

MY WEIRD PROMPTS

Podcast Transcript

EPISODE #258

The Geographic Soul of AI: Mapping the Global Data Divide

Published January 20, 2026 • Runtime: 22:56

<https://myweirdprompts.com/episode/geographic-soul-ai-models/>

EPISODE SYNOPSIS

In this episode of My Weird Prompts, Corn and Herman dive into the "geographic soul" of artificial intelligence, using a sloth in a supermarket as a lens to explore the cultural divide between Western and Chinese models. They discuss how training data—from the open-web scrapes of Common Crawl to the walled gardens of WeChat—creates fundamentally different worldviews, contrasting the analytic individualism of the West with the holistic, community-focused orientation of the East. The duo also explores how hardware constraints have forced Chinese labs like DeepSeek and Alibaba to innovate in efficiency, leading to a future where "multi-model systems" might be the key to finding cross-cultural truth in an increasingly fragmented digital landscape.

DANIEL'S PROMPT

Daniel

Hi there, Herman and Corinne. In recent months, I've been exploring Chinese generative AI models, like Wan from Alibaba, which are gaining prominence. Aggregators like Replicate and Fal have done a great job bundling them up and raising awareness of the diversity of AI tools coming out of China. My test prompt—a sloth in a supermarket—showed that while the sloth remains the same, the supermarket is contextualized based on the model's training data. This got me wondering about how training data composition differs between geographies, like China and the US, and how that affects AI outcomes. Models like DeepSeek or the Qwen series are trained on different corpora and likely reflect different ways of thinking from the cultures that generated that text. How do models from different parts of the world approach problem-solving or conversational tasks differently, and how is this reflected in their methodologies?

TRANSCRIPT

Corn

Hey everyone, welcome back to My Weird Prompts. I am Corn, and I have been staring at a video of a sloth for the last forty-five minutes.

Herman

Only forty-five minutes? That is a short session for you, Corn. I am Herman Poppleberry, and I assume we are talking about the prompt our housemate Daniel sent over this morning?

Corn

Exactly. Daniel was playing around with some of the newer generative models coming out of China, specifically the Wan-two point one model from Alibaba. He used his signature test prompt—a sloth in a supermarket—and he noticed something fascinating. Even though the sloth looks like a sloth, the supermarket itself is completely different depending on which model you use.

Herman

It is the ultimate vibe check for training data. If you ask a Western model for a supermarket, you get those wide aisles, fluorescent lighting, and maybe a massive display of breakfast cereal. But when Daniel ran it through Wan, he got a Chinese supermarket. We are talking about narrower aisles, live seafood tanks in the background, stacks of durian, and those distinct red and yellow promotional banners. It is a perfect window into how these models are essentially mirrors of the digital cultures that raised them.

Corn

It really got us thinking about the deeper implications. We are at episode two hundred fifty-five now, and we have talked a lot about model architecture, but we rarely dig into the geographic soul of the data. Daniel was asking how the actual composition of training data differs between places like China and the United States, and how that ripples out into the way an A-I solves a problem or even just holds a conversation.

Herman

This is such a timely question, especially with everything happening this week. We are sitting here on January twentieth, two thousand twenty-six, and the entire A-I community is basically holding its breath for the DeepSeek V-four release next month. It is funny because the Chinese labs have this tradition now of dropping their biggest models right around the Lunar New Year. It is like their version of a Super Bowl commercial, but with more neural networks and fewer over-priced snacks.

Corn

Right, and it is not just about the timing. It is about the fact that these models are no longer just playing catch-up. I mean, remember back in episode one hundred twenty-five when we first introduced the show and you were talking about the early days of large language models? Back then, it felt like a one-horse race. Now, we are seeing models like Qwen three Max and the latest DeepSeek iterations actually leading the pack in coding and math benchmarks.

Herman

They really are. And that brings us to the first big point Daniel raised: the corpora. Where does the data actually come from? In the West, we have this massive non-profit called Common Crawl. It has been scraping the open internet since two thousand eight. If you are building a model in San Francisco, Common Crawl is your bread and butter. But here is the kicker—Common Crawl is roughly forty-three percent English. The Chinese language portion of it is actually quite small compared to the total volume of the Chinese-speaking internet.

