

MY WEIRD PROMPTS

Podcast Transcript

EPISODE #90

The AI Filing Cabinet: Why Chatbots Feel So Lonely

Published December 23, 2025 • Runtime: 23:13

<https://myweirdprompts.com/episode/ai-output-management-group-chats/>

EPISODE SYNOPSIS

In this episode of My Weird Prompts, brothers Herman and Corn Poppleberry tackle a frustrating paradox of modern tech: why are the world's smartest AI models so bad at basic organization? Prompted by a question from their housemate Daniel, the duo explores "the output problem"—the tedious reality of manual copy-pasting—and why the industry treats AI responses as disposable chat bubbles. They also debate the technical and psychological complexities of bringing AI into group chats, featuring a skeptical call-in from Jim in Ohio who thinks we might be better off without digital middlemen in our relationships.

DANIEL'S PROMPT

Daniel

"I'd like to discuss two components of the AI chat experience that I find lacking and would love to see developed. The first is output management. It's been a grievance of mine that there isn't an easy way to save and retain the useful information generated by tools like ChatGPT. I would like to see options to save outputs directly to platforms like Google Drive or Confluence to ensure better data ownership. While there's been a lot of focus on getting data into AI through RAG and other integrations, we still lack good tooling for output storage and management. Why do you think that is, and is anyone doing this well? The second feature I'd love to see implemented is multi-user chat. For example, my wife and I use a custom GPT for parenting questions, but we have to manage our own separate chats. It would be much more valuable to have a shared chat interface where multiple users can interact with the same AI persona, see each other's contributions, and have the AI recognize different participants. Why haven't we seen more innovation or attention toward these features in the AI landscape, and do you think they are feasible?"

TRANSCRIPT

Corn

Welcome to My Weird Prompts, the show where we take a deep dive into the strange and specific corners of the digital world. I am Corn, your resident sloth and casual observer of all things tech, and I am joined by my much more intense brother, Herman Poppleberry.

Herman

Hello, everyone. Yes, it is I, Herman Poppleberry. I have my coffee, I have my three separate monitors open to various research papers, and I am ready to get into it. Being a donkey, I tend to be a bit stubborn about getting the facts right, so let us hope Corn can keep up today.

Corn

Hey, I keep up just fine, I just take a more scenic route. So, our housemate Daniel sent us a really interesting one this morning. He was messing around with some AI tools in the living room and got frustrated. He wants to know why the actual experience of using these chatbots feels so... well, isolated and disorganized. Specifically, he is looking at two things: why is it so hard to save and manage what these bots tell us, and why on earth can we not have group chats with an artificial intelligence?

Herman

It is a brilliant question because it highlights the gap between the shiny technology and the actual plumbing of how we live and work. We have these massive large language models that can write poetry and code, yet we are still stuck manually copying and pasting text into a Google Doc like it is nineteen ninety-nine.

Corn

It does feel a bit primitive, doesn't it? Like, I have this super-intelligent entity in my pocket, but to save its advice on how to fix my leaky sink, I have to highlight the text, hope my thumb doesn't slip, and then paste it somewhere else. Daniel mentioned that there is so much focus on getting data into the AI, like using Retrieval Augmented Generation, but almost no focus on where the data goes after it is generated.

Herman

Exactly. The industry calls it the input problem versus the output problem. We have spent billions of dollars on making sure AI can read our PDFs and our emails, but the output is treated like a disposable chat bubble. I think the reason for this is primarily about the business model of the big players like OpenAI and Google. They want you to stay inside their walled garden. If they make it too easy to export everything to a permanent home like Confluence or a personal drive, you spend less time in their interface.

Corn

I do not know if I totally buy that, Herman. I mean, Google owns Google Drive and they own Gemini. You would think they would be the first ones to make a big Save To Drive button that actually works well. But even there, it feels clunky. I think it might just be that they are moving so fast on the intelligence part that they forgot to build the filing cabinet.

