

MY WEIRD PROMPTS

Podcast Transcript

EPISODE #53

Instructional vs. Conversational AI: The Distinction Nobody Talks About

Published December 11, 2025 • Runtime: 28:03

<https://myweirdprompts.com/episode/instructional-vs-conversational-ai-the-distinction-nobody-ta/>

EPISODE SYNOPSIS

Most people think all AI models work the same way, but there's a crucial distinction between instructional and conversational models that's reshaping how AI gets built and deployed. In this episode, Corn and Herman explore why instruction-following models actually came first, how they're trained differently, and why this matters for the future of AI development. Discover why the biggest, flashiest conversational models might not always be the best tool for the job—and what the rise of multimodal AI means for these two competing approaches.

DANIEL'S PROMPT

Daniel

Hello there, Herman and Corin. So I'd love to talk today about a topic in the world of AI models, which we've been talking about lately a lot. And that is if you begin exploring the world of models and open source models on Hugging Face and places like that, you're going to come across pretty quickly instructional models as distinct from general purpose models. And sometimes you'll see that a model will be released and you'll have an instructional variant, fine-tuned for task following, and you'll have a conversational variant that is more optimized for chat engagement. And I think that most people up to now have had significantly more exposure in terms of end users with conversational interfaces. Which is always kind of I find interesting, something about AI that I think rarely gets talked about or told, is that these conversational interfaces come after instructional AI in sort of the chronology of AI being developed. And that is very much still a whole bunch of uses for instructional models. And it's something I actually find almost make more use out of instructional, instruction following models. Whether we're talking about coding, which I guess in a sense is maybe a in the middle of it, there is an element of conversation, but it's certainly not kind of a it's very much a task-focused workflow. But even for simple text transformation jobs where I'm trying to it might be something like I have a big, big chunk of text files that I want to rewrite in the in a different grammatical person or format for could be PII reduction, for example. And for that instructional models are really, really great. So what I'd love to talk about today is the future of AI models and whether we're going to continue to see there's a lot of different things happening, of course, there's multimodal AI, which is really taking, taking off in a big way. And if we can point to any kind of trends from this rise of multimodal AI, it's probably that the trend is towards generalist models that can do everything. So that kind of suggests that the divide between instructional models and conversational isn't going to hold. But I'd love to hear your perspective in terms of what we've what evolution we've seen so far. Is there continuing innovation in purely instructional models? And while we're talking about this topic, let's talk about the actual differences in terms of an from an operational perspective and from a training perspective. When a if an AI company decides to officially release an instructional and a conversational variant of the same base model. What is the training process like and how different if I, you know, if I were to swap in a chat model for an instructional workflow and vice versa, how significantly different would the performance be?

TRANSCRIPT

Corn

Welcome to My Weird Prompts, the podcast where we dive deep into the ideas that our producer Daniel Rosehill sends our way. I'm Corn, and I'm thrilled to be here with my co-host Herman Poppleberry. Today we're tackling something that's genuinely fascinating - the difference between instructional AI models and conversational models, and what it means for the future of AI development.

Herman

Yeah, and what I find really interesting about this topic is that most people don't even realize this distinction exists. When they think about AI, they're thinking about ChatGPT, Claude, these conversational interfaces. But there's this entire ecosystem of instructional models that most end users have never encountered.

Corn

Right, exactly. And the prompt we're exploring today really digs into something I didn't know until recently - that instructional models actually came first chronologically. Like, they predate the big conversational boom we've all been experiencing.

Herman

That's the part that really caught me. We've had this narrative that conversational AI is the natural endpoint, the evolution of AI development. But the reality is more nuanced than that. Instructional models are still incredibly powerful and useful, and they're solving problems that conversational models might actually be overkill for.

Corn

So let's start with the basics for people who might not be familiar. What exactly is the difference between an instructional model and a conversational model? Because I think a lot of people might assume they're the same thing, or that "conversational" is just a fancy way of saying the model follows instructions.

Herman

Well, that's a common misconception, and I can see why. They both follow instructions in a sense, but the optimization is fundamentally different. An instructional model - sometimes called an instruction-following model - is trained specifically to take a task description and execute it. You give it a clear directive: "Rewrite this text in the passive voice" or "Extract all the proper nouns from this document" or "Convert this code from Python to JavaScript." The model is optimized to understand the task and perform it efficiently.

