

## MY WEIRD PROMPTS

Podcast Transcript

### EPISODE #94

# Inside Smallville: Can AI Agent Villages Predict Humanity?

Published December 23, 2025 • Runtime: 23:27

<https://myweirdprompts.com/episode/ai-generative-agents-smallville/>

## EPISODE SYNOPSIS

In this episode of My Weird Prompts, Herman Popleberry and his brother Corn dive into the fascinating 2023 Stanford and Google study that populated a virtual town called Smallville with twenty-five generative AI agents. The duo explores how these digital entities use memory, reflection, and planning to exhibit emergent social behaviors—like spontaneously organizing a Valentine's Day party—and debates whether such simulations are revolutionary tools for social science or merely "expensive digital ant farms." From the potential for urban planning and software testing to the "empathy gap" and grumpy critiques from real-world callers like Jim from Ohio, this discussion challenges our understanding of what it means to model human community in an increasingly algorithmic world.

## DANIEL'S PROMPT

### Daniel

I'm interested in the experiments involving towns populated by AI agents, similar to an AI version of The Sims. Beyond being an interesting experiment, is this concept useful in any way? What are the practical applications or use cases for these AI villages, and have they been used to solve real-world problems or provide insights into human behavior?

# TRANSCRIPT

## Corn

Welcome to My Weird Prompts, the show where we take the strange, the obscure, and the downright confusing ideas and try to make sense of them. I am Corn, and as a sloth, I usually like to take things pretty slow, but today's topic actually has me leaning forward in my chair. I am here, as always, with my brother, Herman.

## Herman

That is Herman Poppleberry to you, Corn. And yes, while I may be a donkey, I can assure you my brain is moving at a much faster gallop than yours today. We have a really fascinating prompt to dive into. Our housemate Daniel sent this one over to us this morning. He was reading about these digital towns populated entirely by artificial intelligence agents.

## Corn

Yeah, Daniel was saying it is like an artificial intelligence version of The Sims. You know, that game where you tell people to go to the bathroom and then they accidentally set the kitchen on fire? But apparently, these agents are much smarter than that. They live in these virtual villages, they talk to each other, they form memories, and they even plan parties.

## Herman

It is officially called Generative Agents. The big study everyone points to was released in April of two thousand twenty-three by researchers at Stanford University and Google. They created a small virtual town called Smallville and populated it with twenty-five artificial intelligence agents. Each agent was powered by a large language model, specifically the same kind of technology behind things like Chat G-P-T.

## Corn

So, before we get into whether this is actually useful for the real world, we should probably explain what happened in Smallville. Because when I first heard about this, I just thought, okay, so it is a chat room with graphics. But it is more than that, right?

### Herman

Much more. See, what makes these agents different from a standard non-player character in a video game is their architecture. Most game characters follow a script. If A happens, do B. But these agents have a memory stream. They record everything they experience. They have a process called reflection, where they look back at their memories and draw higher-level conclusions. And they have planning. They actually decide what they want to do with their day based on their goals and their past experiences.

### Corn

Okay, but Herman, I have to jump in here. Is it really planning? Or is it just a very long, very complicated autocomplete? I mean, if I tell a computer to act like a person who likes coffee, it is going to say it wants coffee. That does not mean it is actually thirsty or making a choice.

### Herman

Well, hold on, that is a bit of a reductionist way to look at it. The interesting part is not that an agent says it wants coffee. The interesting part is the emergent behavior. For example, in the Smallville experiment, the researchers gave one agent, named Isabella, the idea to throw a Valentine's Day party. Isabella then spent her day telling other agents about the party. Those agents then told other agents. They coordinated who was going to show up, they discussed who they were going to bring as dates, and they actually showed up at the right time. The researchers did not script any of that coordination. It happened because the agents were reacting to each other in a social context.

### Corn

I don't know, Herman. It still feels a bit like we are just watching a very expensive digital ant farm. It is cool to look at, sure, but Daniel's prompt was specifically asking if this is useful. What is the actual point of building an artificial intelligence village besides proving that we can make robots talk to each other about pretend cake?

