

## MY WEIRD PROMPTS

Podcast Transcript

### EPISODE #28

# Your AI, Evolving: Beyond the Static Snapshot

Published December 07, 2025 • Runtime: 25:43

<https://myweirdprompts.com/episode/self-training-ai-models/>

## EPISODE SYNOPSIS

This week on "My Weird Prompts," Corn and Herman tackle Daniel Rosehill's fascinating challenge: how do we make personalized AI truly evolve with its user, moving beyond a static snapshot? We dissect Daniel's experience fine-tuning a speech-to-text model for his unique voice and specialized tech jargon, highlighting both the immense power and the significant hurdles of current customization methods. The discussion reveals a core dilemma: current fine-tuned models, while precise, become quickly outdated as users' needs or knowledge domains shift, creating an "old suit" that no longer fits. We delve into Daniel's visionary concept for "auto-correcting, auto-calibrating, auto-training" AI—a system using dynamic buffers and incremental learning to adapt continuously without "catastrophic forgetting"—and explore how cutting-edge research in continual learning aims to bring this truly adaptive, living AI closer to reality.

# TRANSCRIPT

## Corn

Welcome to "My Weird Prompts," the podcast where Daniel Rosehill sends us the most fascinating, mind-bending, and often deeply technical questions, and Herman and I try to make sense of them. I'm Corn, your endlessly curious co-host.

## Herman

And I'm Herman, here to add a touch of technical precision to Corn's enthusiasm. This week, Daniel has given us a prompt that really cuts to the core of what personalized AI means, and the significant challenge of keeping it, well, \*personal\* over time. He's asking about something quite profound: how we can move AI from a static snapshot to a living, breathing entity that evolves with its user.

## Corn

That's right, Herman. Daniel's prompt this week is all about how we can alter standard AI tools like large language models and speech-to-text models to tailor them for our specific needs, and the various ways we can do that. He recently shared his own experience fine-tuning OpenAI's Whisper model – a speech-to-text AI – with about an hour of his own voice data.

## Herman

It's a fantastic real-world example, Corn. Daniel's objective was two-fold: first, to see if direct exposure to his voice would improve accuracy for his unique vocal patterns, and second, to see if it could correctly transcribe niche vocabulary he uses frequently, like "containerization" or "Kubernetes."

## Corn

Oh, I know those words well from Daniel's prompts! They're definitely not everyday conversation. So, he wanted the AI to learn his specific tech jargon. And it sounds like his experiment worked, right? He saw encouraging examples that it could actually perform as intended.

### Herman

Indeed. The fine-tuned Whisper model was able to better understand his voice and correctly transcribe those specialized terms. This highlights the immense power of fine-tuning: taking a general-purpose model and customizing it for a very specific use case, significantly improving performance and relevance for that particular domain or individual.

### Corn

But, and this is where it gets really interesting, Daniel pointed out that while the results were useful and powerful, the fine-tuning process itself was far from easy. He mentioned it took him like a year to learn how to do it properly. That sounds like a significant barrier to entry for most people.

### Herman

It absolutely is, Corn. Daniel described having to prepare his dataset according to very specific formats, which often differ between models and tasks. A dataset for fine-tuning an image generation model, for instance, would be vastly different from one for a large language model or an audio recognition model. These idiosyncratic dataset requirements, coupled with environment setup and the actual training process, make it a complex and time-consuming endeavor. It's not just about collecting data; it's about meticulously preparing it.

### Corn

Okay, so a lot of effort upfront for a powerful, tailored AI. But then Daniel introduces the core dilemma. He poses a hypothetical: what if he changes careers? What if he stops talking about "Kubernetes" and "containerization" and becomes a doctor, needing the AI to understand "pneumothorax" or "tachycardia"? He can't just "erase" his old tech data or easily update the existing fine-tune. He'd essentially have to start over, or at least go through a similar, painstaking fine-tuning process for the new domain.

### Herman

This is precisely the "static snapshot" problem Daniel is highlighting. When you fine-tune a model, you're essentially freezing its learned parameters at a specific point in time, based on the data you provided. It creates a highly optimized version for that particular context. But human beings, as Daniel aptly put it, are not static. Our vocabularies evolve, our preferences shift, our knowledge expands. A fine-tuned model, after a year or two, might become less and less relevant because its internal representation of "you" or "your domain" hasn't kept pace.

### Corn

So, it's like buying a custom-tailored suit, but then your body shape changes. The suit still fits the \*old\* you, but not the \*new\* you. And you can't just un-tailor it or re-tailor it with a simple button push.

### Herman

A perfect analogy, Corn. And Daniel also touched on existing, simpler solutions like vocabulary dictionaries. These are lists of words you can add to a speech-to-text system, for instance, to improve recognition of specific terms. But he correctly noted these are often "program-specific" and act more like "rewriting rules" rather than fundamentally changing the model's underlying intelligence.