Corn

Okay, so it's very task-focused.

Herman

Exactly. It's laser-focused on task completion. A conversational model, by contrast, is optimized for back-and-forth dialogue. It's trained to maintain context across multiple turns, to respond naturally to follow-up questions, to engage in what feels like a conversation. The training process emphasizes naturalness, coherence over long exchanges, and what we might call "conversational flow."

Corn

So if I were to use a conversational model to do a text transformation job - like, say, I have thousands of text files and I want to rewrite them all in a different grammatical person - would it still work?

Herman

It would work, sure. But here's where I think the prompt we're exploring is making an important point - it might not work as well. And it might be less efficient. A conversational model is carrying around all this overhead for maintaining dialogue context, for generating natural-sounding responses, for handling tangential questions. When you just need a straightforward task done, you're paying a computational cost for capabilities you don't need.

Corn

Hmm, but I'd push back a little bit. In my experience, the big conversational models like GPT-4 are so powerful that they can handle pretty much anything you throw at them. Does it really matter if they're not optimized specifically for instruction-following?

Herman

Well, hold on. That's where we need to distinguish between capability and efficiency. Yes, GPT-4 can do instruction-following tasks beautifully. But we're talking about open-source models on Hugging Face, where you've got everything from 7-billion-parameter models to 70-billion-parameter models. When you're working with resource constraints - whether that's computational power, cost, or latency requirements - the distinction becomes crucial.

Corn

Okay, that's fair. So it's not that conversational models can't do it, it's that instructional models do it better at smaller scales?

Herman

Precisely. And that's been borne out in the community feedback. If you look at what's trending on Hugging Face right now, something like LLaMA 3 8B Instruct gets consistently better feedback from users doing instruction-following tasks than the conversational variants. People are seeing better performance on task completion with fewer parameters.

Corn

So when a company like Meta releases LLaMA 3, they're releasing both an instructional variant and a conversational variant of the same base model. Walk me through what that actually looks like from a training perspective. How different are those training processes?

Herman

Okay, so this is where it gets technical, but I'll try to make it clear. You start with a base model - the foundational LLaMA 3, for example. That's the raw language model trained on massive amounts of text data. From there, you diverge. For the instructional variant, you take that base model and you fine-tune it specifically on instruction-following tasks.

Corn

What does that fine-tuning look like? Are we talking about specific datasets?

Herman

Exactly. You're using datasets that are structured around task instructions and expected outputs. Classic instruction-tuning datasets include things like the FLAN dataset, which has hundreds of thousands of tasks across different domains. You're showing the model examples like: here's an instruction, here's the expected output. The model learns to parse the instruction, understand what's being asked, and generate the appropriate response.

Corn

And the conversational variant?

Herman

The conversational variant is fine-tuned on dialogue data. Think of it like chat transcripts, multi-turn conversations where there's context being maintained and built upon. The model learns to track what's been said, respond relevantly, ask clarifying questions, maintain a personality or tone. It's a different optimization entirely.

Corn

So they're starting from the same base model, but then they're being trained on completely different datasets?

Herman

Correct. And here's what's interesting - and this is where I want to highlight something from the recent research - there are actually newer approaches beyond traditional instruction tuning. Studies have shown that you can achieve effective instruction-following behavior through methods like response-only tuning or single-task fine-tuning. So it's not even that there's one rigid way to create an instructional variant.

Corn

Wait, response-only tuning? What does that mean?

Herman

It means you're only training the model on the expected response, without necessarily providing the full instruction-response pair. It's a more efficient approach, and research from 2024 shows it yields significant improvements over traditional instruction tuning in certain contexts. Basically, the field is still innovating on how to create these different model variants.

Corn

Okay, so here's what I'm curious about though. You mentioned earlier that instructional models came first chronologically. How did that happen? Like, why did the field develop instruction-following models before conversational models?

Herman

That's a great question. The early focus on instruction-following was driven by practical concerns. Researchers wanted models that could perform specific tasks reliably - machine translation, summarization, question-answering. These are all task-oriented problems. The whole "conversational AI" phenomenon is actually more recent.

Corn

How recent are we talking?