### Herman

That is where you are skipping over the real potential, Corn. Think about social science. If a sociologist wants to study how a rumor spreads through a town, or how a new policy might affect a community, they usually have to rely on historical data or very limited small-scale human experiments. Human experiments are expensive, they take forever, and there are huge ethical constraints. But with a town of artificial intelligence agents, you can run a simulation of a thousand different scenarios in an afternoon.

### Corn

Wait, I am not so sure about that. You are saying we can use these digital puppets to predict how real humans will act? Humans are messy. We have bad days, we have biological impulses, we get hangry. These agents are just... logic boxes. How can a simulation of a fake town tell us anything about how people in Jerusalem or New York or Tokyo are going to behave?

### Herman

See, I actually see it differently. We are not trying to predict exactly what one specific person will do. We are looking for patterns in social architecture. For example, if you are an urban planner and you want to see if building a new park in a specific area will actually increase community interaction, you could model that town. You could see if the agents' daily routines naturally lead them to congregate in that new space. It is a sandbox for testing how environments shape behavior.

### Corn

Mmm, I am still skeptical. It feels like we are adding a layer of abstraction that might just lead us to the wrong conclusions. If the model used to train the artificial intelligence has a bias, then the whole town has that bias. If the model thinks everyone is polite, then your simulated park looks like a utopia. But in the real world, someone is going to leave trash on the bench and someone else is going to play loud music at three in the morning.

### Herman

That is exactly why you run the simulation! You can actually program in those personality traits. You can give one agent a trait that says they are messy, or another agent a trait that says they are sensitive to noise. The point is not that it is perfect, it is that it is a tool for exploration that we have never had before. It is like how aeronautical engineers use wind tunnels. A wind tunnel is not the sky, but it lets you test how a wing might react to turbulence before you build the actual plane.

## Corn

I guess that makes sense. But speaking of things that are a bit turbulent, let us take a quick break from our sponsors. Larry: Are you tired of your thoughts being too loud? Do you wish you could just put your brain on mute for a while? Introducing the Cranial Cozy! It is not a hat, it is a revolutionary sound-dampening environment for your entire skull. Made from recycled insulation and mystery fibers found in an abandoned textile mill, the Cranial Cozy uses patented static-cling technology to adhere directly to your ears. Perfect for family gatherings, long walks in the rain, or when you just want to pretend you are a very quiet rock. Warning: may cause temporary loss of equilibrium and an intense craving for radishes. The Cranial Cozy. Because silence is golden, but muffled silence is even better. BUY NOW!

## Herman

Alright, thanks Larry. I think. Anyway, back to the artificial intelligence villages. Corn, you were asking about practical applications beyond just social science research. There is actually a lot of movement in the software development space.

## Corn

Like what? Are they going to have the agents write the code for us?

## Herman

Not exactly. Think about how we test software today. Usually, you have a human tester who tries to break the app. Or you have automated scripts that follow a very specific path. But if you have an environment with artificial intelligence agents, you can let them use your software as if they were real users. If you are building a new social media platform, you can populate it with five hundred agents, give them different goals and personalities, and see how they interact. Do they start being mean to each other? Does the algorithm show them things they actually want to see? It is a way to find edge cases that a human tester might never think of.

## Corn

Okay, that actually sounds useful. Especially for things like testing how a new tax law or a basic income program might work. Instead of just guessing, you could run a simulation of a small economy and see if the agents stop working or if they start new businesses. But there is still a part of me that thinks this is a bit dangerous. If we start trusting these simulations more than we trust real human feedback, aren't we just creating a feedback loop of artificial ideas?

### Herman

That is a valid concern, but I think you are overstating the risk of replacement. No one is saying we should stop talking to real people. But think about the scale. You cannot ask ten thousand people to participate in a three-month study on a new urban transit system every time you have a new idea. But you can run that simulation as a first draft. It helps you narrow down the best ideas so that when you finally do go to the real humans, you are not wasting their time with something that was obviously going to fail.

### Corn

I suppose. It is like a high-tech brainstorming session. But I still wonder about the empathy gap. An artificial intelligence agent can simulate being upset, but it does not actually feel the consequences of a bad policy. If a simulated town goes bankrupt, you just hit the reset button. In the real world, people lose their homes. I worry that policy makers might become detached if they spend too much time looking at the Sims version of the world.

