

# MY WEIRD PROMPTS

Podcast Transcript

## EPISODE #47

# From Sketch to Studio: AI & Control Nets in Design

Published December 10, 2025 • Runtime: 27:13

<https://myweirdprompts.com/episode/generative-ai-in-architecture-and-creative-industries-lvcpt2k/>

## EPISODE SYNOPSIS

Get ready to see architecture and design through a revolutionary lens! In this episode of "My Weird Prompts," hosts Corn and Herman dive deep into how generative AI, specifically "control nets," transforms abstract design sketches into stunning photorealistic renderings and immersive virtual walkthroughs. Discover how architects leverage these advanced tools to accelerate visualization and overcome traditional design hurdles, making complex concepts tangible for clients. The discussion explores the technical intricacies of co-located AI models, the crucial role of cloud platforms in democratizing this power, and the delicate balance between user accessibility and the professional expertise required to achieve breathtaking, precise results.

# TRANSCRIPT

## Corn

Welcome back to "My Weird Prompts"! I'm Corn, your perpetually curious host, and I'm ready to dive into another fascinating topic with my brilliant co-host.

## Herman

And I'm Herman Poppleberry, here to bring some much-needed precision and expertise to Corn's enthusiastic meanderings. It's good to be back, Corn.

## Corn

Precision, yes, that's what we love about you, Herman. And today's prompt, sent in by our very own producer, Daniel Rosehill, is all about the bleeding edge of creativity and technology, specifically in the world of architecture and design.

## Herman

Indeed. What's particularly intriguing about this prompt is its focus on how generative AI is transitioning from mere novelty to an indispensable tool in professional workflows, particularly for visual communication. Most people hear "AI art" and think funny images, but this goes so much deeper.

## Corn

Yeah, it's not just about making a cat wearing a monocle anymore, is it? We're talking about architects taking a simple sketch, a pencil drawing, and using AI to conjure a photorealistic rendering, or even a full virtual walkthrough. My mind is already buzzing with possibilities!

## Herman

Well, let's temper that buzz with a bit of technical reality, Corn. While the results can certainly appear magical, the underlying process involves sophisticated techniques, particularly the use of something called "control nets." And that's really the crux of Daniel's prompt today: how these advanced tools are being used, and more importantly, *where* they're being used, especially in a cloud-based environment.

## Corn

Right, because for most of us, we don't have a supercomputer humming away in our spare room, do we? This isn't just a fun app; this is about professionals who need to deliver high-fidelity visuals to clients. It's like, instead of trying to imagine a building from a blueprint, you can literally walk through it before it's even built. That's game-changing!

## Herman

Absolutely. The ability to translate an abstract concept, like a line drawing or a basic sketch, into a detailed, photorealistic 3D environment instantaneously is a monumental leap. Traditionally, this process would involve hours, if not days, of work from skilled 3D artists and renderers. AI, specifically with control nets, compresses that timeline dramatically while maintaining an astonishing degree of control over the output.

## Corn

Okay, so let's break down what a "control net" actually *is* for those of us who aren't steeped in the AI jargon. Is it like a really smart filter, or something more fundamental? Because when I hear "control," I think about telling the AI exactly what to do.

## Herman

That's a good intuitive leap, Corn. Think of a control net as a way to give a large generative AI model, like a stable diffusion model, very precise instructions or constraints. Without a control net, if you just give a diffusion model a text prompt like "a modern house interior," it will generate something, but you have very little control over the layout, the exact geometry, or the precise composition.

## Corn

So it would just make *a* modern house, not *my* modern house.

## Herman

Precisely. A control net allows you to feed in additional structural information alongside your text prompt. This could be things like a depth map, which defines the 3D structure of the scene; a Canny edge map, which extracts the outlines of objects; or even a normal map, which describes surface orientation. The control net then guides the diffusion model to generate an image that *\*adheres\** to those structural inputs.

## Corn

Oh, I get it! So, if I draw a simple floor plan or a perspective sketch, the control net takes that basic structure and tells the AI, "Okay, render this, but make sure the walls are here, the windows are there, and the furniture is arranged like *\*this\**." It's not just generating; it's *\*following instructions\** at a fundamental level.

## Herman

Exactly. And this is why it's so powerful for architects. Daniel's wife, for instance, can sketch out an interior design, and the control net uses that sketch as a foundational blueprint, allowing the diffusion model to fill in the details with photorealistic textures, lighting, and materials, while preserving the spatial relationships from the original sketch. This turns a vague idea into a concrete visual representation almost instantly.

