




Your architectural prompt guide.



Prompt



What's in it for you?

Mastering AI prompt engineering turns Veras into an incredibly fast, highly capable digital studio assistant that brings your architectural visions to life.

This guide provides a clear blueprint for controlling AI output so you can consistently generate accurate, high-impact visuals.



Key takeaways.

- **Master the prompt anatomy:**
Learn how to structure prompts systematically.
- **Protect your architectural vision:**
Discover how to use Veras to protect structural geometry, keep realistic textures, and eliminate unwanted AI artifacts.
- **Achieve publication-quality realism:**
Move past generic renders by controlling design context.
- **Practical hands-on workflows:**
Follow practical exercises directly inside Veras AI.






Meet your new studio sidekick: **Veras AI.**

- The absolute best way to approach writing a prompt is to imagine you just hired a brand-new assistant for your studio. They are ridiculously fast, super eager to please, and hold a massive library of architectural history in their brain.

But there is a catch: they walked into the office exactly five minutes ago and know absolutely nothing about your current project.

- If you turn to this newbie and simply say, “Give me a cool modern house,” you are going to get something completely generic and uninspired.

But, if you sit down and hype them up with the real context—the crazy steep topography of the site, your obsession with board-formed concrete and warm timber, and the client’s absolute need for morning light in the kitchen—your assistant can actually crush the assignment.



The golden rule of prompting is simple:

The better you explain and the more context you drop, the more mind-blowing the output will be.

Prompt anatomy.




+



+



=



Task

This is the objective or the goal - what Veras AI is asked to do.

Attributes

Components, Style, Light & Color, Camera - equal-weight inputs.


Constraints

This is where you define what must stay true - affirmative and negative rules.

Generated image

Your vision comes to life as a final visual.


Prompt anatomy.



The task.

The Task is the very first step in writing a Veras prompt. It simply tells the AI exactly what you want it to do.

Think of it as a clear, straightforward project title rather than a detailed description.



Key ingredients.

- **Start with an action:**
Lead with a strong verb. Words like enhance, visualize, add or transform instantly tell the AI what operation it needs to perform.
- **Stick to one goal:**
Focus on a single objective. Cramming multiple requests into the first step can confuse the AI and give you mixed results.
- **Keep the tone neutral:**
Save the mood, style, and artistic flair for the later. Adding other details here will only water down your main instruction.





Prompt:

Emphasize the intricate foreground figure, house serving as a background. Rich tactile detail, close-up, creamy bokeh. Ensuring crystal-clear focus, the house should be softly out of focus.

Scene attributes | Components


The Components stage is where you describe exactly what exists in your image, where you actually build the scene—defining both the star of the show and everything surrounding it.



What to include.

- **The main subject:**
Clearly identify the "hero" of your image. The AI needs one unambiguous focal point to organize the composition around.
- **The environment:**
Give your subject a believable home—context brings realism.
- **Material details:**
Describe physical textures (e.g., wood grain, concrete pores, glass reflections). This is the secret sauce that turns an "AI" render into something real.
- **A sense of scale:**
Add things like people, cars, or furniture. Visual cues help the viewer (and the AI) understand the true size and proportions of your space.





Prompt:

Architecture photography of the building facade, showing intricate joinery. **Highly detailed, tactile material focus.** Natural afternoon lighting.



Scene attributes | Style

The Style is the part where you tell the AI exactly how the image should look and feel.

While the Components cover what is physically in the frame, the style dictates the overarching mood, aesthetic, and visual quality.



What to include.

- **Editorial references:**

Anchor your prompt to a known visual culture. Referencing a real publication (like **ArchDaily**, **Dezeen**, or **Architectural Digest**) or a specific photography genre gives the AI a concrete benchmark to aim for, rather than guessing what a vague word like "nice" means.

- **Quality benchmarks:**

Tell the AI what level of production you expect. Words like "ultra-detailed," "hyper-realistic" or "implementation-feasible" signal the level of realism and refinement you need.

- **Stylistic consistency:**

Make sure the whole image plays by the same visual rules. If left unguided, AI models might blend conflicting styles—like part illustration, part cinematic, and part documentary—resulting in a messy, confused image.





Prompt:

Transform the image to **ArchDaily style**.

Photorealistic, hyper-detailed, with realistic light bounce and shadows.



Scene attributes | Light & color

The Light & Color dictates exactly how light interacts with your scene, both physically and emotionally.

If the Components establish "what is there" and the Style sets the "visual language," this part answers: what is the light doing to everything?



What to include.

- **Time of day:**
This single choice dictates the sun's position and how shadows behave. "**Golden hour**" provides long, warm shadows, while midday gives you short, harsh ones. Getting this right is crucial for physical realism.
- **Quality of light:**
Define the softness, direction, and volume of the light (such as "**volumetric sunlight filtering through trees**"). This separates a beautifully lit render from one that is merely bright.
- **Interior vs. exterior balance:**
For buildings with windows, contrast a warm interior glow against a cooler exterior. This adds depth and makes the architecture feel occupied and alive rather than staged.





Prompt:

Enhance the provided render. **Natural afternoon lighting. Volumetric sunlight filtering through trees**, softly blurred background. Ensure a clear visual contrast.

Scene attributes | Camera

The Camera translates your photographic vision into AI-friendly instructions.

