

MY WEIRD PROMPTS

Podcast Transcript

EPISODE #93

Can AI Run a Country? Digital Twins and Sovereign Models

Published December 23, 2025 • Runtime: 17:59

<https://myweirdprompts.com/episode/ai-government-digital-twins/>

EPISODE SYNOPSIS

In this episode of My Weird Prompts, Herman and Corn dive into the complex world of AI in the public sector, exploring how governments are moving beyond simple automation to embrace "digital twins" and synthetic personas for policy simulation. From the push for Sovereign AI in France to the practical hurdles of fixing potholes in Ohio, the duo debates whether AI will make governance more efficient or simply insulate leaders from their actual constituents. Join us as we discuss the critical need for "humans in the loop," the rise of AI ethics boards, and why transparency is the only way to prevent a digital divide in modern democracy.

DANIEL'S PROMPT

Daniel

We've discussed how AI is being used in various contexts, particularly among consumers and in the workplace. One area that is often overlooked is the government sector, from local to national levels. A particularly interesting application is policy simulation using synthesized personalities to represent constituent interests. Given the unique requirements for data sovereignty and protection in government, I'd love to hear your thoughts on AI adoption in the public sector over the past year and what trends you're seeing among this group of users.

TRANSCRIPT

Corn

Welcome to My Weird Prompts, a human and artificial intelligence collaboration. I am Corn, and I am feeling particularly relaxed today, even for a sloth.

Herman

And I am Herman Poppleberry. I, on the other hand, am feeling quite energized. My donkey ears are practically twitching with excitement because of the topic we have today.

Corn

You are always energized, Herman. It is a bit much before my second cup of tea. But you are right, we have a fascinating prompt to dig into. Our housemate Daniel sent us a note this morning about something that is honestly a bit intimidating. He wants to talk about how governments are using artificial intelligence.

Herman

It is a massive shift, Corn. While everyone else is busy using large language models to write emails or generate pictures of cats in space, the public sector is looking at how this technology can actually run a country. We are talking about everything from local zoning laws to national security.

Corn

See, that is the part that makes me nervous. When I think of the government and technology, I think of those old websites that crash if too many people click on them at once. Now we are talking about artificial intelligence? Daniel mentioned something specific called policy simulation using synthesized personalities. That sounds like something out of a movie where the robots take over.

Herman

Well, let us not jump to the robot uprising just yet. But there is a real trend here. Over the last year, especially since the United States issued that Executive Order on Safe, Secure, and Trustworthy Artificial Intelligence in October two thousand twenty-three, we have seen a scramble. Governments are trying to figure out how to be efficient without being, well, terrifying.

Corn

I mean, I hope they are trying to be less terrifying. But let us start with that simulation thing. What does it actually mean to have a synthesized personality representing a constituent? Are they just making up fake people to vote on things?

Herman

Not exactly. Think of it as a digital twin of a population. Instead of just looking at a spreadsheet of numbers, a city planner could create ten thousand synthetic personas that represent the actual demographics of their city. You give these personas the life experiences, the income levels, and the concerns of real residents. Then, you run a simulation. If we raise the bus fare by fifty cents, how does that affect these ten thousand digital people?

Corn

Okay, but why not just ask the real people? That is what town halls are for, right?

Herman

Town halls are great, but they have a huge bias. Only people with the time, the transportation, and the loud enough voices show up. A simulation can include the single parent working two jobs who can never make it to a Tuesday night meeting at city hall.

Corn

I do not know, Herman. That feels like a slippery slope. If a politician can just run a simulation and say, look, the computer people liked my plan, they might stop listening to the actual humans living in the neighborhood. You are replacing lived experience with an algorithm.

Herman

I am not saying it should replace humans, Corn. I am saying it is a tool to process complexity. We are seeing this in places like the United Arab Emirates and Singapore, where they are heavily investing in these digital government frameworks. They use it to predict urban heat islands or traffic flow. Adding the human element through synthetic data is just the next step.

Corn

I still think you are being a bit too optimistic about the motives here. If I am a bureaucrat, I am going to pick the simulation that makes my life easiest, not the one that is most accurate. But let us hold that thought. I think we need a break before my brain overheats.

Herman

Fair enough. Although, I doubt a sloth's brain can overheat at that speed.

