

PILLAR

INSIGHTS

When AI image generators don't live up to their promises.

A field test across Freepik, Nano Banana Pro, Midjourney, and DALL-E — what these tools actually deliver, where the marketing breaks down, and why the ambiguity persists.

PILLAR MEDIA & ENTERTAINMENT · PILLARME.COM/INSIGHTS

01 — From DALL-E 2 to a working ecosystem: the 2022–2026 leap.

AI image and video generation services are now used in all sorts of occasions. Where graphic design and artistic vision and craftsmanship used to be a requirement, these services can now replicate the same results within a (mostly) reasonable degree of similarity.

However, this wasn't always the case. Back in 2022, services like **DALL-E 2** and **Midjourney v4** were the only ones available and they showed very underwhelming and inconsistent results. Today, these services are complete ecosystems featuring adjustable camera angles through flexible prompting, while being overall of much higher quality and with much more consistent results than ever before.

Like we just mentioned, these services can produce very similar results to work that is typically very labour intensive — perfect for marketing and commercial design, architecture and conceptual art, VFX, animation, and recreational uses (like AI slop reels).

However, not all these services live up to the way that they present themselves. And those that do still show a few inconsistencies in their generations that, for now, only the human eye can identify.

The pattern is worth stating plainly. Four years ago, a generated image was an obvious novelty — misshapen hands, melted text, lighting that made no physical sense. Today, a generated image can pass for a stock photo in three out of four glances. The remaining glance is the one that matters: the human reviewer who notices that the highway traffic is going the wrong way, that the storefront sign reads as nonsense letters, that the shadow falls on the wrong side of the figure. The visible glitches have moved from the surface of the rendering down into the world model behind it.

02 — Freepik's deliberately vague positioning.

Starting with **Freepik**. This service is ambiguous regarding what it can actually do according to its own description, leading some to believe that it can generate videos from prompts (which it cannot do in the free version).

According to their own LinkedIn page, Freepik (now **Magnific**) claims to be able to bring together “image generation, video creation, upscaling, and a library of 250M+ assets. All designed for people who need control, speed, and quality in the same place.”

This description is quite hard to visualize mentally and that seems to be the intention here. Did they say “video creation” instead of “video generation” in order to avoid repeating the word “generation” twice, or does this mean that “creation” and “generation” have two different meanings here? And what is upscaling in this context? And who doesn’t need control, speed, and quality all in the same place? It seems like that description includes everyone.

It is precisely this vagueness that leads people like me to give this service a try when in reality it does not offer the services that I seek.

03 – A field test: the “AI excursions” advert.

In my own experience, I was attempting to generate a video that was supposed to be a 10-second-long advertisement aimed at high school management, staff, and student body. The goal was to promote these hypothetical “AI excursions” in which a hypothetical company would sponsor school excursions for these high schools to learn about AI. The end goal is to show students how they can use AI as a vehicle that will concretize their creativity and ingenuity into something fascinating that can be shared with the whole world.

When using the free version to try to generate this advert, I was relieved when I saw the textbox for writing up a prompt. Little did I know that it would be completely useless to me, since what I was trying to do was completely outside of the scope of what the free version allowed.

Magnific (formerly Freepik)

Upscale Conf SF June 3-4

Servicios de diseño · 97 mil seguidores · 501-1 mil empleados

+ Seguir

Enviar mensaje



Inicio

Acerca de

Publicaciones

Empleos

Personas

Resumen

Magnific is the pro platform for creative visionaries who set the bar high.

We build AI-powered creative tools for professionals, studios, and brands making their best work across every medium. Our AI Suite brings together image generation, video creation, upscaling, and a library of 250M+ assets. All designed for people who need control, speed, and quality in the same place.

This is where the thinking behind the work lives. Tutorials, deep-dives into real workflows, and honest conversations about what it actually takes to push creative work further with AI. No hype. No hand-waving. Just the process inside the process.

You might know us as Freepik. Magnific is what a decade of building for the creative industry looks like when you stop holding back. 110M+ users. 1M subscribers. A global team obsessed with reducing friction and increasing flow for the people who make things for a living.

We're not here to show you what AI can do. We're here to show you what you can do with it when you push further.

The Freepik free-tier video flow surfaces a prompt textbox first, then requires a reference image — the prompt itself does almost none of the generative work.

Once I was done writing up this detailed script of how I wanted everything to be, I found out that I first needed to provide a reference image. After a quick search, I landed on an image depicting a crowd attending a TED talk. Eager to finally be able to generate my advert, I clicked on the "Generate" icon and waited approximately one minute for the image to generate.

To my dismay, all it did was animate the one frame whilst also lowering the image quality. In essence, all it did was give the people in the crowd of the talk a bit of movement, and that's pretty much it.

Things are even worse off now, since you can no longer generate any videos without having a paid subscription plan. Even the underwhelming video I generated is no longer a thing that can be done.

04 — Nano Banana Pro: honest scope, sharper results.

Looking at **Nano Banana Pro** (Gemini), this service cannot generate videos — however, it never claimed to. It is great for image generation based on prompts and has all sorts of options.