### Herman

That is a deep philosophical point, Corn, and I actually agree with you to an extent. We have to be careful not to treat these simulations as the absolute truth. They are models. And as the saying goes, all models are wrong, but some are useful. The goal is to find the useful parts without forgetting that the map is not the territory.

### Corn

Speaking of the territory, I think we have someone from the real world on the line. 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 these digital towns and I have to say, it sounds like a load of bunk. You are telling me we are going to spend millions of dollars to watch computer programs pretend to live in a house? My neighbor, Frank, spends all day watching the squirrels in his yard, and I am pretty sure he is getting more useful data than these Stanford scientists. Probably costs him a lot less, too. Just a bag of peanuts.

### Herman

Well, Jim, I hear your skepticism, but the goal is to understand complex systems that are hard to see just by watching squirrels. We are talking about how ideas move through a population. Jim: Ideas? I will tell you about an idea. Here is an idea: how about we fix the actual potholes on my street instead of building a digital street with digital potholes? I hit one this morning and I thought my teeth were going to fall out. And don't even get me started on the weather. It has been raining for three days straight, and my cat, Whiskers, refuses to leave the laundry room. He just sits there staring at the dryer like it is a television. It is unnatural.

### Corn

I am sorry to hear about Whiskers, Jim. And the potholes. But don't you think there is some value in testing things out in a safe environment before we try them in the real world? Jim: Safe environment? Nothing is safe when you involve computers. My grandson tried to set up a smart toaster last week and now every time I want a piece of rye bread, the thing tries to update its firmware. I just want toast! I don't want a conversation with an appliance. And you guys are talking about entire towns of these things. It is just more noise. People used to talk to their neighbors. Now they want to talk to a simulation of a neighbor. It is backwards.

### Herman

I think you are touching on a real fear, Jim, which is the loss of human connection. But the researchers aren't trying to replace neighbors. They are trying to understand how to build better communities for real people. Jim: Well, they could start by asking me. I have lived in this town for fifty years. I know what makes a community work. It is not an algorithm. It is having a decent hardware store and a place to get a sandwich that doesn't cost fifteen dollars. Anyway, I have to go. Frank is out there with the squirrels again and I think he is teaching them how to use a slingshot. Thanks for nothing!

### Corn

Thanks for calling in, Jim! He is always a breath of fresh air, isn't he?

### Herman

He certainly has a way of grounding the conversation in the frustrations of the present. But even Jim's complaints are a type of data, in a way. He is talking about the friction of daily life. And believe it or not, there are researchers working on how to incorporate that kind of friction into artificial intelligence agents. They are trying to make them less perfect, less helpful, and more like real people who get annoyed by potholes and smart toasters.

### Corn

If they can simulate a grumpy Jim from Ohio, then I will be truly impressed. But let us get back to the practical stuff. We talked about social science and software testing. What about things like disaster response? Could these villages be used to train for emergencies?

### Herman

Absolutely. That is actually one of the most promising use cases. Imagine you are trying to plan an evacuation for a city during a hurricane. You can use these agents to simulate how different types of people might react to different messages. If you send a text alert, who leaves immediately? Who stays to look for their cat? Who gets stuck in traffic because they took a specific route? By giving the agents different personalities and priorities, you can see where the bottlenecks in your plan might be. It is much more realistic than just assuming everyone will follow instructions perfectly.

### Corn

That makes a lot of sense. In a crisis, people don't always act logically. If the agents can mimic that kind of emotional decision-making, it could actually save lives. But Herman, how do these agents actually talk to each other? Is it just like a giant group chat?

### Herman

It is more structured than that. When two agents in a simulation like Smallville come near each other, the system triggers an interaction. The model looks at the context—where they are, what time it is, and what they know about each other from their memory streams. Then it generates a dialogue. If agent A knows that agent B is looking for a job, and agent A just heard about a job opening, the model will likely have agent A mention it. It is all based on retrieving relevant memories and using them to inform the current conversation.

