



FIELD GUIDE / 2026 EDITION

What AI Rendering **Still Can't** Do

A working prompt pack for the ten tools architects actually use — and the eight checks that decide whether AI output survives a client meeting.

[START HERE](#)

What this is, and what it is not

You have read the tool comparison. You have probably signed up for two or three of them. This guide covers the part that comes after: getting output you can actually put in front of a client, and knowing the exact point where AI stops being the right tool.

Three parts

01 The prompt pack. Six prompts that work, written the way architectural images are actually briefed — materials, light angle, lens. Copy them, swap the specifics.

02 The Handoff Test. Eight checks to run before AI output leaves your desk. Most failures are invisible until a client points at them.

03 The cost line. What each route costs, so the decision is arithmetic instead of instinct.

Where this comes from

Ten years of commissioned work, and our own project mix rather than an industry survey. Where we quote a number, it is a number from that work or a market range we are willing to stand behind in writing.

Where we do not have a defensible number, we say so instead of inventing one. You will see that twice in this document.

10+

YEARS RENDERING
VISUAL STORYTELLING

350+

BRANDS & AGENCIES
EMPOWERED GLOBALLY

95%

CLIENT
SATISFACTION RATE

40%+

OF OUR COMMISSIONS NOW
INCLUDE MOTION OR 360



Maverick Frame production render. Note what carries this frame: water, rock, vegetation. Remember that — it decides where AI helps you most.

PART 01 — THE PROMPT PACK

Anatomy of a prompt that does not look like stock

Most architectural AI output fails for one reason: the prompt describes a feeling instead of a building. Every prompt in this pack is built from the same six slots, in this order.

01 Subject and program	What it is and how big. "Two-storey private house", "120-room hotel lobby". Not "modern building".
02 Materials, named	The single highest-leverage slot. "Board-formed concrete, vertical larch cladding, slim black aluminium frames" beats any adjective you could add.
03 Light, with an angle	Time of day, direction, weather. "Overcast, no direct sun" is the most reliable lighting in AI rendering, and conveniently it is what most client presentations look like.
04 Camera	Height, focal length, perspective type, and the words <i>verticals parallel</i> . Without them AI converges your vertical lines and the image reads as amateur.
05 Context and landscape	Ground surface, planting, season. This is where AI is strongest — let it carry weight.
06 Reference and negatives	A photographic reference, not an art style. Then the exclusions.

Five phrases to delete from every prompt

- "8K, ultra-detailed, hyperrealistic" — produces plastic over-sharpening, not detail.
- "Stunning, beautiful, award-winning" — pulls straight toward the stock look everyone recognises.
- "Modern house" with no materials — you get the same Californian box as everyone else.
- A famous architect's name — a planner or a competition jury will recognise the transfer.
- "Sunset" with no sun angle — blown-out sky, no readable façade.

The negative string

Append to every exterior and interior prompt. Syntax varies by tool; the content does not.

--no text, watermark, signage, people looking at camera, fisheye, tilted verticals, converging verticals, oversaturated sky, lens flare, bokeh, extra floors, duplicated windows, floating furniture, warped door frames

Half of these are not aesthetic preferences. Duplicated windows and extra floors are the two errors clients catch fastest, because they can count.

PROMPT 01 & 02 — EXTERIOR

Exterior

01 Overcast daylight — the workhorse

Architectural photograph of a two-storey private house. Board-formed concrete base, vertical larch cladding above, floor-to-ceiling glazing with slim black aluminium frames, flat roof with a thin fascia. Overcast soft daylight, high thin cloud, no direct sun, ground damp after rain. Camera at 1.6 m eye level, 35 mm lens, two-point perspective, verticals parallel. Mature birch, wild grasses and a gravel path in the foreground. Photorealistic architectural photography, natural colour, no HDR halos. --ar 3:2

Why it works. Overcast light has no hard shadows, and hard shadows are where AI gives away that it does not understand the geometry it just drew. This is also the least glamorous prompt in the pack, which is exactly why its hit rate is the highest.