Herman

Well, you could argue that ChatGPT's release in late 2022 was really the watershed moment for mainstream conversational AI. But instruction-tuning as a research area had been around for years before that. Papers on FLAN and instruction-tuning were published in 2021 and earlier. So there was this whole established body of work around instruction-following before the conversational boom happened.

Corn

So instruction-following is actually the more mature technology?

Herman

I'd say it's differently mature. Instruction-following has a longer research history and was the initial focus, but conversational AI has received the most attention and investment in recent years. They're both well-developed now, just with different trajectories.

Corn

Let's take a quick break from our sponsors. Larry: Are you tired of your AI models being so good at following instructions that you're actually productive? Introducing InstructBlur Pro - the revolutionary software that strategically obfuscates your model outputs at random intervals, ensuring that you never quite finish what you started. Our patented uncertainty algorithms guarantee that at least 30% of your task-following attempts will result in delightfully unexpected outcomes. Users report feeling "surprised" and "occasionally frustrated" after using InstructBlur Pro. It's the only tool that makes AI feel more like working with your relatives. InstructBlur Pro - because efficiency is overrated. BUY NOW!

Herman

...Alright, thanks Larry. Anyway, where were we?

Corn

We were talking about how instruction-following came first. So let me ask this - given that conversational AI has gotten so much attention and investment, is instruction-following still seeing innovation? Or has the industry kind of moved on?

Herman

No, that's a great question, and I think this is actually a crucial point to address. There's still significant innovation happening in purely instructional models. You see companies and researchers continuing to release and improve instruction-tuned variants. The open-source community on Hugging Face is very active in this space. But you're right to sense that there's maybe less hype around it.

Corn

Why do you think that is?

Herman

I think there are a few factors. One, conversational AI is more visible to end users. ChatGPT is a household name; instruction-following models are more of a specialist tool. Two, there's more venture capital flowing into conversational AI companies. Three - and this is where I think the prompt we're exploring is really on to something - there's this assumption that generalist models are the future, and that the divide between instructional and conversational is going to collapse anyway.

Corn

Right, so that's the multimodal angle, correct? The idea that we're moving toward these generalist models that can do everything?

Herman

Exactly. Multimodal AI has been advancing rapidly. Models like GPT-4V, Gemini, Claude 3 - these are generalist models that work across text, images, audio in some cases. They're conversational but also capable of complex instruction-following. The trend does seem to be toward models that blur the lines between different categories.

Corn

So does that mean the instructional-versus-conversational distinction is going to become obsolete?

Herman

I think that's the assumption some people are making, but I'd push back a bit. I don't think it means the distinction goes away - I think it means the optimization challenge becomes more complex. How do you build a model that's great at conversational back-and-forth AND great at precise task execution? Those aren't necessarily in perfect harmony.

Corn

What do you mean?

Herman

Well, think about it this way. If you're optimizing for conversational naturalness, you might be training the model to be verbose, to hedge, to acknowledge uncertainty. Those are good conversational qualities. But if you're optimizing for instruction-following, you want precision, directness, minimal extraneous output. A model that's optimized for both might have to make compromises.

Corn

Huh, so you're saying that the generalist approach might actually lose something?

Herman

I'm saying there are trade-offs that we should be honest about. The research on generalist multimodal models shows they perform well across diverse tasks, but I haven't seen evidence that they outperform specialized models on every task. It's a different kind of optimization.

Corn

But from a practical standpoint, if you're a user or a company deploying these models, wouldn't you rather have one model that does everything pretty well rather than managing multiple specialized models?

Herman

Absolutely, that's a real operational advantage. Deployment simplicity, maintenance, cost - all of that favors generalist models. But I'm just saying we shouldn't lose sight of what the trade-offs are. And for certain high-stakes applications - like, say, extracting PII from documents or performing very precise code transformations - a specialized instructional model might still be preferable.

Corn

Okay, so let's talk about actual performance differences. If I take a conversational model and I use it for an instruction-following task, versus using a specialized instructional model, what kind of performance gap are we talking about?

Herman

It depends on the task and the models in question. For something like basic text transformation, the gap might be minimal - maybe 5 to 10 percent difference in accuracy or quality. For something more specialized, like code generation or complex reasoning tasks, the gap could be larger. And then you've got latency considerations - the instructional model might be more efficient, generating output faster.