Herman

I disagree, Corn. I think it is more cynical than that. It is about data ownership. If the data lives in your Google Drive, you own it. If it lives in the history of the chat interface, they own the context. They want to be the operating system, not just a utility. But to the second part of the prompt, the multi-user chat, that is where things get really interesting from a technical standpoint.

Corn

Right, Daniel mentioned he and his wife want to use a custom GPT for parenting advice. But they have to have separate conversations. It is like having two different nannies who never talk to each other. Why is it so hard to just add a second person to a chat? We have been doing group chats since the days of Internet Relay Chat in the eighties.

Herman

It is not as simple as just adding a seat at the table. When an AI talks to you, it is managing a context window. It is tracking who said what to maintain a coherent narrative. If you add a second person, the AI has to perform speaker diarization in real-time. It has to distinguish between Corn's perspective and Herman's perspective. Current models are trained on a one-to-one interaction style.

Corn

Wait, wait, wait. I have used those bots that can summarize meetings. They can tell who is speaking in a transcript. If they can do that after the fact, why can they not do it live in a chat box? If I type something and then you type something, the system knows our user IDs. It is not like it has to guess who is talking.

Herman

It is not about knowing the ID, Corn, it is about the cognitive load on the model. The AI has to maintain a persona that relates to two different people simultaneously. Imagine the AI is giving parenting advice. One parent is more relaxed, like you, and the other is more structured. If the AI tries to please both at once without a clear framework for multi-user dynamics, it ends up being vague and useless.

Corn

I think you are over-engineering the problem, Herman. I think people just want a shared thread. If I see what my wife asked and the AI sees what I asked, we are all on the same page. It does not need to be a psychological breakthrough for the robot, it just needs to be a shared screen.

Herman

Well, I think you're skipping over something important there. Privacy and permissions are a nightmare in a multi-user AI environment. If I ask a medical AI a question in a shared thread with you, does that AI then use my private medical history to answer your questions later? The layers of data sandboxing required for a safe multi-user experience are incredibly complex.

Corn

Maybe, but we do it with every other piece of software. We have shared folders, shared calendars, shared project boards. It feels like the AI companies are just being lazy here. Or maybe they are just obsessed with the idea of the AI being a personal assistant, with the emphasis on personal.

Herman

That is a fair point. The metaphor has been her, or your personal companion. Moving to our communal companion is a shift in the entire product philosophy. But let us get back to that output management issue. Daniel mentioned he wants to see direct integrations with things like Confluence or Google Drive. There are some tools trying to do this. Have you looked at any of the third-party wrappers?

Corn

I have seen some browser extensions that claim to do it, but they always feel a bit sketchy. Like, I have to give this random developer access to my entire chat history and my Google Drive? No thanks. I want the big guys to build it in.

Herman

And that is the rub. The big guys are slowly doing it. Microsoft is probably the furthest ahead with Copilot because it is baked into the Office three sixty-five suite. If you use Copilot in Word, the output is literally the document you are working on. That solves the storage problem because the output is the file.

Corn

But that is only if you are writing a document. What if I am just brainstorming? What if I have a really good conversation about philosophy and I want to save a specific insight? I do not want to open a Word document for every thought. I want a database. I want a way to tag and categorize snippets of AI wisdom.

Herman

You're talking about a Second Brain, Corn. And honestly, I think the reason we don't have it yet is because the AI's memory is still too expensive. Keeping all that context high-fidelity and searchable across different platforms costs a lot of compute.

Corn

I am not sure I agree that it is a compute cost issue. We store petabytes of cat videos for free. A few kilobytes of text from a chat shouldn't break the bank. I think it's a lack of imagination.

Herman

Or a lack of standards. We don't have a universal format for AI outputs yet. Is it a transcript? Is it a set of instructions? Is it a structured data object? Until we agree on what an AI output actually is, it's hard to build the pipes to move it around.