In my experience with this service, all I did was copy-paste the exact same prompt I had given Freepik, and the results were vastly different. This time I was only generating an image, and in the space of **30 seconds** it yielded pretty much exactly what I had in mind.



Nano Banana Pro returned the requested scene in roughly 30 seconds. The composition reads as photographic on first glance — but a human reviewer can still spot real-world errors the model missed.

However, there were some inconsistencies that AI couldn't pick up but humans can, such as a car driving in the wrong direction and into incoming traffic on a highway. There isn't much else to be said about Gemini other than that.

05 — The structural question: why does the messaging stay vague?

To recap, not all services provide the results that they claim to yield. Services like Freepik deliberately use ambiguous messaging to make it seem like they can do anything, when in reality the offering is very limited — especially if you do not have a paid version. In addition, such services occasionally produce results that are inconsistent with reality despite the high image quality and accuracy in relation to the prompt.

This now leads us to ask the following questions:

- Why do the companies that operate these services deliberately use ambiguous and super-generalized messaging?
- What are the incentive structures that have led us to this current state?
- Will these services eventually overcome these minor inconsistencies, just like they overcame the initial hurdles in 2022?

The honest answer is that vague positioning widens the top of the funnel. Generalized claims like “control, speed, and quality in the same place” describe every buyer’s stated needs. Precision would deter many of the trial signups that fund growth. Until the incentive flips — until buyers reward clarity at the top of the funnel — the ambiguity is the product strategy, not an oversight.

There is a second factor worth naming. The capabilities of these models are moving fast enough that any specific claim made today may be true in three months even if it isn’t true today. The marketing copy is written for a roadmap, not the current build. That makes the ambiguity feel almost rational from inside the company — until you sit in the chair of the user who pulled out their credit card on the strength of a sentence the product can’t actually deliver.

ABOUT THE AUTHOR

Guillem Clibbon

Guillem Clibbon is a contributing writer following the practical edges of generative AI tooling — what these services advertise, what they actually deliver, and where the gap is widening or closing. This is a LATAM Ambassador contribution to [Pillar Insights \(/insights\)](#).

Frequently asked questions.

How much have AI image and video generators improved since 2022?

Substantially. In 2022, the only credible options were DALL-E 2 and Midjourney v4, and both produced underwhelming and inconsistent results. Today's services are complete ecosystems featuring flexible prompting, adjustable camera angles, much higher image quality, and far more consistent results. They can plausibly replicate work in marketing and commercial design, architecture and conceptual art, VFX, animation, and recreational content. The gap between professional output and generative output has closed enough that human craftsmanship is no longer a strict requirement for many production tasks, though humans still catch errors machines miss.

Does Freepik actually generate video from a text prompt?

Not in the way the marketing implies. Freepik (now Magnific) describes its product as bringing together image generation, video creation, upscaling, and a 250M+ asset library. In practice, the free version cannot generate video from a prompt at all. It requires a reference image, and the output is essentially an animation of the single frame you provided with reduced quality. Video generation from prompts now requires a paid subscription. The deliberately ambiguous distinction between "video generation" and "video creation" appears designed to widen the funnel.

How does Nano Banana Pro (Gemini) compare to Freepik for image work?

Nano Banana Pro is strictly an image generator and never claimed to produce video. Given the same prompt that failed on Freepik, Nano Banana Pro returned a usable result in roughly 30 seconds that closely matched the requested scene. The honest scoping of capability and the prompt fidelity are its strongest features. The caveat is that even high-quality generated images still contain real-world inconsistencies that only a human reviewer can catch, such as vehicles driving the wrong way into oncoming traffic.

Why do AI generation services use deliberately vague marketing?

Ambiguity widens the addressable funnel. A description that mentions image generation, video creation, upscaling, and a massive asset library covers nearly every visual workflow a buyer might imagine. Generalized claims like “control, speed, and quality in the same place” describe everyone’s stated needs. The vagueness draws curious users into the free tier, where they discover the real product is narrower or paywalled. The incentive structure rewards top-of-funnel breadth over precise capability disclosure, because precision would deter many of the trial signups that fund growth.

Where does AI image generation still fail in ways humans notice?

In contextual coherence. A modern image generator can produce a photorealistic highway scene that looks correct on first glance but contains a car driving the wrong direction into oncoming traffic. The model gets surface texture, lighting, and composition right while missing a real-world rule that any human driver would catch instantly. Similar errors appear in physics, anatomy, signage, and culturally specific details. Until generators reliably enforce world-model constraints, every commercial output still needs a human reviewer.

Will AI generators eventually overcome these inconsistencies?

The 2022 to 2026 trajectory suggests yes. The leap from DALL-E 2 and Midjourney v4 to today’s tools cleared hurdles that experts said would persist for years. The remaining issues are world-model issues rather than rendering issues, which means progress depends on different research bets than the ones that drove the last wave. Expect the visible glitches to shrink each model generation. The harder question is whether the marketing will catch up with reality or whether ambiguity will remain a deliberate growth lever.

Premium addresses for AI-native brands.

Pillar manages a deep inventory of premium domains built for the generative-AI era. Whether you're launching a new image tool, a video model, or an honest alternative to the incumbents, the right address makes the marketing claim easier to back up.