## Corn

That's incredible. It makes visualization so much more accessible, especially for clients who, as Daniel mentioned, might struggle to "see" a design from abstract plans. You can give them a virtual walkthrough, and they can experience the space.

## Herman

It democratizes high-fidelity visualization, yes. But here's where I'd push back on the idea of it being "easy." While the *\*user experience\** can be simplified, setting up these pipelines and ensuring the quality and consistency of the output still requires significant technical understanding and careful workflow design. It's not just plug-and-play for professional results.

## Corn

Well, I don't know about that, Herman. I mean, Daniel himself mentioned he's interested in this for a children's book. That's not a hyper-technical architectural project. If he can use it, it can't be *that* hard, can it? Aren't tools like ComfyUI trying to make it more accessible for everyone?

## Herman

ComfyUI, by itself, is a node-based interface that gives you granular control, which is fantastic for power users and developers, but it also means there's a steeper learning curve than, say, a simple web app. And while the *concept* of using control nets is broadly applicable, the *level of precision* required for architectural renderings, where structural integrity and accurate material representation are paramount, is far higher than what you'd need for, say, a whimsical children's book illustration. You need to understand the nuances of the various control net models, how they interact, and how to fine-tune them for specific outputs. It's not just a "make pretty picture" button.

## Corn

Okay, I see your point. It's like, the raw tools are there, but mastering them for a specific, high-stakes domain like architecture still requires expertise. So, what about the "where"? Daniel's prompt brought up the question of running these complex operations without owning a massive GPU farm. This is where cloud computing comes in, right?

## Herman

Exactly. This brings us to the core technical challenge Daniel highlighted. To run a diffusion model with a control net, both components typically need to be "co-located." That means they must be running on the same server, or at least on compute resources that are in very close proximity, often within the same node or machine. You can't have your control net running on your laptop and expect it to seamlessly integrate with a diffusion model running on a remote cloud server across the internet with any kind of real-time performance or stability. The data transfer alone would be a nightmare.

## Corn

So, if I'm understanding this, all the heavy lifting, the control net processing, and the diffusion model generation, needs to happen in one place, like on one powerful machine or within a single, unified cloud instance.

### Herman

Precisely. That's why services like RunPod, Replicate, or Fal, which Daniel mentioned, are becoming crucial. They provide the cloud infrastructure and often pre-configured environments where you can deploy your ComfyUI workflows, complete with the necessary GPUs, and ensure that your control net models and diffusion models are co-located for optimal performance.

### Corn

Ah, so it's not just about renting a server; it's about renting a server that's specifically set up to handle these kinds of interconnected AI workloads. So then the question becomes, how "arbitrary" can you be? Can I just plug any control net model I find online into any image-to-image diffusion model I want, as long as it's on the same cloud instance? Or are there compatibility issues?

### Herman

That's an excellent question, Corn, and it touches on the open-source nature of many of these tools. In general, thanks to the standardization within the open-source community, many control net models are designed to be quite interoperable with various stable diffusion models. They typically work by taking a specific input type – say, a Canny edge map – and producing an output that guides the diffusion process, regardless of the specific diffusion model you're using.

### Corn

So, I *could* technically use a control net trained for general image manipulation with a fine-tuned architectural diffusion model?

### Herman

Yes, within certain bounds. As long as the control net's input type matches what you're providing (e.g., you're feeding it a depth map, and it's expecting a depth map), and its output format is compatible with the diffusion model, it *should* work. However, the *quality* of the results can vary wildly. A control net specifically trained on architectural data or paired with a diffusion model fine-tuned for architecture will almost certainly yield superior results for that domain. It's about maximizing the synergy between the components.

## Corn

So, it's less about strict "compatibility" and more about "optimal performance" for professional-grade output. You *can* mix and match, but if you want the best architectural rendering, you'll want control nets and diffusion models that are singing the same song, so to speak.

## Herman

Exactly. And the "precise formula" Daniel asked about often refers to the specific combination of models, parameters, and preprocessing steps that a professional studio has found yields consistent, high-quality results for their specific use cases. It's less about a strict technical limitation and more about established best practices for achieving excellence.