02 Dusk hero — higher reward, higher failure rate

Architectural photograph of the same house at blue hour, sun 8 degrees below the horizon. Warm interior lighting at low output spilling through the glazing, no interior fixture visible as a bright point. Deep blue gradient sky, faint cloud. Camera at 1.5 m, 35 mm lens, two-point perspective, verticals parallel. Wet gravel reflecting the interior glow. Photorealistic architectural photography, natural colour, no lens flare. --ar 3:2

Where it breaks. Dusk shots are the ones clients ask for and the ones AI ruins most often: window mullions dissolve into the glow, and interior light gets applied to rooms that do not exist in the plan. Run the Handoff Test on page 07 twice as carefully on any dusk image.



Maverick Frame production render. The car, the step geometry and the reflection in the glazing are the three elements AI most reliably gets wrong at this camera height.

PROMPT 03 – 06

Interior, sketch, context, material study

03 Interior daylight

Interior photograph of an open-plan living room, 3.2 m ceiling. Micro-cement floor, oak slatted feature wall, warm white plaster. North-facing full-height window, soft directional daylight from the left, no visible sun patch on the floor. Camera at 1.3 m, 24 mm lens, one-point perspective, verticals parallel. Low linen sofa, travertine coffee table, one olive tree in a clay pot. Photorealistic interior photography, natural colour, no wide-angle distortion. --ar 3:2

Note. The ceiling height and the camera height are doing real work. Leave them out and AI defaults to a cathedral ceiling shot from a child's eye level, which is why so much AI interior output feels subtly wrong without anyone being able to say why.

04 Sketch-to-render — for the tools that take your drawing

Keep the geometry exactly as drawn. Materials: dark standing-seam zinc roof, lime render walls, oak window reveals, basalt paving. Late afternoon sun from the left at 25 degrees, long soft shadows. Photorealistic architectural photography. Do not add or remove windows, do not change the roof pitch, do not add storeys.

Different job, different prompt. Here the sketch carries the geometry, so the prompt should be short and mostly prohibitions. Long descriptive prompts fight your own drawing and the drawing usually loses.

05 Massing and context

Aerial architectural photograph, 60 m altitude, 50 mm lens, looking down at 35 degrees. Four-storey mixed-use block, brick and precast concrete, flat roof with visible plant screening. Existing low-rise context on three sides, mature street trees, wet asphalt. Overcast light. Photorealistic, natural colour, no tilt-shift. --ar 16:9

06 Material study — one surface, many options

Close-up architectural detail photograph, 90 mm lens, 1 m from the surface. Corner junction of vertical charred timber cladding meeting board-formed concrete, 10 mm shadow gap between them. Overcast daylight, raking from the left. Photorealistic, natural colour, neutral white balance, no vignette. --ar 1:1

The most underused prompt in the pack. AI is genuinely good at surfaces and genuinely bad at buildings. Material studies play entirely to that strength — and a sheet of six material options is often more useful in a client meeting than one hero image.

MATCHING

Which prompt goes into which tool

The ten tools from the comparison split into three families. Feeding a prompt to the wrong family is the most common reason people conclude "AI rendering doesn't work".

FAMILY	TOOLS	HOW TO PROMPT IT
Geometry-led Your model or view is the input	Veras, D5 Render, Enscape, ArkoAI	Use Prompt 04 . Short, material-and-light only, heavy on prohibitions. Your model already holds the geometry — every descriptive sentence you add is a chance for the tool to overrule it.
Sketch-led A drawing is the input	Vizcom, ReRender AI, Gendo, LookX AI	Start from Prompt 04 , then add two lines of material detail. Raise the tool's own strength setting only after you have a version that respects the drawing, never before.
Text-led Only words are the input	Midjourney, ArchiVinci	Use Prompts 01, 02, 03, 05, 06 in full. All six slots matter here because nothing else is constraining the output. Expect to run twenty variations for one usable frame.