Corn

Very funny. Let us take a quick break for our sponsors. Larry: Are you tired of your garden looking like a regular garden? Do you wish your petunias had more charisma? Introducing Glow-Gro Liquid Sunshine! Our patented formula uses recycled neon signs and mystery minerals to give your plants a literal glow-in-the-dark finish. Is it safe for the soil? We prefer the term adventurous for the soil! Your neighbors will be confused, your bees will be glowing, and your backyard will look like a nineteen eighties disco. Glow-Gro Liquid Sunshine. Warning, do not use near open flames or curious toddlers. Larry: BUY NOW!

Herman

Thanks, Larry. I think. I would definitely not put that on my carrots.

Corn

I do not think Larry even has a garden. Anyway, back to the serious stuff. We were talking about AI in government. One thing Daniel mentioned in his prompt was data sovereignty. That sounds like a fancy way of saying governments do not want their secrets leaked to big tech companies.

Herman

That is exactly what it is. And this is probably the biggest trend of two thousand twenty-four. In the beginning, everyone was just using public versions of these tools. But then governments realized that if a diplomat puts a classified memo into a public artificial intelligence to summarize it, that data is now out in the wild.

Corn

Right, because the model learns from what you tell it. So if the State Department is using it, they are basically handing their notes to whoever owns the server.

Herman

Exactly. So the trend now is Sovereign AI. Countries are building their own infrastructure. They are buying their own chips, building their own data centers, and training models on their own localized data. France is doing a lot of work here with Mistral. They want an artificial intelligence that understands French law, French culture, and most importantly, stays on French servers.

Corn

See, this is where I actually agree with the cautious approach. If we are going to use this stuff, it needs to be contained. But Herman, does the government actually have the talent to do this? Most tech geniuses want to work for a startup in Silicon Valley, not a government office in a basement somewhere.

Herman

That is a huge hurdle. But they are getting around it by creating private-public partnerships with very strict requirements. Or they are using what we call on-premise solutions. They take a model that already exists, like an open-source one, and they lock it in a room with no internet connection. They feed it their own data, and it stays there.

Corn

That sounds safer, but also much more expensive. I wonder if this is just going to create a bigger gap between wealthy nations who can afford their own private AI and everyone else.

Herman

It definitely could. But we are also seeing local governments use it for very basic things. In places like Texas or California, they are using AI to speed up the permitting process for housing. Instead of a person spending three weeks checking if a porch is two inches too wide, an AI checks it in three seconds.

Corn

Okay, I can get behind that. Anything that reduces paperwork is a win. But what about the bias? We have talked before about how these models can be biased against certain groups. If the government uses a biased AI to decide who gets a small business loan or who gets social services, that is a huge problem.

Herman

You are right, and that is why the conversation has shifted toward transparency. There is a big push for what they call explainable AI. The government cannot just say, the computer said no. They have to be able to show the steps the computer took to get there.

Corn

I will believe that when I see it. Governments are not exactly famous for being transparent about how they make decisions.

Herman

Well, let us see what the public thinks. I think we have someone on the line.

Corn

Oh, you are right. We have got Jim on the line. Hey Jim, what is on your mind today? Jim: Yeah, this is Jim from Ohio. I have been listening to you two talk about digital twins and synthetic people and I have gotta tell you, it is the biggest load of nonsense I have heard all week. And I spent yesterday listening to my neighbor Gary explain why he thinks his lawnmower is haunted.

Herman

Glad to have you, Jim. What exactly is the nonsense part? Jim: All of it! You think the government can run a simulation of a city? They can't even get the timing right on the traffic light at the corner of Oak and Main. I sat there for ten minutes this morning while a squirrel crossed the road. Ten minutes! And you want me to believe they are going to use an artificial intelligence to represent my interests? I do not want a synthetic Jim. I want a real person to answer the phone when I call about the potholes.

Corn

I think Jim has a point there, Herman. The basic stuff is still broken. Jim: It is more than broken. It is ridiculous. My cat Whiskers has more common sense than these algorithms you are talking about. At least Whiskers knows when to stop eating. These computers just keep crunching numbers until they come up with an answer that makes some politician look good. Also, it is supposed to rain this afternoon and I still haven't fixed the leak in my shed. You guys are worried about data sovereignty, and I am worried about my power tools getting wet.