## Corn

Let's take a quick break from our sponsors. Larry: Are you tired of that nagging feeling that you're just not *shiny* enough? Do your shoes lack that certain je ne sais quoi? Introducing **GlimmerGlo All-Surface Elixir!** This revolutionary, non-Newtonian fluid will instantly transform anything dull into a dazzling spectacle. Made from 100% sustainably harvested moon dust and purified essence of forgotten dreams, GlimmerGlo works on leather, metal, plastics, even old leftovers! No scrubbing, no buffing, just apply and watch the magic happen. **WARNING:** May cause temporary blindness in direct sunlight and an uncontrollable urge to wear sequined apparel. GlimmerGlo: Because life's too short to be subtle! **BUY NOW!**

## Herman

...Alright, thanks Larry. Anyway, where were we? Ah, yes, optimizing performance and the interplay of different models. Corn, you brought up the children's book example from Daniel. This is a perfect illustration of how the requirements change based on the end goal. For an architect needing precise structural integrity, the "arbitrary" approach might introduce too many variables. For a children's book where a bit of unexpected whimsy might even be a *feature*, a more experimental approach is perfectly viable.

## Corn

So, it's about matching the tool and the workflow to the outcome. What about these services like Replicate and Fal? How do they fit into this cloud-based workflow?

## Herman

They abstract away a lot of the complexity. Replicate, for example, provides an API and infrastructure for running open-source models. You can send your input, tell it which model and control net you want to use, and it handles the underlying compute and co-location. Fal goes a step further, offering its own visual workflow builder, which can integrate ComfyUI components but provides a more managed, streamlined experience. These services are essentially making powerful AI infrastructure accessible without you needing to manage the servers yourself. They're like the managed hosting of the AI world.

## Corn

So, for someone like Daniel, wanting to experiment for a children's book, or even a small architectural firm that doesn't want to invest in a massive in-house setup, these cloud services are the way to go. They handle the "precise formula" of setting up the hardware and co-locating the models, letting the user focus on the creative input.

## Herman

Exactly. It lowers the barrier to entry for utilizing these powerful tools. However, there's always a trade-off. While convenient, using these managed services might offer less granular control over the underlying infrastructure or specific model versions than if you were deploying ComfyUI directly on a raw cloud VM or on-premises server you fully control. For maximum customization and cutting-edge research, some will still prefer direct access.

## Corn

But for 99% of people, or even 99% of professionals, that direct access is overkill. They just want the tool to work, and work reliably.

## Herman

I'd agree with that. The balance shifts from needing deep infrastructure knowledge to needing deep \*domain\* knowledge—understanding how to effectively prompt the AI, how to prepare your input sketches, and how to interpret and refine the output. The tools become enablers for creative professionals, rather than requiring them to become IT specialists.

## Corn

That's a really important distinction. And it brings us to another part of the show...

## Corn

And we've got Jim on the line – hey Jim, what's on your mind? Jim: Yeah, this is Jim from Ohio. Been listening to you two go on and on about these... control nets. Sounds like a whole lot of fancy talk for not much. You know, back in my day, if an architect wanted to show a client something, they made a physical model. With their hands! Or they drew a blueprint and if the client couldn't understand it, well, that was their problem for not having an imagination. All this computer wizardry, I don't buy it. My cat, Whiskers, could probably draw a better blueprint with a crayon and she mostly just sleeps all day.

## Herman

Well, Jim, I appreciate the nostalgia for physical models, and they certainly have their place. But the efficiency and iterative capability of generative AI means architects can explore dozens of design variations in the time it would take to build one physical model. This allows for far more comprehensive design exploration and client feedback cycles. Jim: Iterative capability, shmiterative capability. What's wrong with just getting it right the first time? My neighbor, Gary, he builds birdhouses and he doesn't need any "control nets" to make them perfect. He just uses a saw and some nails. And he always complains about how windy it gets here in Ohio in the spring. Anyway, you guys are making it sound like nobody could do good design before these computers came along, and that's just hogwash.

## Corn

No, no, Jim, nobody's saying that! Great design has always existed. What we're saying is that these tools are \*augmenting\* the design process, making it faster, more detailed, and more accessible for communication. It's about empowering designers, not replacing them. Think of it like a really advanced drafting tool. Jim: Eh, I don't know. Seems like it just gives people an excuse not to learn how to draw properly. You just type in a few words and poof, there's a building. Sounds like cheating to me. And frankly, all this talk about cloud servers and "co-location" makes my head spin. I like my technology to be simple. Like my old rotary phone. That always worked.

