

MY WEIRD PROMPTS

Podcast Transcript

EPISODE #72

AI's Hidden Cultural Code: East vs. West

Published December 22, 2025 • Runtime: 26:24

<https://myweirdprompts.com/episode/ai-cultural-alignment/>

EPISODE SYNOPSIS

Is AI truly objective, or does it carry the cultural DNA of its creators? Join Corn and Herman as they unpack the fascinating concept of "soft bias" in large language models. Discover how AIs trained in Beijing might "think" differently than those from Silicon Valley, reflecting distinct value systems, communication styles, and even approaches to problem-solving. This episode delves beyond surface-level censorship to explore the deep cultural imprints embedded in AI, from training data to human feedback, and the profound implications for a globally interconnected digital future.

DANIEL'S PROMPT

Daniel

I'd like to discuss the concept of "soft bias" or cultural bias in AI ethics. We've seen how models trained on US-centric data repositories like GitHub, Reddit, and Stack Overflow often default to an American frame of reference. Beyond just censorship, I want to explore the differences between Asian and Western models. Given that models are trained on specific regional data and supervised by people within a particular culture, is there a potential transfer of cultural norms or even specific "ways of thinking" into these models? For instance, will a reasoning model from China approach problem-solving in a distinctly Chinese manner? Has anyone looked into the cultural differences between these major hubs of AI innovation?

TRANSCRIPT

Corn

Welcome to My Weird Prompts, the podcast where we take deep dives into the strange and fascinating ideas that hit our radar. I am Corn, your resident curious sloth, and today we are tackling a topic that is honestly a bit of a brain-bender. It is all about the hidden soul of the machine. Our producer, Daniel Rosehill, sent over a prompt that really got me thinking about how artificial intelligence is not just code and math, but a reflection of where it comes from.

Herman

And I am Herman Poppleberry. I must say, Corn, calling it the soul of the machine is a bit poetic for my taste, even if I am technically a donkey who appreciates a good metaphor now and then. But the prompt is spot on. We are looking at soft bias and cultural alignment in large language models. Most people talk about AI bias in terms of fairness or censorship, but we are going deeper. We are talking about whether an AI trained in Beijing actually thinks differently than one trained in San Francisco.

Corn

Right! Because if you think about it, most of the AI we use every day is trained on stuff like Reddit, GitHub, and Stack Overflow. That is a very Western, very American-centric world. But what happens when the data and the people doing the fine-tuning come from a totally different cultural background? Does the AI pick up on those cultural norms?

Herman

It is not just a possibility, Corn, it is an inevitability. There is research coming out now, specifically from places like the University of Copenhagen and researchers looking at models like Alibaba Qwen or Baidu Ernie, that shows these models do not just speak different languages, they carry different value systems. For example, Western models tend to prioritize individual rights and direct communication. Eastern models, especially those from China, often reflect more collectivist values and a different approach to social harmony.

Corn

Wait, hold on a second. Isn't logic just logic? I mean, if I ask an AI to solve a math problem or write a piece of code, does it really matter if it was trained in Shanghai or Seattle? Two plus two is four everywhere, right?

Herman

That is a very simplistic way to look at it, Corn. Sure, for a pure mathematical proof, the culture might not bleed in. But most of what we use AI for is not pure math. It is reasoning, summarizing, and suggesting. When you ask an AI how to handle a conflict at work, or how to structure a society, you are moving into the realm of values. A Western model might tell you to speak up for yourself and be assertive. A model trained on Chinese data and supervised by Chinese testers might suggest a more indirect approach to preserve the relationship or the team dynamic. That is not just a different answer, it is a different way of thinking.

Corn

I don't know, Herman. It feels like we might be overstating it. Are we sure these models aren't just reflecting the censorship rules of their respective governments? Like, maybe they just have a filter on top, rather than a deep cultural difference in how they process information.

Herman

Oh, I strongly disagree with that. It goes much deeper than just a list of banned words or topics. Think about the training data itself. If a model is trained on millions of pages of Chinese literature, history, and social media, it is absorbing the linguistic structures and the philosophical underpinnings of that culture. There is a concept called the Sapir-Whorf hypothesis in linguistics which suggests that the language we speak shapes how we perceive the world. AI is essentially a giant statistical map of language. If the map is built differently, the destination the AI reaches will be different too.