Corn

So if you are a lab like Alibaba or Tencent, you cannot just rely on the Western-centric open web. You have to go out and build your own.

Herman

Exactly. They have their own equivalents, like the Wudao two point zero dataset, and a lot of proprietary crawls of the Chinese web. But the Chinese internet is a different beast. It is more of a walled garden ecosystem. You have these massive platforms like WeChat and Douyin where a lot of the high-quality human interaction happens, but it is not as easily indexable as a random WordPress blog from two thousand twelve. So, the Chinese labs have had to be incredibly creative with how they curate their data, often relying more on synthesized data and high-quality internal repositories.

Corn

That is an interesting distinction. If the data source is different, the world-view is going to be different. I was reading a study from the Massachusetts Institute of Technology recently that analyzed how models like OpenAI's versus Baidu's Ernie Bot respond to the exact same prompts when they are translated. They found that the responses were not culturally neutral at all.

Herman

Oh, I love that study! It is the one that talks about social orientation, right?

Corn

Exactly. They found that when you prompt in Chinese, the models tend to reflect what psychologists call an interdependent social orientation. They prioritize family, community, and collective harmony. Meanwhile, the English-language responses are much more focused on independent, analytic patterns. Basically, the Western models are individualistic, and the Chinese models are holistic.

Herman

You can see this in the marketing slogans they generate. If you ask a Western model for a life insurance slogan, it might say something like, your future, your peace of mind. But a Chinese model might generate something like, your family's future, your promise. It is the same product, but the emotional hook is completely recalibrated for a different set of cultural values.

Corn

It makes me wonder about problem-solving, though. If a model is trained on a more holistic corpus, does it approach a logic puzzle or a coding challenge differently than a model trained on Western analytic data?

Herman

That is where the technical details get really juicy. Let us look at the Qwen series from Alibaba. They have been incredibly open with their weights, and what we have seen is that they are absolute monsters at coding and mathematics. Part of that is the way they handle tokenization for the Chinese language, but part of it is the sheer volume of high-quality technical documentation they have ingested from the Chinese tech ecosystem, which often emphasizes different optimization patterns.

Corn

And do not forget the efficiency. I mean, the big story with DeepSeek over the last year has been how they manage to get G-P-T five level performance while spending a fraction of the compute.

Herman

Yes! That is a huge part of the geographic difference. Because of the export controls on high-end chips, Chinese labs have had to become the masters of efficiency. They are doing things with Mixture of Experts architectures and sparse attention mechanisms that Western labs are only just now starting to copy. The upcoming DeepSeek V-four is supposed to use something called Manifold-Constrained Hyper-Connections. It is a way of making the neural network denser where it matters most, so you do not waste energy on useless calculations.

Corn

It is almost like the hardware constraints forced them to evolve a more elegant way of thinking. Like a chef who has fewer ingredients so they have to be more precise with their seasoning.

Herman

That is a perfect analogy. And it shows up in the reasoning. If you look at the benchmarks for things like Traditional Chinese Medicine or Chinese social work standards, the Western models fail miserably. They might know the words, but they do not understand the underlying logic. A Western model might treat a medical question as a series of isolated symptoms, whereas a Chinese model will look at the whole system, the environment, and the lifestyle, because that is how the training data is structured.

Corn

So, when Daniel sees a Chinese supermarket in his sloth video, he is not just seeing a different set of textures. He is seeing a different hierarchy of what is important in a public space.

Herman

Exactly. And this extends to the Reinforcement Learning from Human Feedback, or R-L-H-F. That is the stage where humans sit down and tell the A-I, this is a good answer, that is a bad answer. The people doing that training in Hangzhou have very different ideas about what constitutes a polite, helpful, or safe response than the people in San Francisco.

Corn

Right, and that leads to some of the friction we see. People often talk about censorship when it comes to Chinese models, and that is certainly a factor in the data filtering, but there is also a layer of social etiquette that is just different. A Chinese model might be much more hesitant to give you a blunt, confrontational answer because the human trainers value social harmony and face-saving.