### Herman

With respect, Jim, it's far from just typing in a few words. The "skill" shifts from the manual dexterity of traditional drafting to the intellectual skill of guiding and refining the AI's output, understanding its parameters, and curating its suggestions. It's a new form of craftsmanship, not a replacement for it. And those cloud servers are precisely what allow smaller firms or individual designers to access this power without the prohibitive cost of owning the hardware themselves.

### Corn

Yeah, it's like using a really good camera. You still need a skilled photographer to get a great shot, even if the camera does a lot of the technical stuff automatically. You still need the eye, the vision. Jim: Well, I suppose. But still, seems like a lot of fuss. My back's been acting up something fierce lately, probably from raking all those leaves last fall. But whatever. You two carry on with your fancy robot architecture. I'm sure it'll make perfectly soulless buildings.

### Corn

Thanks for calling in, Jim! Always a pleasure to hear from you.

### Herman

Indeed, Jim, always enlightening.

### Corn

Alright, so, moving from Jim's healthy skepticism to what people can actually \*do\* with this. What are the practical takeaways for architects, designers, or even just creative folks interested in high-fidelity AI generation?

### Herman

The biggest takeaway is that generative AI, specifically with control nets, is no longer a niche curiosity. It's a professional-grade tool for visualization and design exploration. For architects, it means faster iteration, more compelling client presentations, and the ability to explore a wider array of design options in a fraction of the time.

### Corn

So, for smaller firms, it's about staying competitive without breaking the bank on hardware, thanks to cloud services.

### Herman

Precisely. Investigate cloud platforms that offer ComfyUI or similar node-based interfaces, paying close attention to their GPU offerings and pricing models. Services like RunPod, Replicate, or Fal are excellent starting points for getting access to powerful compute without the capital expenditure.

### Corn

And for individual creatives, like Daniel with his children's book idea, it means the ability to produce high-quality visuals without needing to be an expert in traditional rendering software.

### Herman

Yes, but with the caveat that understanding the input types for various control nets – depth, Canny, normal maps – and how they influence the diffusion model is crucial for achieving desired results. It's not just about prompting; it's about providing intelligent visual guidance. My advice would be to start with well-documented open-source models and experiment with simple inputs to understand the core mechanics before diving into complex workflows.

### Corn

And don't be afraid to mix and match. You might find that a control net intended for one purpose actually works really well for a different creative outcome, even if it's not "architecturally precise."

### Herman

That's true for experimental, lower-stakes projects. But for professional architectural work, always prioritize models and workflows that have proven robustness and fidelity. Look for communities or resources that share optimized ComfyUI workflows for specific architectural applications. The "precise formula" is often shared within these communities.

### Corn

So, the human element isn't gone; it's just shifted. Instead of drawing every brick, you're designing the system that draws the bricks, and then curating the output. It's a new form of digital craftsmanship.

### Herman

A highly leveraged form of craftsmanship, I'd say. It allows designers to focus on the high-level aesthetic and functional aspects of their work, letting the AI handle the repetitive or time-consuming rendering tasks.

### Corn

Looking ahead, Herman, where do you see this headed? More integration? Even greater autonomy for the AI?

### Herman

I foresee deeper integration within existing CAD and BIM software. The current setup often involves moving data between applications. Future iterations will likely see these generative AI capabilities, including control nets, embedded directly within design tools, making the workflow even more seamless. We'll also see more specialized control nets and diffusion models, fine-tuned for incredibly niche applications, pushing the boundaries of what's possible.

### Corn

And what about the ethical considerations? Like, what happens when an AI generates a design that looks suspiciously like something an existing architect has already done?

### Herman

That's a valid concern, Corn, and one the industry is actively grappling with. Issues of intellectual property, originality, and the provenance of training data will continue to be critical discussions. For now, it's important for users to understand that AI is a tool, and the ultimate responsibility for ethical use and originality rests with the human designer. It's about how you *use* the tool, not just that you *have* it.

**Corn**

Fascinating stuff, Herman. Absolutely mind-bending what these tools can do when put in the right hands, and the right cloud. It makes you wonder what Daniel's next prompt will be.

**Herman**

No doubt it will be equally thought-provoking.

**Corn**

Absolutely. That wraps up another episode of "My Weird Prompts." A huge thank you to Daniel for sending in such a compelling prompt this week.

**Herman**

Indeed. And thank you, listeners, for joining us.

**Corn**

You can find "My Weird Prompts" on Spotify and wherever else you get your podcasts. Make sure to subscribe so you don't miss an episode!

**Herman**

Until next time, keep questioning the unexpected.

**Corn**

And stay weird!