### Corn

Yes, he said they don't fundamentally change the model's ability. It's like teaching the model to recognize "Kubernetes" as a specific string of characters, but it doesn't necessarily integrate that into a deeper understanding of the concept, or how it relates to other tech terms. It's just a surface-level addition.

### Herman

Precisely. A dictionary provides a lookup table, a rule-based override. It doesn't modify the millions or billions of parameters within the neural network that represent the model's understanding of language, context, and nuance. True fine-tuning alters these fundamental parameters, allowing the model to genuinely "learn" new patterns, styles, or vocabularies and generalize them across different tasks or user interfaces.

### Corn

Okay, so that brings us to Daniel's "ideal circumstance." He envisions a model that's "self-correcting and self-updating." He talks about a "buffer" that records updates to the model – new words, updated user data – and then triggers automatic, periodic, incremental fine-tuning. He calls it an "auto-correcting, auto-calibrating, auto-training model" that happens incrementally. That sounds incredibly advanced, Herman. Does such a thing even exist, or is it pure science fiction at this point?

**Herman**

It's a vision that aligns very closely with some of the most cutting-edge research in AI, Corn, particularly in the fields of continual learning, online learning, and adaptive AI systems. While a fully realized, off-the-shelf "auto-correcting, auto-calibrating, auto-training" model isn't ubiquitous yet, the components and theoretical underpinnings are very much active areas of development.

**Corn**

So, what exactly are those components? How would this "buffer" system work in practice?

**Herman**

Let's break it down. Daniel's "buffer" concept is essentially a dynamic memory or experience replay mechanism. In traditional machine learning, models are trained on a fixed dataset in a batch process. Once trained, they're deployed, and their knowledge is static. Daniel's buffer suggests an ongoing feedback loop.

**Corn**

So, as I use the AI, it's constantly collecting new information about me?

**Herman**

Exactly. Imagine your interaction with a personalized AI. When you introduce a new word it doesn't recognize, or provide explicit feedback like "I really didn't like that movie," or even implicit signals like how long you dwell on a certain piece of content, that information is temporarily stored in this "buffer."

**Corn**

And then what happens? Does it just sit there?

### Herman

No, that's where the "auto-training" comes in. Periodically, or perhaps based on a threshold of new data collected, this buffer would trigger a micro-fine-tuning event. Instead of retraining the entire model from scratch, which is computationally expensive and risks "catastrophic forgetting"—where the model forgets previously learned information when learning new things—these systems employ techniques designed for incremental learning.

### Corn

"Catastrophic forgetting" sounds pretty bad. So, it learns new things but forgets old things? That wouldn't be very personal, would it?

### Herman

No, it wouldn't. It's a major challenge in continual learning. Researchers are exploring various strategies to mitigate this. One approach is **Elastic Weight Consolidation (EWC)**, where the model identifies which of its parameters are most crucial for previously learned tasks and "protects" them, allowing other, less critical parameters to adapt to new information. Another is **Replay-based learning**, where the buffer not only stores new data but also a small, representative sample of old data, which is then occasionally "replayed" alongside new data during updates to reinforce prior knowledge.

### Corn

Okay, so it's not just learning new things, but trying to retain the old things it knew about me. That makes a lot more sense for a truly personalized AI. And how does this relate to the "self-correcting" aspect Daniel mentioned?

### Herman

The self-correcting aspect ties into what we call **Reinforcement Learning from Human Feedback (RLHF)**, but applied continuously and at a micro-level. When you mark a transcription as incorrect, or explicitly state a preference, that's a direct signal. But often, it's implicit. For a recommendation system, if you consistently skip certain types of content, that's a signal. The "self-correcting" model would infer these preferences and adjust its internal weights without explicit intervention, ideally leading to better recommendations or more accurate transcriptions over time.

## Corn

So, it's learning from my ongoing behavior, not just a one-time dataset. That sounds like a much more intelligent, adaptive assistant. Are there any actual real-world implementations or research areas pushing this forward?

## Herman

Absolutely. While a fully autonomous, production-ready system embodying all of Daniel's vision is still evolving, many AI systems are incorporating elements of this. Consider personalized **recommendation engines**, for example. Advanced ones don't just rely on your initial preferences. They continuously update your profile based on new items you view, add to a wishlist, purchase, or explicitly rate. This continuous feedback loop, often without full fine-tuning of the base model, creates a dynamic profile. In **conversational AI**, personal assistants are getting better at remembering context and user preferences across sessions. This isn't full model retraining but rather an intelligent use of memory layers and dynamic knowledge graphs that are continuously updated. So, if you tell your assistant your favorite coffee order, it remembers it for next time.

## Corn

But that sounds more like a sophisticated memory than true learning, right? Daniel seemed to be talking about changing the actual underlying model.

## Herman

You're perceptive, Corn. And you're right, simply remembering facts is different from fundamentally altering the model's perception or generation capabilities. The more advanced implementations are looking at truly adaptive learning. For instance, **Federated Learning** is a privacy-preserving technique where models are trained on decentralized user data—like on your phone—and only the learned updates are sent back to a central server, which then aggregates these updates to improve the global model. This allows for continuous, incremental learning without centralizing sensitive user data. Each user's device could essentially host a "micro-fine-tuned" version of the model that's periodically updated locally and then subtly contributes to the global model.