The one setting most people get wrong

Every geometry-led and sketch-led tool has a slider for how far it may depart from your input. It ships somewhere near the middle, where output looks impressive and is dimensionally useless. Start it low, make the image boring and correct, and raise it only until the first thing you can count changes.

On seeds and variation

If a tool exposes a seed, record it with the prompt. It is the only mechanism that lets you return to an image later and change one variable. Without it, a client asking for "the same shot but darker cladding" is a new lottery, not an edit — which is check 08 on the next page.

PART 02 — THE HANDOFF TEST

Eight checks before AI output leaves your desk

Run these in order. They take about four minutes. A client will run checks 02 and 03 without knowing they are doing it.

01 Verticals

Hold a straight edge against the image. Are vertical lines parallel, or do they converge toward the top? AI converges them almost by default. Nobody names this error — they just say the image "looks off".

02 Count the openings

Count windows and doors on each elevation against your plan. **AI adds and drops them silently.** This is the single most common error a client catches, because counting requires no expertise.

03 Floor-to-floor

Find a door in the image and use it as a 2.1 m ruler. Does the storey height read correctly against it? Drift here makes a four-storey building look like three and a half.

04 Material continuity at corners

Follow each material around the corner of the building. Does the cladding keep its direction? Does the brick bond survive the turn? AI treats each façade as a separate picture.

05 Reflections

Look at what the glazing reflects. It should be the context in front of it. AI routinely invents a different, more interesting building in the reflection — which is a problem when the real context is the reason the scheme looks the way it does.

06 The named product

If the brief specifies a real window system, a real tile, a client's own furniture or a brand's own colour: is it **that** one, or a convincing lookalike? A lookalike is where an image stops being a marketing problem and becomes a contractual one.

07 Resolution at final size

Scale to the size it will actually be used — A2 board, printed brochure, a 4 m trade-stand panel — and look at the edges at 100%. AI edges fall apart well before a rendered image does.

08 Can you change exactly one thing?

Ask for the same image with different cladding, and nothing else moved. If you cannot hold the rest fixed, you do not have an asset. You have one lucky frame, and the next round of client comments will cost you the whole image.

How to score it

One or two failures on checks 01–06 — fix it in post, you are fine. **Three or more, or any failure on 06 or 08** — the geometry is not there. What you are holding is a concept image, and presenting it as a deliverable is how the conversation with the client goes wrong two weeks later rather than today.



THE LIMITS

Four walls AI has not moved in 2026

"The neural net can just make the model now" — it can't. Modelling isn't there. The tools that promise a finished render straight from your drawings read construction drawings badly. AI is a genuinely useful assistant, and nothing like the populist marketing version.

DIM KUZMENKO, CO-FOUNDER, MAVERICK FRAME STUDIO

01 It does not produce a model	Every image is generated fresh. There is no geometry behind it to re-light, re-angle, or hand to anyone downstream. The moment a project needs a second view of the same building, the route changes.
02 It reads drawings badly	Drawing-to-render tools handle a clean single-storey massing sketch. Give them a real construction set — dimension strings, hatching, section marks, revision clouds — and accuracy collapses in ways that are hard to spot in the output.
03 Nature carries the frame	Our biggest AI wins come where greenery, water and planting do the visual work. On a bare urban block with a tight context and no landscape to hide behind, the gap between AI output and a modelled render is at its widest.
04 Revisions have no memory	A modelled scene takes a change and keeps everything else. A generated image takes a change and re-rolls the world. This is the wall that turns a cheap image into an expensive week, and it arrives after the client has already seen the first version.



Bare frame, no landscape to lean on — the hardest case for AI, and the easiest for a client to pick apart.

And one more, about motion

More than 40% of the work commissioned from us now includes animation, a 360 tour, VR or panoramas — instead of stills, or alongside them. That share keeps rising. AI will give you a frame. It will not give you a coherent walkthrough of *your* building, and that is increasingly what the client is actually asking for.