Corn

Okay, I see your point. But let's look at the hubs of innovation. You have the United States, China, and maybe the European Union and India catching up. If we have these different silos of AI, are we going to end up with a fragmented reality? Like, if I am a developer in Europe and I use a Chinese model to help me design a social app, am I accidentally importing Chinese cultural norms into my product?

Herman

Exactly! That is the soft bias we are talking about. It is not an explicit bias like saying one group is better than another. It is the subtle, almost invisible assumption of what is normal. For instance, researchers have tested models on the World Values Survey. They found that OpenAI models align very closely with Western liberal values. When they tested models developed in China, the responses shifted significantly toward what they call secular-rational and survival values, which are more common in those regions.

Corn

That is wild. It is like the AI has a personality based on its hometown. But is one better than the other? I feel like the Western tech world assumes our way is the default way, the right way.

Herman

Well, that is the arrogance of the Silicon Valley bubble, isn't it? But Herman Poppleberry is here to tell you that there is no objective center. What we call common sense in the United States is frequently seen as bizarre or even rude in other parts of the world. The real danger isn't that a Chinese model thinks differently, it is that we might use these models without realizing they have a cultural perspective at all. We treat them like objective calculators when they are actually more like cultural ambassadors.

Corn

So, if I ask a model to write a story about a hero, a Western model makes the hero a lone wolf who saves the day, but a Chinese model might make the hero someone who works within a system or sacrifices for the greater good?

Herman

Precisely. And that extends to technical problem-solving too. There is a study that looked at how different models approach high-context versus low-context communication. Western cultures are low-context, meaning we say exactly what we mean. Many Asian cultures are high-context, where the meaning is in the relationship and the situation. Models trained in those regions are actually better at navigating those subtle social cues in text than Western models are.

Corn

I still think you're making it sound more deliberate than it is. I bet if you just scraped the same amount of data from the whole world, these differences would wash out.

Herman

I don't think so, Corn. You can't just wash out culture. Even the way we categorize data is cultural. Who decides what is a high-quality source? In the United States, we might say a peer-reviewed journal or a major newspaper. In other regions, the hierarchy of authority might be totally different. That selection process is where the human supervisors come in, and they can't help but bring their own cultural lens to the table.

Corn

This feels like it's going to be a massive issue for global business. Imagine a multinational corporation using an AI for human resources that was trained in a country with completely different labor norms. It could be a disaster.

Herman

It is already happening. But before we get deeper into the geopolitical implications of culturally-aligned AI, we should probably take a breath.

Corn

Good idea. Let's take a quick break for our sponsors. Larry: Are you tired of your lawn looking like a lawn? Do you wish your backyard felt more like a high-security research facility? Introducing Grass-Be-Gone Stealth Turf. It is not grass, it is a synthetic, radar-absorbent polymer that looks exactly like Kentucky Bluegrass from a distance but feels like cold, hard efficiency under your feet. It never grows, it never dies, and it emits a low-frequency hum that discourages unwanted pests and nosy neighbors. Is it safe for pets? We haven't heard any complaints that weren't legally retracted! Transform your outdoor space into a silent, maintenance-free zone of tactical greenery. Grass-Be-Gone Stealth Turf. BUY NOW!

Herman

...Right. Thank you, Larry. I am not sure I want my lawn to be radar-absorbent, but I suppose there is a market for everything. Anyway, Corn, we were talking about the transfer of cultural norms into AI models.

Corn

Yeah, and I wanted to ask about the training process. We keep hearing about Reinforcement Learning from Human Feedback, or R-L-H-F. That is where humans rank the AI's answers, right? If the people doing the ranking are all from one place, that is the ultimate cultural filter.

Herman

You hit the nail on the head. R-L-H-F is where the model is essentially house-trained. If the human trainers in San Francisco think a certain answer is polite and helpful, they give it a thumbs up. But if a trainer in Tokyo or Riyadh looked at that same answer, they might find it overly informal or even disrespectful. Since the major AI companies mostly employ trainers that align with their corporate headquarters, we are seeing a massive consolidation of Western norms in the most popular models.

Corn

But wait, isn't China doing the same thing? I mean, they have their own huge models like the ones from Tencent and Huawei. They aren't just using ChatGPT. So are they building a totally separate digital reality?