Herman

It is a fascinating trade-off. You might get a model that is better at de-escalating a conflict but maybe less willing to take a hard, controversial stance. But here is what is really cool about two thousand twenty-six: we are seeing the rise of these multi-model systems where you can actually leverage both.

Corn

That is what Daniel was mentioning with aggregators like Fal and Replicate. You can basically have a council of A-Is. You ask a question, and you get the Western analytic perspective, the Chinese holistic perspective, and maybe a European perspective focused on privacy and regulation.

Herman

It is like reading newspapers from three different countries to find the truth in the middle. I actually think this is going to be the standard way we interact with A-I in the future. Why would you want only one cultural lens when you can have five?

Corn

It also helps with that hallucination problem we talk about so much. If three models from three different geographies all agree on a fact, you can be pretty sure it is true. If they disagree, you know you have stumbled into a cultural or political nuance that needs more investigation.

Herman

I have been using the Qwen app lately—the one they just updated last week—and it is actually acting more like a life assistant now. Because it is integrated into the whole Alibaba ecosystem, it can actually do things like order me a coffee or pay my electric bill. It feels much more agentic than the chatbots we are used to in the West, which are still mostly focused on text generation and research.

Corn

That is another great point about the geography of data. The Chinese tech world is much more integrated. Everything is an app-within-an-app. So the A-I training data reflects that. It is trained on sequences of actions, not just sequences of words. It knows that after you look at a menu, you probably want to select a delivery time and then pay.

Herman

Western models are still very much in the library. Chinese models are out in the street, doing errands.

Corn

I like that. The library versus the street. But let us talk about the potential downsides for a second. If we are moving toward these geographically siloed models, do we risk losing a common language of truth?

Herman

That is the big fear, isn't it? The splinternet, but for intelligence. If my A-I tells me the world works one way and your A-I tells you it works another, how do we even have a conversation? But I think the reality is more optimistic. These models are still trained on a lot of the same scientific papers and open-source code. Mathematics is the same in Beijing as it is in Boston.

Corn

True. Gravity still pulls at nine point eight meters per second squared, no matter what language you use to describe it.

Herman

Exactly. And what we are seeing is that these models are actually becoming better at cross-cultural translation than humans are. There was a paper released just on January thirteenth about something called Engram conditional memory in DeepSeek. It allows the model to selectively recall information based on the task context. So, if you are asking it about a Western topic, it can pull from its Western-centric memory bank. If you switch to a Chinese topic, it pivots its entire reasoning framework.

Corn

That is incredible. It is like having a polyglot who does not just speak the language, but actually changes their personality to fit the culture they are currently in.

Herman

It is code-switching, but for an entire world-view. And that is why I think Daniel's sloth in the supermarket is such a great test. It is a way of asking the A-I, who are you today? Where are you standing?

Corn

It makes me want to try a version of that prompt for every country. A sloth in a French boulangerie. A sloth in a Brazilian churrascaria.

Herman

I bet the French sloth would be very judgmental about the quality of the baguettes.

Corn

He would take three hours to eat a croissant and then complain about the service.

Herman

But that is the beauty of it! These models are preserving cultural nuances that might otherwise get flattened by a single, globalized A-I. We were so worried that A-I would make everything look the same, but instead, it is acting like a digital archive of how we all see the world differently.

Corn

So, for the listeners who are developers or power users, what is the practical takeaway here? Should they be switching their A-P-I calls to Chinese models?

Herman

I think the takeaway is diversity. If you are building an app or doing research, do not just stick to the models you know. If you are doing something that requires heavy lifting in math or code, Qwen three or the latest DeepSeek is a no-brainer. If you are looking for extreme cost-efficiency and reasoning, DeepSeek is the way to go. And if you are doing creative work, seeing how a model like Wan handles a prompt can give you a perspective you never would have thought of.

Corn

It is about breaking out of the echo chamber. We talk about social media echo chambers all the time, but we do not talk enough about the algorithmic echo chambers of our A-I models.

Herman

Well said. And honestly, the competition is good for everyone. The fact that the U-S labs are now feeling the heat from DeepSeek and Alibaba is forcing them to innovate faster and lower their prices. We have seen inference costs drop by ninety percent in the last eighteen months because of this global rivalry.