If the Components define what is in the scene, and Light & Color defines how it is lit, the Camera defines exactly how it is seen—the literal viewpoint from which your image is captured.



What to include.

- **Lens specifications:**

This determines your field of view. A wide-angle lens stretches depth and makes spaces feel larger and more dramatic, while a longer lens compresses distance and flattens the perspective. This single choice dictates whether a space feels vast or intimate. Instead of vague terms like "wide shot," name an actual focal length (e.g., **24mm**, **35mm**, or **50mm**). This gives the AI a concrete physical reference.


- **Camera position:**

Define the height and angle of the camera. An eye-level shot feels neutral and documentary; a low angle makes buildings look monumental and imposing; an aerial view emphasizes the layout and surrounding site.

- **Perspective control:**

Tell the AI to keep architectural lines correct. AI models frequently struggle with converging or "keystoning" vertical lines, so explicitly asking for perfectly straight verticals is crucial here.





Prompt:

Populate the scene with human figures walking outdoors. **Eye-level perspective, clear depth. Shot on 50mm lens.** Strictly preserve original architectural design geometry.



Constraints.

Constraints are your safety net. While every other part focuses on creative freedom and building your scene, this part is all about protection.

It tells the AI exactly what must remain true, locking down the crucial elements that cannot change, regardless of the style or lighting you chose earlier.



What to protect.

- **Structural integrity:**
Ensure the AI doesn't "creatively reinterpret" your design.
The goal is to show this specific building, exactly as intended.
- **Material authenticity:**
Block those fake, plastic-looking textures that immediately scream "AI-generated."
- **Geometric accuracy:**
Enforce straight lines, proper proportions, and zero architectural warping.
- **Artifact prevention:**
Eliminate common AI glitches, like melting geometry, nonsensical background objects, or people with extra limbs.



Because AI reads text much like a human, telling it only what to avoid can sometimes backfire (if I say "don't think of warped windows," the AI still has to think about warped windows).

To fix this, split your constraints into two parts:

- **The affirmative (The "Do's"):**
Tell the AI what it **must** actively preserve.
This gives it a concrete, positive target to hold onto.
- **The negative (The "Don'ts"):**
Tell the AI what it **must not** include to act as an exclusionary guardrail.





Prompt:

A blurred person walking in the space, clear human scale. Photorealistic, tactile material textures, realistic light bounce and shadows. **Strictly preserve the original design geometry and all interior elements. Don't add any extra elements to the project.**



Unscramble the prompt.

The architectural AI prompt on the right is jumbled up. The anatomy structure we learnt needs to be applied here.

Unscramble the prompt using the **DIY Sheet** in the next slide.

"Shot on a 35mm wide-angle lens, from a low, scale-emphasizing angle. Avoid plastic-looking materials, warped geometry, or lens distortion. Enhance this building design into a photorealistic architectural image, lit at golden hour, with warm interior light spilling through windows, rendered in a high-end, ArchDaily-style editorial photography look. Maintain the exact structural silhouette and geometric accuracy; set on a lush hillside, with wood grain and glass reflections visible."

DIY sheet.

Task:
What you want
the AI to do.



Attributes:
Key aspects
of the image.



Constraints:
Boundaries to preserve
quality and accuracy.



Components:
The elements and
details to include.



Style: The look
and visual tone.



Light & color:
Lighting, mood
and color feel.



Camera: Viewpoint
and
lens setup.



Affirmative: What
must be preserved.



Negative: What
must be avoided.



Sample prompt.

Task:
What you want
the AI to do.



Enhance this building design into a photorealistic architectural image,

Attributes:
Key aspects
of the image.



set on a lush hillside, with wood grain and glass reflections visible,



rendered in a high-end, ArchDaily-style editorial photography look,



lit at golden hour, with warm interior light spilling through windows,



shot on a 35mm wide-angle lens, from a low, scale-emphasizing angle.



Components:
The elements and
details to include.



Style: The look
and visual tone.



Light & color:
Lighting, mood
and color feel.



Camera: Viewpoint
and
lens setup.

Constraints:
Boundaries to preserve
quality and accuracy.



Maintain the exact structural silhouette and geometric accuracy;



avoid plastic-looking materials, warped geometry, or lens distortion.



Affirmative: What
must be preserved.



Negative: What
must be avoided.




Human scale | close-up

"Close-up of a person interacting with the architecture, clear human scale. Photorealistic, tactile material textures. Cinematic lighting, realistic light bounce and shadows. Shot on 85mm lens, shallow depth of field, bokeh, hyper-detailed. Strictly preserve original architectural design geometry."





Replace the keywords.




Use the base image above in Veras. Copy the following prompt in the prompt field. Replace the highlighted keywords to change the output and compare the results with the images on the right.


Replace the keywords.




"Close-up of a **person** interacting with the architecture, clear **human** scale. Photorealistic, tactile material textures. Cinematic lighting, realistic light bounce and shadows. Shot on 85mm lens, shallow depth of field, bokeh, hyper-detailed. Strictly preserve original architectural design geometry."



Person
Human



Dog
Animal



Robot
Humanoid

Resources



chaos.com



cgconnect.chaos.com



academy.chaos.com





chaos
Veras

chaos.com/veras