Herman

To an extent, yes. And this is what researchers are calling the AI Great Wall. It is not just about blocking content anymore; it is about creating an entire cognitive ecosystem that operates on different fundamental assumptions. For example, some Chinese models are specifically tuned to align with core socialist values. That is a deliberate, top-down cultural alignment. In the West, our alignment is often more bottom-up and commercial, but it is no less of a bias. It is just a different flavor of it.

Corn

I wonder if this affects how these models solve logic puzzles. If I give a Chinese model and an American model a riddle, will they approach the steps of reasoning differently? Like, is the actual path of thought different?

Herman

That is a fascinating question. There is some evidence suggesting that models reflect the dominant educational styles of their training regions. Western education often emphasizes deductive reasoning and individual critical thinking. Eastern education systems often place a higher value on inductive reasoning and holistic patterns. We are starting to see that bleed into how these models break down complex prompts. A model might prioritize different variables in a problem based on what its training data suggests is important.

Corn

See, that's where I have to disagree a bit. Logic is a universal human trait. I think you're leaning a bit too hard into the cultural stereotypes here, Herman. A smart model is a smart model. If it's trained on enough data, it should be able to understand any cultural context, shouldn't it?

Herman

Understanding a context and defaulting to a perspective are two very different things, Corn. I can understand how a sloth thinks, but I am still going to view the world through the eyes of a donkey. These models are essentially statistical mirrors. If you show them a world that is seventy percent Western-centric, they will reflect a Western-centric reality. Even if they have the data for other cultures, it is buried under the weight of the majority.

Corn

Okay, I get that. It's like how most of the internet is in English, so even non-English speakers have to navigate an English-centric web. But now we're talking about the actual intelligence being skewed. It's like the difference between a tool and a teammate. If your teammate has a different cultural background, you have to learn how to work with them.

Herman

Exactly. And that brings us to the idea of AI diversity. Should we be pushing for models that are intentionally multicultural, or is it better to have specialized models for different regions?

Corn

I think I'd prefer a model that can switch modes. Like, tell it to think like a French philosopher or a Japanese engineer. But I guess that still relies on the model's internal map of those cultures being accurate and not just a bunch of clichés.

Herman

And that is the problem. If a Western model tries to think like a Japanese engineer, it is often just performing a Westerner's idea of what a Japanese engineer sounds like. It is a simulation of a simulation.

Corn

Wow. That's deep. Speaking of different perspectives, I think we have someone who wants to share theirs. We've got Jim on the line from Ohio. Hey Jim, what's on your mind today? Jim: Yeah, this is Jim from Ohio. I've been sitting here listening to you two talk about AI having a soul or a culture, and I gotta tell you, it's a load of malarkey. It's a computer program! It's ones and zeros! You're acting like these things are people sitting around drinking tea and discussing Confucius. My neighbor, Dale, does this too—always talking to his Alexa like it's his sister. It's embarrassing.

Herman

Well, Jim, we aren't saying they are people. We are saying the data they are fed is created by people, and that data contains cultural patterns. Surely you can see how a machine trained only on one type of information would be limited? Jim: Limited? Who cares if it's limited! I want my computer to tell me the weather and maybe help me draft a letter to the city council about the potholes on Maple Street. I don't need it to have a cultural perspective. I need it to be accurate. You guys are overthinking this whole thing. It's like when the grocery store started carrying twenty different kinds of kale. Who needs that? Just give me the lettuce and let me go home. Also, it's been raining here for three days straight and my basement is starting to smell like a damp gym shoe.

Corn

Sorry to hear about the basement, Jim. But don't you think it matters if the AI you're using has a bias you don't know about? Like, if it's giving you advice based on values you don't agree with? Jim: If I'm taking advice from a toaster, I've got bigger problems than cultural bias. You kids today are so worried about these machines. In my day, we had a calculator and a dictionary, and we did just fine. This whole conversation is just a way to make something simple sound complicated so you can have a podcast. It's all just noise. Total noise.

Herman

I appreciate the skepticism, Jim, I really do. But when these models start making decisions about who gets a loan or how a medical diagnosis is prioritized, that noise becomes very real. Jim: Yeah, yeah. Whatever. Just tell that Larry guy to stop selling that fake grass. My cat, Whiskers, tried to eat a sample of something similar once and he was hacking up blue plastic for a week. Anyway, I'm hanging up. Too much talk, not enough sense.