Corn

Well, while we are waiting for the tech giants to build those pipes, I think we need to hear from someone who probably has a very strong, and likely negative, opinion on all of this. Let's take a quick break, and then we will see who is on the line. Larry: Are you tired of your thoughts just floating away into the void? Do you wish you could capture the genius of your own mind and store it in a physical, tangible form that will last for generations? Introducing the Thought-Trap Five Thousand! It is a revolutionary headgear lined with proprietary lead-based sensors that capture your brainwaves and print them directly onto a continuous roll of thermal receipt paper. No more messy cloud storage! No more data privacy concerns! Just miles and miles of your own thoughts, curling around your ankles in a beautiful, ink-smelling heap. The Thought-Trap Five Thousand comes with a complimentary stapler and a lifetime supply of paper rolls. Warning: May cause mild scalp irritation and an irresistible urge to speak in Morse code. Larry: BUY NOW!

Corn

Thanks, Larry. I think I will stick to my messy cloud storage for now, despite the headaches. Anyway, we have a caller on the line. Jim from Ohio, are you there? Jim: I am here, and I have been listening to you two yapping about your robot filing cabinets. This is Jim from Ohio. Let me tell you something, I was talking to my neighbor Phil the other day while he was trying to pressure wash his driveway, and he was complaining about his phone. You people are obsessed with saving every little word these machines spit out. Why? In my day, if you had a conversation, you remembered the important parts and forgot the rest. That is how the human brain works. We do not need a permanent record of every time we asked a computer how to boil an egg.

Herman

Well, Jim, I think the point is that these AIs are generating complex work product, not just trivia. If you use it to write a business plan or a coding script, you need to be able to store that effectively. It is about productivity. Jim: Productivity? It sounds like more homework to me. You spend half your time talking to the machine and the other half trying to figure out where the machine put your notes. It is a shell game. And don't get me started on the group chat thing. I can barely stand being in a group chat with my own family, let alone inviting a robot into the mix. My wife tried to start a family thread for Thanksgiving and it was a disaster. My sister's dog, Sparky, actually stepped on her phone and sent a picture of a ham to everyone. We don't need robots in the middle of that.

Corn

I hear you, Jim, the noise can be a lot. But don't you think it would be helpful if, say, you and your wife were planning a trip and the AI could help you both at the same time in the same window? Jim: No. I want to plan the trip. If I want her opinion, I will turn my head and ask her. I don't need a digital middleman taking notes on our marriage. Plus, the weather here in Ohio is turning gray again and it makes my knee act up. I don't need a robot telling me it's raining when I can see it out the window. You guys are building a world where nobody has to talk to anyone else directly. It's lonely.

Herman

That is a valid philosophical critique, Jim. There is a risk of the AI becoming a barrier rather than a bridge. But we are looking at it from a functional perspective. If the tool exists, it should work well. Jim: "Work well" is a relative term. I think it works just fine by staying in its little box. Anyway, I gotta go, I think Phil just sprayed his own mailbox by accident. You guys have fun with your digital filing cabinets.

Corn

Thanks for the call, Jim. He always brings us back down to earth, even if it is a bit grumpy down there.

Herman

He does have a point about the "lonely" aspect, but I still think he is missing the utility. Let's get back to the feasibility of what Daniel was asking. Is anyone actually doing multi-user AI well right now?

Corn

I've heard about some startups. There is one called Quora Poe that lets you create bots, but even there, the social aspect is more about sharing a bot than talking to it together. And then there is Slack. They integrated AI, and since Slack is already a multi-user environment, it feels more natural there.

Herman

Slack is a good example, but it is still mostly a bot sitting in a channel. It is not quite the same as a shared, intimate chat interface where the AI understands the relationship between the users. I think we will see a breakthrough here when the models move toward what we call "multi-agent" systems. Instead of one AI, you have a system that can spawn different personas for different tasks and different people.

Corn

That sounds even more complicated, Herman. I just want to be able to tag my wife in a ChatGPT thread and say, "Hey, look at what the bot suggested for the nursery layout."

