

I Tried to Replace an Interior Designer with Gemini. Here's What Happened.

Published: October 8, 2025 at 5:46 PM Author: Richard Djarbeng 10 min read

When I finally got access to Gemini Pro, I had one main goal: to test its new image generation capabilities, which I'd heard referred to by the intriguing codename "Nano Banana." After opening the familiar chat interface in Chrome, I decided to skip the basic "draw me a cat" prompts. I wanted to give it a real, practical challenge to see what this new model could actually do. So, I grabbed a hand-drawn sketch of a room that was actually under construction, complete with specific, messy dimensions. The experiment: Could Gemini follow my instructions, understand a crude drawing, and iterate on a design like a real collaborator? This is a step-by-step log of what happened when I put the new "Nano Banana" model to the test.

Initial test

I decided to put it to the test. First I prompted it 'Hello, can you show me what you can do'. Gemini gave a pretty standard, comprehensive text-based answer, listing its abilities like answering questions, writing code, and solving math problems. Here's a snippet of the answer:

Hello! Of course. I'd be happy to show you what I can do.

I am a large language model, trained by Google. Think of me as a versatile tool for information, creativity, and productivity. My capabilities are broad, but here's a taste of what I can help you with:

1. Answer Your Questions (from simple to complex)

....

2. Create and Brainstorm

3. Help with Practical Tasks

4. Role-Playing and Scenarios

It was a good summary, but it was missing the one thing I was there for.

"Can you generate images?" I asked, "if yes why didn't you mention it?"

That's when it came clean. It confirmed it could generate images and apologized for the omission. Then it generated this:




I then asked it about the technology, specifically mentioning "Nano Banana," the codename I'd seen floating around. Gemini confirmed that this was indeed the internal name for the powerful model it uses. Now we were getting somewhere. From Gemini

*Here's the story: "**Nano Banana**" was the internal codename Google used for a powerful new image generation model while it was being tested. This codename became public and caught on, creating a lot of buzz online. The official name for this technology is **Gemini 2.5 Flash Image**.*

The Real Challenge: Designing a Room

To move beyond simple prompts, I uploaded a hand-drawn floor plan of a room with specific dimensions for the room itself, a bed, a table, a cupboard, and two doors. The task was simple: generate a design.

Picture 1 - Design my room



The first attempt came back. It was a nicely rendered photorealistic room, but it had a few issues. I'll add it below so you can see for yourself.



Picture 2 - Remove the TV


A couple of problems with this photo. It added a TV I never mentioned and had placed the doors on the same wall, this looks more like a hotel room. Not a great start. I corrected it.

There was no tv in the room and the doors were not on the same side of the wall



Picture 3- Fix the Measurements

The second version removed the TV and fixed the doors, but now the scale was off. The table looked larger than the cupboard, despite my dimensions clearly stating the cupboard was wider. Some text from the sketch in purple ink with an arrow was also somehow in the photo. I also pointed out that there were no lamps or pictures in my simple plan.




This attempt had a plainer slightly less visually appealing room that was better than the previous generations.

Picture 4- Bird's eye-view

The third attempt was closer, but it was clear that the 3D perspective was making it hard for the AI to get the layout just right. So, I switched tactics.

"Can you generate it from a bird's eye view from the top?"

This was the breakthrough.



Room dimension 3m x 4m

The top-down view was much more accurate.

Picture 5- Adding windows

Now we could really get into the details. The top-down view has a window close to the table, while this addition was not intentional on my part, I wanted to work with it since the room had a window anyway. I decided to add a window, specifying its exact size and location: centered on the wall directly opposite the entrance. The AI tried, but it put the window on the wrong wall.


My prompt:

Good, but you added a window. I would like the window to be on the wall opposite the entrance. The window is in the middle of that wall opposite the entrance and is 1.5 m wide

Gemini's output:

Thank you for the precise feedback! My apologies for that stray window. You're right, it needs to be on the wall opposite the entrance, centered, and 1.5m wide...

I corrected it again, and it moved the window... to the *other* wrong wall.




Room dimension 3m x 4m

It was a fascinating failure. It showed the AI was trying to follow instructions but was struggling with spatial relationships described in text. I decided to scrap the window idea altogether.

Picture 6- Removing windows

"Go back to the previous design before you added the window," I instructed.




Room dimension 3m x 4m

It did, but it also removed the bathroom door in the process! A small but crucial detail.

Picture 7- Bring back the bathroom

After one final correction to bring the door back, we had a perfect, minimalist top-down view of my original layout.




Room dimension 3m x 4m

The Final Test: A New Layout with Constraints

Now for the real test of its reasoning. I gave it a new set of rules: the cupboard and doors were fixed, but the bed and table could be moved. Crucially, I added that the entrance door swings *inward*. Could it create a functional new layout?

Assuming the cupboard position cannot be changed along with the position of the doors, but all other items can be moved, also the door at the entrance swings outward into the room can you. Suggest another possible layout?



Room dimension 3m x 4m

Its first attempt was good, but I spotted a flaw that only someone who had studied the dimensions would see.

"If you look at the dimensions for the door and the cupboard you realize that they take up most of the wall which means that space between the entrance door and the wall will be very small, not very large as shown in the current layout."

This was the final boss of corrections. It required the AI to not just place objects, but to understand the spatial implications of their dimensions on a fixed-size wall.


Gemini's output:

You are absolutely right to point that out! My apologies, I failed to accurately represent the tight space between the cupboard and the entrance door given their combined width and the wall length. The previous layout did not respect those precise dimensions, making the gap appear larger than it would be.

Let's fix that. The cupboard is 1.8m and the entrance door is 1.2m. On a 3m wall, that only leaves 0m ($1.8m + 1.2m = 3m$) if they are side-by-side, or very minimal space if the cupboard is offset slightly. This means the cupboard must start right next to the wall to accommodate the door.

I'll adjust the layout to reflect this tight fit accurately, while still allowing the entrance door to swing inward.

And it nailed it. The final image came back with the cupboard pushed right up against the corner, accurately reflecting the tight space left for the door and its inward swing.



Room dimension 3m x 4m

There were some flaws in the picture, I have to admit, but I considered them to be minor. They included problems with text and a semblance of a window on the correct wall highlighted it in blue. In spite of this, I consider the final image a decent attempt to position and design the room.

Then after playing around to a point where I was satisfied with its ability to generate images, I told it to write this website post.

Write me a website post in markdown about this interaction. It should go something like this... (rest of prompt)

(The post ends like this) Then after playing around to a point where I was satisfied with it's ability to generate images, I told it to write this website post.

Conclusion

This final portion is written by me, Richard. I hope you enjoyed testing out the capabilities of nano-banana in Gemini Pro with me as it designed this room layout. Did this layout come from an actual room to be designed? Yes, yes it did. When I was testing this the room was under construction. Happy to add that the final layout adopted by the occupant of the room was very similar to what was generated by nano-banana. Looks like AI wins this round.

I would like to figure out ways in which this can be extended one day. Also hope to review other models soon. See you in the next post.

PS: I have a confession. I lied a little in the post when I said the final output was this website post. After writing this post I realized I didn't have a cover image for this post so I gave Gemini another task to generate a cover image for this post.