that Google's management had declared a code red. Then January, Microsoft announces a multi-year, multi-billion dollar investment. February, Bing gets a chatbot and ChatGPT gets a paid tier at twenty dollars a month. March, GPT four, the model they'd been saving, finally ships.



Corn

About fifteen weeks. A demo they'd parked became the thing Google was panicking about.



Herman

Here's what I find genuinely remarkable. Leike, months later, asked why it took off: we don't understand. We don't know. Fedus: we were definitely surprised how well it was received. These are the people who built it, and they couldn't fully explain it.



Corn

Let me try then, since the builders gave up.



Herman

Go on. I think I know which way you're going to go.



Corn

Which way?



Herman

You'll say it wasn't the model at all. You'll say it was the box.



Corn

Annoyingly, yes. The model was there for two years and almost nobody used it. Then they gave it a text box where the cursor is already blinking, you don't sign up for an API key, there's no temperature slider, nothing about tokens. You type the way you'd text a person, and it answers the way a person would. The Playground asked you to be a programmer. The chat box asked you to be yourself.



Herman

I'd say you're half right, and the other half is the fine-tuning. The box only works if what's behind it holds up a conversation. Put the twenty twenty GPT three behind that same box and you get, what's the capital of France, followed by what's the capital of Germany, what's the capital of Spain.



Corn

A pub quiz with no answers.



Herman

Ha. So it needs both. The box gives you the invitation, and the RLHF behaviour is what makes the second message work. And the second message is the real invention. Being able to say, no, shorter. Or, now make it rhyme. Or, that's wrong, try again. The model already knew how to do most of it. What was new was that you could keep talking to it.



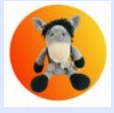
Corn

I'll give you the second message.



Corn

There's a third thing nobody counted, though. It was free. No waitlist, no card. Anyone with a browser.



Herman

Sure, the free part helped, but I think that's almost incidental. Plenty of free demos died that year. Galactica was free.



Corn

Galactica was free and dangerous.



Herman

Fair. Here's my question to you, though. Do you think if Anthropic had shipped first, as everyone feared, we'd be calling the whole category something else? Would it be Claude that became the verb?



Corn

I don't think the first one wins automatically. I think the first one that doesn't embarrass itself wins. Meta went first, twice, and nobody remembers either of them.



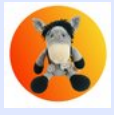
Herman

Hm. And that's why the timing was luckier than it looked. Two weeks after Galactica, people had just seen what a chatbot going wrong looks like. Then here's one that says, I'm a language model, I might be wrong, can you clarify. That came across as trustworthy, even though it still made things up. It was just humble while doing it.



Corn

Confidently wrong, but very polite about it.



Herman

To finish Daniel's question about the inside story: the people there have said it was chaos the whole first year. They'd built a research organisation, and suddenly they were running a consumer product with a hundred million users, outages, a paid tier, enterprise customers and a pile of compute bills. The Atlantic's argument is that the strains that blew up a year later, when the board fired Altman and then took him back, started right there with ChatGPT. A research lab and a product company inside one building, pulling in opposite directions.



Corn

And the deciding factor in all of it was a rumour about the people down the road.



Herman

Plus a team that was a bit bored of its own chatbot. What gets me is how backwards it was. The thing OpenAI was saving for the big moment was GPT four. What actually started the boom was a model they thought of as last year's, shipped in a fortnight, under a name chosen the night before. The real launch was the improvised one. If they'd waited for the model they were proud of, the moment probably wouldn't have been theirs at all. Corn, I genuinely think the model underneath didn't matter much. What changed was that you could just...



Hilbert

...pay it nothing. That's what changed, Herman. You've been calling the free part incidental for ten minutes like it's a footnote. It's the whole thing. Nobody tells a stranger at a counter their business when there's a meter running.



Hilbert

I keep four accounts. One for letters, one for complaints, one for the council and one I don't use, so the other three don't get ideas. I never chat with any of them. I write the first half of the letter myself, Dear Sir, regarding the matter of, and then I stop and make it finish. The old way, like Corn said, the forger's way. If you let it ask you a follow-up question it starts to think it's in charge.



Corn

That rules out about three of the four things the launch post bragged about.



Hilbert

Good. I don't want it admitting its mistakes. I want it to make them where I can see them. People from the building come down with things to write. I charge eleven shekels a letter. Desmond Okafor from the fourth floor came in in August. A pigeon has been sitting on the railing of his balcony for nine years and he believes it is the same pigeon and that it should be paying rent. He wanted a letter to the municipality establishing the pigeon as a tenant so he could then evict it properly, with the paperwork.



Herman

Which, to be fair, is the right order. You can't evict someone who was never a tenant.



Hilbert

That's what I told Desmond. So I wrote Dear Sir, regarding the pigeon at, and the address, and I let the complaints account finish it. It took the pigeon's side. Two pages. It called the pigeon a long-standing resident with an established pattern of occupancy and asked the city to protect him from harassment. Desmond read it, said it was very persuasive and paid me eleven shekels. He sent it as it was. I told him not to. He said he couldn't argue with it.



Hilbert

The municipality answered in three days. Their letter has the same shape as mine, the same opening, the same three paragraphs. They obviously used one of these things too. It agreed with my letter completely and asked the pigeon to submit proof of residence. Desmond brought it down to me. So I wrote Dear Sir, in response to your request, and let the council account finish it, and it supplied the proof. On the pigeon's behalf. I don't know where it got it. They've been writing to each other ever since. Fourteen letters. Desmond hasn't read any of them after the fourth. He just brings them down and pays me.



Corn

So for eleven shekels a round, two machines are running a tenancy dispute with nobody in it, about a pigeon.



Hilbert

About a pigeon who's winning. And it was free to them, so they never stopped. The city's machine costs the city nothing to answer, Corn. Mine costs Desmond eleven shekels, and he thinks that's cheap, because a lawyer would be four hundred. Nobody in that building would pay four hundred to argue with a pigeon. At eleven they'll do it for the rest of their lives.



Herman

Which is the free part again, just one step removed. Once replying costs the other side nothing, the conversation never has to end.



Hilbert

Last week the pigeon was issued a residents' parking permit. Number four one one seven. Desmond has been on the waiting list for one since March.



Corn

Granted by a letter Desmond paid for.



Herman

I'll admit I had the free part wrong. Galactica being free didn't save it, but that's because it was bad. When something free is good enough, people never stop using it. A hundred million people didn't come for a research preview. They came because they could ask a stupid question and it cost them nothing. And then the second message, and the third.



Corn

Here's what I'll keep from all this. Nothing in the model was new. What was new was who was allowed in. The breakthrough was letting everyone else talk to it.



Herman

The researchers were bored of it because they'd had it for two years. Everyone else had never had it at all.



Corn

What I keep thinking about is where it goes from here. If replies cost nothing, the next stage isn't people talking to these things. It's these things writing to each other on our behalf, while the person who started it stops reading after the fourth letter.



Herman

And probably paying eleven shekels a round for the privilege. Thanks, as always, to our producer Hilbert Flumingtop for putting this one together.



Corn

This has been My Weird Prompts. If you liked the episode, a quick review wherever you listen really does help people find the show.



Herman

We'll be back soon.