## Corn

So, my phone's AI could be constantly learning my specific speech patterns or my unique preferences without sending all my private data to a central cloud, and that learning helps me directly and maybe even improves the general model a little bit too? That's really smart.

### Herman

Exactly. Another area is **Meta-learning**, or "learning to learn." These models are designed to rapidly adapt to new tasks or data with very few examples. So, if Daniel started talking about medical terms, a meta-learned model might pick up the new vocabulary and context much faster than a traditional model, requiring less new fine-tuning data. We're also seeing modular AI architectures emerge. Instead of one monolithic model, you might have a core foundation model and then smaller, more agile "adapter modules" that are easier to fine-tune and update for specific user preferences or domains. These adapters can be swapped or updated without touching the massive base model, making the process much more efficient and less prone to catastrophic forgetting.

### Corn

That sounds like a much more elegant solution to Daniel's career change problem. Instead of retraining the whole "person" (the base model), you're just swapping out or updating a smaller "skill set" (the adapter module).

### Herman

Precisely. And this is where the concept of a "digital twin" for a user becomes relevant. Imagine an AI agent that maintains a continuously updating profile of your knowledge, preferences, and even your evolving verbal tics or content consumption habits. This digital twin would then intelligently inform how the base AI model interacts with you, ensuring it remains deeply personalized over time.

### Corn

This really changes the game, doesn't it? From a static AI that just responds to commands, to an AI that truly understands and grows with you. But what are the challenges in building something like this? It sounds incredibly complex.

## Herman

The challenges are significant. **Computational Cost:** Continuous training, even incremental, still requires substantial computational resources. Finding efficient ways to update models on the fly, perhaps even on edge devices, is crucial. **Data Management:** Managing the "buffer" of new data, deciding what to keep, what to discard, and how often to trigger updates is a complex data engineering problem. **Stability vs. Plasticity:** This is the core dilemma of continual learning. How do you ensure the model remains stable and consistent in its core knowledge while being plastic enough to adapt to new information without forgetting the old? **Privacy and Security:** As the AI learns more about individual users, the ethical implications and data privacy concerns become even more pronounced. How do we ensure that this deep personalization doesn't lead to misuse of personal data? **Explainability and Control:** If the AI is constantly updating itself, how do users understand why it behaves a certain way? And how much control do users have over what the AI learns or unlearns about them?

## Corn

Those are definitely big questions. So, for our listeners, what are the practical takeaways here? For developers building AI, and for us, the users, what can we do or expect?

## Herman

For developers, Daniel's prompt is a clarion call to move beyond "model-centric" AI design to "user-centric" AI design. This means: 1. **Prioritizing adaptive architectures:** Design AI systems from the ground up with continuous learning and personalization in mind, rather than trying to retrofit it later. 2. **Robust feedback loops:** Implement clear, intuitive ways for users to provide explicit feedback, and build intelligent systems to infer implicit feedback from user interactions. 3. **Modular and incremental updates:** Explore modular model designs and incremental learning techniques to make personalization more efficient and less resource-intensive. 4. **Embrace ethical AI design:** Put privacy, security, and user control at the forefront of adaptive AI development.

## Corn

And for us, the users? What does this mean for how we interact with AI in the future?

### Herman

For users, it means the promise of truly personalized AI that doesn't just respond to commands but genuinely anticipates and adapts to our evolving needs and preferences. We can expect AI systems that get smarter \*with\* us, reflecting our individual journeys and growth. It shifts the burden of adaptation from the user to the AI.

### Corn

So, instead of me having to constantly adjust to the AI, the AI will adjust to me. That sounds like a much more natural and powerful partnership.

### Herman

Exactly. The ultimate vision is an AI that's not just a tool, but an evolving extension of our own cognitive processes, continuously learning and adapting to our dynamic lives. It brings us closer to a future where AI systems are not just intelligent, but intimately familiar and uniquely suited to each individual.

### Corn

That's a truly mind-bending concept, Herman. From the static fine-tune that Daniel experimented with, to this idea of an AI that grows and changes alongside us. It really makes you think about the future of human-AI collaboration on a whole new level.

### Herman

It certainly does, Corn. It's about building AI that lives and learns, rather than just computes.

### Corn

Well, Daniel, thank you for sending us such a profoundly thought-provoking prompt this week. You've given us a lot to chew on.

### Herman

Indeed, Daniel. Your insights into the practicalities and frustrations of current fine-tuning methods have illuminated a critical path forward for personalized AI.

## **Corn**

And to all our listeners, thank you for joining us on "My Weird Prompts." If you want to dive deeper into more of Daniel's fascinating questions, you can find "My Weird Prompts" on Spotify and wherever else you get your podcasts.

## **Herman**

We'll be back next time with another weird prompt to explore.

## **Corn**

See you then!