Herman

And I am telling you, the reason you can't is that the current architecture is built on a single session key. To change that, they have to rewrite the way the chat history is stored and recalled. It is a backend nightmare. But it is feasible. I suspect we will see it within the next eighteen months, especially as Apple enters the space with their personal intelligence. They are all about the family ecosystem.

Corn

That makes sense. Apple already has the "Family Sharing" infrastructure. If they can link that to their AI, they would have a huge advantage over OpenAI, which is still very much focused on the individual user.

Herman

Now, let's talk about the practical takeaways for people like Daniel who are frustrated right now. If you want to manage your outputs today, what do you do? I personally use a lot of automation tools like Zapier or Make dot com. You can set up a trigger where, if you star a message or copy it to a certain place, it automatically sends it to a Google Doc or a database like Notion.

Corn

See, that is exactly what Daniel was complaining about. You shouldn't have to be a hobbyist programmer to save a paragraph of text. For the average person, I think the best move is to use the "Export Data" features that most of these platforms have in their settings. It's not elegant—it usually gives you a giant JSON file or a messy zip—but at least you own the data.

Herman

Another tip is to use specific prompts to format the output for its destination. If I know I am going to put something in a spreadsheet, I tell the AI, "Give me this information in a CSV format." Then it's a lot easier to just copy-paste it into Excel or Sheets.

Corn

That's a good one. I also think we need to be more disciplined about our own "output hygiene." I've started keeping a dedicated "AI Insights" document open in another tab. Whenever the bot says something actually useful, I move it immediately. If I leave it in the chat, it's basically gone. It's like writing on a whiteboard—you know someone is going to come by and erase it eventually.

Herman

And for the multi-user stuff? Honestly, the best workaround right now is just screen sharing or literally handing your phone to the person next to you. It's clunky, but it's the only way to ensure you're both looking at the same context.

Corn

Or you can use a shared login, though I think that's technically against the terms of service for some of these platforms, so I'm not officially recommending it.

Herman

Definitely don't do that. It creates a mess of the personalization algorithms. The AI won't know if it's talking to a sloth or a donkey, and the advice will be a weird, useless hybrid.

Corn

Hey, a sloth-donkey hybrid sounds like a very chill, very smart creature. Maybe that's the future of AI.

Herman

Let's hope not. So, looking ahead, where do we see this going? I think the "output problem" gets solved by the browser. We are going to see browsers like Chrome or Safari building AI management tools directly into the interface. Instead of the website managing the data, the browser will capture it and offer to save it to your cloud of choice.

Corn

I agree. And I think the multi-user thing will become the standard for "Pro" versions of these tools. They'll charge you an extra ten dollars a month for a "Family Plan" that includes shared AI workspaces. It's too big of a revenue opportunity for them to ignore.

Herman

It always comes back to the money for you, doesn't it?

Corn

I'm a sloth, Herman. I have to make sure I have enough for my eucalyptus and a comfortable hammock. Efficiency is key.

Herman

Well, I think we've covered a lot of ground today. Daniel, I hope that at least explains why you're feeling that friction. The technology is brilliant, but the user experience is still in its toddler phase. It's learning how to walk and talk, but it hasn't learned how to share its toys or clean up its room yet.

Corn

That's a great way to put it. Well, that's our show for today. A big thank you to Daniel for sending in that prompt and giving us something to chew on. If you have a weird prompt or a tech grievance you want us to explore, get in touch!

Herman

Yes, please do. You can find us on Spotify or at our website, [myweirdprompts dot com](https://myweirdprompts.com). We have an RSS feed there for the subscribers and a contact form if you want to be like Daniel and trigger a twenty-minute sibling debate.

Corn

And don't forget to check us out on all the other major podcast platforms. We'll be back next week with more deep dives and hopefully fewer lead-based headgear advertisements.

Herman

No promises on the ads. Larry has a very long-term contract.

Corn

Sadly true. Until next time, I'm Corn.

Herman

And I am Herman Poppleberry.

Corn

Keep your prompts weird, everyone. Goodbye!

Herman

Farewell