PART 03 — THE COST LINE

What each route actually costs

Most people reach for AI because a studio feels expensive and hiring feels permanent. Here are the three numbers side by side, so the decision stops being about feeling.

COMMISSIONED WORK	FREELANCER	STUDIO
Single exterior still	\$100 – 300	\$250 – 600
Single interior still	\$120 – 350	\$300 – 700
Package, 5–8 views	\$600 – 1,500	\$1,500 – 4,000
Animation, per minute	\$1,000 – 3,000	\$3,000 – 8,000

Market ranges we are willing to publish. We have deliberately left turnaround and revision-round figures out of this table: ours vary too much by project type to state as a benchmark, and a made-up number here would be worse than no number.

The in-house arithmetic

Bringing rendering in-house rarely pays off below a certain size, and the software licence is the smallest part of it. One artist needs:

- Around **ten different software subscriptions**
- A workstation that can actually render
- Modelling, lighting, post-production *and* client communication — rarely one person

It works at a headcount that keeps that person busy every day. Below that, outsourcing wins on arithmetic alone, before anyone argues about quality.

The mistake that costs the most

The biggest failure we see is not choosing the wrong tool. It is **timing**. Clients start looking for a vendor only after the first one has already failed and the deadline is on fire — and then choose in a panic.

Line up and vet rendering partners **before** you need one. One, ideally a few, briefed and ready to take work. It costs an afternoon now and saves a project later.



Maverick Frame production render. Fabric behaviour, curtain translucency and floor reflectance of the client's actual specified materials — the category in which check 06 is decided.

THE DECISION

Three routes, one question each

ROUTE A AI alone Concept and mood. Internal conversation. Early sketch options. Material studies. A first look at a massing idea before anyone has committed to it. Question: would it matter if a detail is wrong? Cost: an hour, a subscription.	ROUTE B AI plus manual cleanup Competition boards. First client conversation. Speculative pitch where you are selling a direction rather than a building. Question: does it pass the Handoff Test after you fix it? Cost: AI time plus the hours to repair it — usually underestimated by half.	ROUTE C Modelled CGI Planning submission. Sales and marketing material. Print. Anything with a named product, a contractual material, or more than one view of the same building. Question: will anyone make a decision or sign something because of this image? Cost: the table on page 09.
---	--	---

The short version

AI is now a genuinely useful assistant at the front of the process — the place where you are still deciding what the building is. It has not become the thing that produces what you hand over at the end. If an image will be counted, measured, printed, submitted or sold from, it needs geometry behind it. Everything before that point is fair game, and you should be using it.

If you go with Route C: how to brief it

Shorter is better, and this is the opposite of what most people assume. Four things:

- What you need — **how many images, and of what**
- The plans and elevations, as they are
- A few reference images for mood and light
- One clear line on what the images are *for*

Inflating a small job into a fifty-page specification that pins down every element backfires: taken literally it comes out ugly, and it misleads the people doing the work. Tell them the intent and let the team that does this daily handle the detail.

A note on this document

The prompts here are a starting grammar, not a magic formula. Tools change monthly. The six slots on page 03 and the eight checks on page 07 describe how architectural images are read, not how any one tool works, so they should outlast whichever tool you are using now.

If you run the Handoff Test on a real project and find a failure mode we have not listed, tell us. We update this guide, and we credit what we learn from.



MAVERICK
FRAME
STUDIO

When the image has to hold up, **talk to us**

We are a CGI studio: architectural and product rendering, 3D modelling, animation, 360 tours and VR. Ten years, 350+ brands and agencies, and the kind of work that still looks right when somebody counts the windows.

Send us the plans and a line on what the images are for. You get a scope and a price, and no obligation to use either.

START HERE

maverickframe.com/contacts

SEE THE WORK

maverickframe.com/gallery

WRITTEN BY

Dim Kuzmenko

Co-Founder



Maverick Frame production render. · AI Rendering Reality Check, 2026 edition. Cost ranges are market ranges as of 2026 and are not a quotation.