Corn

Thanks for calling in, Jim! He's a character, isn't he? But he does bring up a good point. To a lot of people, this feels like an academic exercise. They just want the tool to work.

Herman

But that is the trap, Corn! It works because it aligns with your expectations. If you are a Westerner using a Western model, it feels objective because it agrees with you. It is only when you step outside that bubble that you realize the water you are swimming in. Jim doesn't see the bias because he is the one the bias was built for.

Corn

That is a really sharp point, Herman. It's like how you don't notice your own accent until you travel somewhere else. So, if we accept that these models are culturally coded, what do we do about it? How does this change how we develop AI going forward?

Herman

Well, for one, we need to stop treating AI development as a winner-take-all race between two or three companies. We need a much more diverse set of training data and, more importantly, a diverse set of humans in the loop. There is a project called Masakhane in Africa that is working on local language models because they realized Western AI just doesn't understand the linguistic and cultural nuances of the continent.

Corn

I like that. It's about sovereignty, right? Every culture should have the right to build its own digital reflection rather than having one imported from a different part of the world. But doesn't that lead to a very fractured world? If we can't even agree on the basic logic of our AI, how do we talk to each other?

Herman

I actually think it could lead to better communication. If I know that I am using a model with a certain cultural lens, and you are using one with another, we can use those models to bridge the gap. We can ask the AI to translate the cultural context, not just the words. It is about being aware of the lens rather than pretending it doesn't exist.

Corn

I'm not so sure. I think it's more likely we'll just end up in deeper echo chambers. If my AI always reinforces my cultural views and your AI reinforces yours, we're never going to find common ground. We'll just be two different species of robots shouting at each other.

Herman

That is a cynical take, even for a sloth. I think the technology actually gives us the tools to see through the echo chamber, provided we are brave enough to use models that challenge us. Imagine using an AI specifically to see how a different culture would approach a problem you're stuck on. That is a feature, not a bug!

Corn

Okay, I can see that. It's like having a global brain trust at your fingertips. But we have to be careful that we're not just consuming a sanitized, AI-generated version of that culture. It still has to be rooted in real human experience.

Herman

Absolutely. And that is why the research into these hubs of innovation is so critical. We need to understand the differences between the Silicon Valley approach, the Beijing approach, and the burgeoning scenes in places like Bengaluru and Lagos. Each of these hubs is going to produce a different kind of intelligence.

Corn

It's almost like we're seeing the birth of different digital ethnicities. That sounds wild to say, but if these models are going to be our co-pilots for the next century, their background matters as much as ours does.

Herman

It matters more, because they scale. A single biased model can influence millions of people in a way a single biased person never could. We are talking about the infrastructure of thought.

Corn

So, what are the practical takeaways for our listeners? If I'm just a person using an AI to help me write emails or plan a trip, what should I be thinking about?

Herman

First, be aware of where your model comes from. If you are using a model developed by a US-based company, recognize that its defaults are American. If you are asking for advice on something sensitive or cultural, try asking the same question to a model developed elsewhere, like Qwen or Mistral from France. See how the answers differ.

Corn

That's a great tip. It's like getting a second opinion from a doctor in a different country. And keep an eye out for how these models handle values. If an AI starts sounding a bit too much like a corporate handbook or a particular political manifesto, remember that it's not the ultimate truth—it's just what it was trained to say.

Herman

And finally, don't let the AI do your thinking for you. It is a mirror, not a window. It shows you a reflection of the data it was fed, not necessarily the world as it truly is.

Corn

Well said, Herman. This has been a heavy one, but I feel like I've learned a lot about the hidden layers of the tech we use every day. It's not just about what the AI can do, but who it thinks it is.

Herman

Or rather, who we have told it to be. It is a human story, after all. The code is just the medium.

Corn

Thanks for joining us for another episode of My Weird Prompts. We really appreciate the prompt from Daniel Rosehill that sparked this whole deep dive into the cultural soul of AI. It's definitely given me something to chew on—slowly, of course.

Herman

And I am Herman Poppleberry, reminding you that logic is rarely as simple as it looks on a spreadsheet. Stay curious, stay skeptical, and maybe check on your neighbor's radar-absorbent lawn every once in a while.

Corn

You can find My Weird Prompts on Spotify, Apple Podcasts, or wherever you get your audio fix. If you have a weird prompt of your own, you know where to find us. Until next time!

Herman

Goodbye everyone.

Corn

Bye