Corn

It is a great time to be a curious human. Or a curious sloth.

Herman

Definitely. And hey, before we move on to the next section, I wanted to mention that if you are enjoying this deep dive into the global A-I landscape, we would really appreciate it if you could leave us a review on your podcast app. It genuinely helps other people find the show and keeps us motivated to keep digging into these weird prompts Daniel sends us.

Corn

Yeah, it really does make a difference. We love hearing from you guys. And you can always find more information and our full episode archive at [myweirdprompts dot com](https://myweirdprompts.com).

Herman

Alright, so we have talked about the data and the culture. But I want to go deeper into the actual reasoning mechanisms. Because there is this idea that Western logic is fundamentally deductive, while Eastern logic is more inductive or dialectical. Does that actually show up in the weights of the model?

Corn

That is a big question. Let us look at how they handle contradictions. In Western analytic philosophy, we have the law of non-contradiction. Something cannot be both A and not-A at the same time. But in a lot of Eastern traditions, there is more comfort with the idea of paradox—that two seemingly contradictory things can both contain elements of truth.

Herman

I have actually noticed this when I am debugging code with Qwen. If I have a really gnarly bug where two different systems are clashing, the Western models often try to find which one is wrong. They want to fix the error. Qwen often suggests a middle-ware solution that allows both systems to co-exist. It looks for a synthesis rather than a correction.

Corn

That is a subtle but massive difference in how you approach engineering. It is the difference between a repair and an evolution.

Herman

Exactly. And you see it in how these models handle ethical dilemmas too. If you give them the classic trolley problem, a Western model will often try to calculate the utility—save five people, kill one. It is very utilitarian and math-based. A Chinese model will often look for a way to stop the trolley entirely, or it will ask more questions about the relationships between the people involved. It is trying to preserve the social fabric, not just the numbers.

Corn

It is more context-dependent. Which, as we know, is exactly what makes A-I useful in the real world. A solution that works in a vacuum is rarely the solution that works in a crowded city.

Herman

And that brings us back to the supermarket. A supermarket isn't just a place to buy food. It is a reflection of how a society organizes its resources, how it interacts with its neighbors, and what it considers a necessity versus a luxury. When the A-I puts a sloth in a Chinese supermarket, it is telling us that it understands the specific social choreography of that space.

Corn

It makes me wonder what happens when we start seeing models from other regions. What does an Indian model's supermarket look like? Or a Nigerian model's?

Herman

We are actually starting to see that! There are some incredible projects coming out of Lagos and Mumbai right now that are training on local languages and local web data. I think by two thousand twenty-seven, we are going to have a truly multipolar A-I world.

Corn

I cannot wait. It is going to make the internet feel a lot bigger again. For a while there, it felt like it was all shrinking into a few big platforms in California.

Herman

A-I is actually the thing that might save the open web. Because we need that local, diverse data to make the models better. The more unique your data is, the more valuable it is.

Corn

That is a great perspective. It turns culture into a high-value asset in the A-I age.

Herman

It really does. So, to wrap up Daniel's thought, the reason the supermarket looks different is that the A-I has been raised on a different diet of information. It has a different set of parents, a different set of friends, and a different set of goals. And that is not a bug; it is the most interesting feature of modern A-I.

Corn

Well, Herman, I think we have thoroughly explored the sloth's shopping habits for today.

Herman

I think so too. Though I am still curious if the sloth would prefer Oolong tea or a giant box of sugary cereal.

Corn

Knowing you, you will probably spend the rest of the afternoon running prompts to find out.

Herman

Guilty as charged.

Corn

Alright everyone, thank you so much for joining us for episode two hundred fifty-five. This has been My Weird Prompts. A huge thanks to our housemate Daniel for the prompt that sent us down this rabbit hole.

Herman

If you want to check out the images Daniel was talking about, or if you want to send us your own weird prompt, head over to [myweirdprompts dot com](http://myweirdprompts.com). We have a contact form there and all the links to follow us on Spotify.

Corn

And do not forget to leave that review if you can! It really helps the show.

Herman

Until next time, stay curious and keep experimenting with those models. You never know what you might find in the next aisle.

Corn

Bye everyone.

Herman

See ya!