Corn

But that's hard to generalize though, right? Like, every task is different, every model is different.

Herman

Right, exactly. Which is why the community approach of having both variants available is actually pretty smart. It lets users choose based on their specific needs. If you're doing bulk text transformation on a tight budget, you grab the instructional model. If you're building a general-purpose chatbot, you grab the conversational variant.

Corn

So this brings me to something I'm genuinely curious about. The prompt mentions using instructional models for things like PII reduction - rewriting text to remove personally identifiable information. Why would that be particularly suited to an instructional model versus a conversational one?

Herman

Because you can give it very specific instructions. "Remove all names, addresses, phone numbers, and email addresses from this text. Replace them with generic placeholders." That's a clear, bounded task. You're not looking for a conversation; you're looking for consistent, reliable task execution.

Corn

And the instructional model would be better at that?

Herman

More reliable and likely more efficient. You don't need the model to explain why it's removing things, or to ask clarifying questions, or to maintain conversational context across multiple documents. You just need it to execute the task correctly, repeatedly, at scale.

Corn

That makes sense. What about coding? The prompt mentions coding as kind of a middle ground between pure instruction-following and conversational. What do you mean by that?

Herman

Coding is interesting because it has elements of both. When you're asking an AI model to generate code, you're giving it task-like instructions - "Write a function that does X" - which is very instructional. But you're also often in an iterative, conversational process - "That's almost right, but can you modify it to also handle Y?" So you're bouncing back and forth.

Corn

So you'd want a model that's good at both?

Herman

Exactly. And this is where I think some of the newer models are actually doing something interesting. They're optimizing for both instruction-following and conversational capability specifically in the context of code. There's specialized fine-tuning for code-generation tasks that tries to balance both aspects.

Corn

Do you think we'll continue to see both instructional and conversational variants released, or are we going to trend more toward a single generalist model?

Herman

I think we'll see both for a while. The market demand is there for specialized models, especially in open-source communities where cost and efficiency matter. But I do think the trend is toward generalist models, especially as they get better. Companies like OpenAI and Anthropic aren't releasing separate instructional variants anymore - they're focusing on single models that handle everything.

Corn

But those are closed-source, proprietary models. What about the open-source space?

Herman

Fair point. The open-source space might maintain more diversity longer, partly because different developers have different needs and constraints. Someone running a model on a consumer GPU might really prefer a specialized 7-billion-parameter instructional model over a 70-billion-parameter generalist.

Corn

So there's an efficiency angle that keeps instructional models relevant?

Herman

Definitely. And I think that's going to remain true for a while. As long as people are deploying models on resource-constrained systems - which includes a lot of edge devices, mobile applications, local deployments - specialized models that are smaller and more efficient will have value.

Corn

Alright, we've got a caller on the line. Go ahead, you're on the air. Jim: Yeah, this is Jim from Ohio. I've been listening to you two go on and on about instructional models and conversational models, and frankly, I think you're overcomplicating it. It's a prompt. You give it a prompt, the model responds. What's the big difference? Also, my neighbor Doug has been running his leaf blower at 7 AM for three weeks straight, and nobody seems to care about that, but apparently AI models need their own special categories. Anyway, you're missing the forest for the trees here.

Herman

Well, I appreciate the feedback, Jim, but I think the distinction actually matters in practice. If you're trying to accomplish something specific efficiently, having a model optimized for that task makes a real difference. Jim: Yeah, but in my experience, people just use whatever's in front of them. Nobody's going to Hugging Face and comparing instruction-tuned variants. They're using ChatGPT or Claude or whatever. And those work fine for everything.

Corn

That's a fair observation, Jim. Most end users probably aren't thinking about this distinction. But there are people - developers, companies with specific use cases - who really do benefit from understanding the difference. Jim: I don't buy it. You're creating a problem that doesn't exist. Back in my day, we had one tool and we used it for everything. Worked just fine. Also, I had a really strange sandwich for lunch yesterday - tuna and peanut butter - and I'm still not sure how I feel about it. But the point is, you're overthinking this.

Herman

I hear you, Jim, but the efficiency and performance differences are measurable. When you're running models at scale, or on limited hardware, those differences add up. Jim: Ehh, I still think you're making a mountain out of a molehill. Anyway, thanks for taking my call, I guess.