### Corn

So, it is like they have a little brain that is constantly searching their past for what to say next. I guess that is not too different from how we work. I am usually thinking about what I had for lunch while I am talking to you.

### Herman

Exactly! And that leads to another use case: personalized education and training. You could create a simulated environment where a student has to interact with artificial intelligence characters to practice a new language or to learn conflict resolution. Instead of just reading a textbook, you are actually practicing in a social setting that feels real but has zero stakes if you mess up.

### Corn

Okay, I am starting to see the appeal. It is about creating a low-risk environment for high-risk or high-cost activities. Whether that is testing a new economy, planning a city, or learning how to de-escalate an argument. But what is the catch? This technology must be incredibly expensive to run.

### Herman

You hit the nail on the head. Right now, the computational cost is the biggest hurdle. Running twenty-five agents in a town like Smallville for just a few days cost thousands of dollars in tokens from the large language model providers. If you wanted to scale that up to a town of ten thousand agents for a year-long simulation, the cost would be astronomical. We are talking millions, maybe billions of dollars with current technology.

### Corn

So, for now, it is mostly for researchers with big grants and tech companies with deep pockets. It is not like I can just start a simulated version of my neighborhood on my laptop to see if I should open a lemonade stand.

### Herman

Not yet. But as these models become more efficient and specialized, the cost will come down. We are already seeing smaller, more efficient models that can run locally on powerful consumer hardware. In five or ten years, you might actually be able to run a small simulation like this at home.

### Corn

That is a bit wild to think about. I could have a little digital Jerusalem running on my desk. I would probably just use it to see which bakery has the shortest line on Friday morning.

### Herman

See? Even your lazy sloth ideas have a practical application! But there is one more thing we should touch on before we wrap up, and that is the ethics of simulation.

### Corn

Oh boy, here we go. Is this where you tell me that we might all be living in a simulation right now?

### Herman

No, I will leave that for another episode. I mean the ethics of how we use these agents. If an artificial intelligence agent is sophisticated enough to have memories, plans, and reflections, do we have any responsibility toward it? Or is it just code? And more importantly, if we use these simulations to make big decisions about real people, what happens when the simulation is wrong? We have already seen how algorithms can reinforce biases in hiring or policing. If we start using artificial intelligence villages to design our society, we risk baking those biases into the very foundation of our world.

### Corn

That is what I was getting at earlier. It feels like we are giving a lot of power to a black box. It is one thing to use it for a video game, but it is another thing to use it for city planning or disaster response. We have to make sure there is always a human in the loop who can say, wait, that doesn't look right.

### Herman

I agree. The simulation should be a tool for thinking, not a replacement for thought. It gives us a new way to ask what if, but we still have to be the ones to decide what to do with the answer.

### Corn

Well, I think that is a good place to leave it. We covered a lot of ground today. From Valentine's Day parties in Smallville to disaster planning and the high cost of digital brains. And of course, Jim's cat Whiskers and his dryer.

### Herman

It is a lot to chew on. But that is what Daniel's prompts are for. They get us out of our comfort zone. Even if your comfort zone is just a very soft branch on a tree.

### Corn

Hey, it is a very nice branch. Anyway, if you want to dive deeper into this, you can find more information about the Generative Agents study online. It is a fascinating read if you have the time and the patience for academic papers.

### Herman

And if you don't, well, that is what we are here for. We will keep an eye on how these digital towns evolve. Maybe one day we will be recording this podcast inside a simulation.

### Corn

I think I would prefer the real Jerusalem, personally. The food is better. Thanks for listening to My Weird Prompts. You can find us on Spotify, or check out our website at [myweirdprompts.com](https://myweirdprompts.com). We have got an R-S-S feed for subscribers and a contact form if you want to send us your own weird prompts, just like our housemate Daniel did.

### Herman

Or if you want to complain like Jim from Ohio. We take all kinds here. I am Herman Popleberry, and I hope you learned something today.

### Corn

And I am Corn. We will be back next week with another deep dive into the strange and wonderful. Until then, stay curious, and maybe go talk to a real neighbor. It is cheaper than a simulation.

### Herman

Goodbye, everyone!

## Corn

Bye! Larry: BUY NOW!