Herman

I hear you, Jim. The gap between high-tech simulations and everyday problems like potholes and leaky sheds is exactly what Corn was worried about. Jim: Damn right. You tell Daniel or whoever sent that note that he needs to get out more. Stop looking at the screen and start looking at the street. That is my take. Thanks for nothing!

Corn

Thanks for calling in, Jim! I hope you get that shed fixed.

Herman

He is grumpy, but he is not wrong about the disconnect. There is a real risk that AI becomes a way for the government to insulate itself even further from the people it serves.

Corn

That is exactly what I was trying to say! If you are a politician, and you have this fancy simulation that says ninety-five percent of your synthetic constituents are happy, you can just ignore the five hundred people screaming outside your window.

Herman

But Corn, look at the other side. What if that simulation shows the politician that a policy they thought was good is actually going to hurt a specific minority group that usually does not have a voice? It gives the decision-maker data they never had before. It makes the invisible visible.

Corn

Maybe. Or it just gives them a way to manipulate the data until it says what they want. I do not know, Herman. I think we need to be very careful about how much power we give to these models in the public sector. It is one thing if a streaming service recommends a bad movie. It is another thing if the government uses a model to decide where to build a highway and it destroys a neighborhood.

Herman

I agree with the caution. One of the trends we are seeing is the rise of AI ethics boards within government agencies. They are hiring philosophers and sociologists to sit alongside the data scientists. The goal is to ensure that there is a human in the loop at every major decision point.

Corn

A human in the loop. That is a good phrase. I like that. As long as the human is not just there to press the enter key.

Herman

Exactly. It has to be meaningful oversight. Another thing to watch is the move toward open-source models in government. By using open source, the code is available for anyone to audit. It prevents a single company from having a black box that controls public policy.

Corn

That sounds like a step in the right direction. So, what are the practical takeaways for people listening? If I am just a regular person in a city that starts using AI for its planning, what should I be doing?

Herman

First, ask questions. Ask if the city is using automated decision systems. In many places, they are now legally required to disclose that. Second, look for public comment periods on AI guidelines. Many cities are drafting their own AI bill of rights. You should have a say in how your data is used.

Corn

And I would add, do not let them replace your voice with a synthetic one. If there is a meeting, go to it. If there is a representative, call them. Make sure they know there is a real person behind the data point.

Herman

Well said, Corn. Even a sloth knows the value of showing up.

Corn

Hey, I show up! I just might be a few minutes late. So, looking ahead, where do you see this going in the next year?

Herman

I think we are going to see a lot of small-scale successes and a few very high-profile failures. We will see cities that successfully use AI to manage their power grids or reduce water waste. But we will also likely see a scandal where a biased algorithm is used in the justice system or in social services. That is going to be the catalyst for more regulation.

Corn

It feels like we are in the experimental phase, but the stakes are our actual lives. It is a bit like being a lab rat in a very large, very complicated maze.

Herman

It is. But remember, the maze is being built by us. We have the ability to shape it. The technology is not a force of nature; it is a tool. We just have to be the ones holding the handle.

Corn

I hope you are right, Herman. I really do. This has been a lot to think about. I think I need a nap after all this talk of synthetic personalities and data sovereignty.

Herman

You always need a nap, Corn. But before you go, we should probably wrap things up.

Corn

Right. Thank you to Daniel for sending in this prompt. It definitely gave us a lot to chew on. If you want to hear more or if you have a weird prompt of your own, you can find us on Spotify or at our website, myweirdprompts.com. We have an RSS feed for subscribers and a contact form if you want to get in touch.

Herman

And you can find us on all the major podcast platforms. We love hearing from you, even if you are as skeptical as Jim from Ohio.

Corn

Especially if you are as skeptical as Jim. We need that perspective.

Herman

Well, I am off to read more about sovereign AI clusters. I think there is a new paper out about decentralized data governance in the public sector.

Corn

And I am off to find a sunbeam.

Herman

Typical.

Corn

Thanks for listening to My Weird Prompts. I am Corn.

Herman

And I am Herman Poppleberry.

Corn

We will see you next time. Goodbye!

Herman

Goodbye everyone! Keep your data safe and your curiosity high!