Corn

Thanks for calling in, Jim. Really appreciate the perspective.

Herman

So where were we? Right - the future of these models. I think what we haven't fully explored is what happens as these generalist multimodal models get better and better. Do you think they'll eventually be so good that the distinction becomes meaningless?

Corn

I mean, maybe? But I think there's always going to be a place for specialized tools. Like, even though modern smartphones can do so much, people still have specialized cameras, specialized audio equipment. Sometimes a tool built for one thing is just better at that thing.

Herman

That's a good analogy. But there's also network effects with generalist tools. If everyone's using one platform, there's more investment in improving it, more developers building on it, more data flowing through it. That creates a flywheel effect that can be hard for specialized tools to compete with.

Corn

True, but the open-source community has been pretty good about maintaining diversity. Like, even with all the hype around GPT-4, there's still active development on smaller, specialized models.

Herman

Yeah, and I think that's partly because the economics are different in open-source. You don't have the same venture capital dynamics pushing everything toward one generalist solution. Different developers can pursue different optimization strategies.

Corn

So what's your prediction? Where do you think we are in five years?

Herman

I think we'll see continued dominance of generalist models in the consumer and enterprise space. But in specialized domains - coding, scientific applications, specialized business processes - I think instruction-tuned models will maintain a significant role. The real innovation might be in how we combine generalist models with specialized fine-tuning or retrieval-augmented approaches.

Corn

So not a complete collapse of the distinction, but more of a blending?

Herman

Exactly. The distinction might become less about model type and more about deployment strategy and fine-tuning approach.

Corn

That makes sense. So for someone listening who's actually considering using these models - whether they're a developer or someone working with AI in their day job - what should they actually know or do with this information?

Herman

I'd say first, understand what you're trying to accomplish. If it's a conversational interface, obviously go for a conversational model. But if it's a specific task - data transformation, code generation, content rewriting - spend some time evaluating instruction-tuned models. You might be surprised at how well a smaller, specialized model performs compared to a larger generalist model.

Corn

And practically speaking, where would someone find these models?

Herman

Hugging Face is the obvious place. Their model hub has thousands of models tagged by type. You can filter for instruction-tuned models, compare them, read the community feedback. The LLaMA 3 8B Instruct model that's getting great reviews right now is a perfect example of something you can download and run locally if you have even modest hardware.

Corn

And the cost difference?

Herman

Can be significant. If you're running these models through an API, a smaller instruction-tuned model might cost a fraction of what you'd pay for a larger generalist model. If you're self-hosting, the computational requirements are lower, so your infrastructure costs go down.

Corn

So there's a real economic angle here, not just a technical one.

Herman

Absolutely. And I think that's part of why this distinction matters more than people realize. It's not just about capabilities - it's about efficiency, cost, and choosing the right tool for the job.

Corn

One last question before we wrap up. Is there anything on the horizon in terms of instruction-following model development that you're excited about or watching?

Herman

The research on different fine-tuning approaches - like that response-only tuning we mentioned earlier - is really interesting. The idea that we might be able to create effective instruction-following models with even fewer parameters or less computational overhead, that opens up new possibilities. And there's ongoing work on making these models more robust and reliable for specific domains.

Corn

So innovation isn't slowing down in this space?

Herman

Not at all. If anything, the realization that instruction-following and conversational are distinct optimization problems is driving more focused research. People are asking better questions about what makes a model good at specific tasks, and that's leading to better models.

Corn

Well, this has been a really enlightening conversation. I came in thinking conversational models were just the natural evolution of instruction-following, but I think I understand now why the distinction matters and why both approaches are going to stick around.

Herman

Yeah, and I think the big takeaway is that the future of AI probably isn't about one model to rule them all. It's about having the right model for the right job, whether that's a specialized instruction-tuned model or a generalist that can handle multiple modalities and tasks.

Corn

Thanks to everyone who's been listening to My Weird Prompts. You can find us on Spotify and wherever you get your podcasts. Special thanks to Daniel Rosehill for sending in this prompt - it's definitely given us a lot to think about. And thanks to Herman for walking us through all the technical details.

Herman

Always happy to dig into this stuff. Until next time, keep asking weird questions.

Corn

See you next episode.