

Product and Packshot Fidelity Playbook

Key principles for keeping product fidelity when crafting image ads with AI.



Contents

1. Model and resolution
2. Lock the setting in first
3. Reference images
4. Re-attach at each iteration
5. Fine print
6. Add a reference for each product
7. Multiple products: two approaches
8. Describe Physics, not Vibes
9. Where you can do it in Pencil



MYTHOPOEIC




Product Fidelity

With AI in Pencil

The product is the source of truth.

- * Colours, labels, typography, proportions, surface detail. When generating images that include real products, everything needs to stay consistent and on-brand. This guide outlines the key principles to bear in mind when you need to ensure product fidelity when crafting ads with AI.



Resolution and model

Generate at 4K from the get go.
Use the best available model.

Resolution sets the ceiling for everything downstream. Generate at **4K** from the start, not as a fix afterwards.

This means your image should be at least 4000 pixels long, at its longest edge. A square image at 4k would be 4096 x 4096, for example.

content. Example Use Case: "Swap the can in this lifestyle photo with our new flavor design while keeping the hand, lighting, and background unchanged."

Select model



Best Available

Nano Banana 2: Gemini 3.1 Flash Image

Aspect ratio ⓘ

1:1 4096 x 4096

The best models for product placement shots and on-location packshots are currently **GPT Image 1.5**, **Nano Banana Pro** and **Nano Banana 2**, all of which are optimised for product fidelity.



A close-up photograph of a falcon perched on a leather-gloved hand. The falcon has dark feathers with light-colored spots and a sharp beak. The hand is wearing a brown leather glove with a silver chainmail bracelet and a blue gemstone. In the background, a person in a fur-lined hood is blurred. The word "MYTHOLOGIC" is written in white capital letters across the center of the image.

MYTHOLOGIC


Scene first. Product last

Lock the scene in, before you introduce the product.

If you're generating the background or environment as well, get that working first. Get your composition, lighting, textures as you want them.

Then attach the product once the scene is stable.

Why? Every time you regenerate the full image to tweak a background detail, the model has an opportunity to drift from the product reference. Fewer regenerations with the product present means higher fidelity.



Product placement example

References:



Reference 1: Model in location

Reference 2: The product

Model:

Nano Banana 2

Product placement example

Prompt:

"The model in attachment 1 is displaying the necklace from attachment 2, clearly showing it to the viewer."

Model:

Nano Banana 2



Reference images

Provide a clean product shot and ask the agent or model to reference it directly.

Upload a high-resolution image of the actual product into the chat. Explicitly tell the agent to use it as a reference.

Clean means: isolated product, plain matte background, no motion blur, good lighting, no glare.



PROMPT: "The woman holds out her hand, palm facing inward, back of hand facing outward toward the viewer, holding up her fingers to show the rings she is wearing (please see product reference image of the stack rings attached). She wears the three rings stacked on one finger."



MYTHOPOEIC

Re-attach, re-reference and specify

When iterating, it can help to re-attach your product reference to avoid product drift.

The model doesn't carry a perfect memory of your reference image across multiple chat turns. Re-attach your source of truth - the original product reference - regularly to refresh the visual anchor, and prevent gradual product drift.

Tip: When creating ads with the latest image models in **Universal Chat**, you can mention the exact pixel ratio or format at the end of your prompt (e.g. 1920 x 1080), and the model will try to render in that format, if possible. This can help prevent the AI from making assumptions that lead to cropping, distortion, or unwanted composition choices, when iterating on a concept.





Small text

Very small text won't be perfect

Fine print, ingredients lists, barcodes, and regulatory text are a known limitation in today's image generation models, even at higher resolutions.

Bold headline text survives 4K well; but dense label copy, for example, often doesn't get preserved.

Some level of post-production is normal in AI-augmented workflows. If you need legible fine print, compositing that element separately in post-production may be the most reliable method, especially if the text is very small.



Detailed product placement example prompt:

"Museum photography. A South African rounded stylised African art statue found object - 'objet trouve' - sculpture figure - carved dark wood, roughly 40cm tall, The Mythopoeic hammered copper cuff bracelet [product reference attached] is placed on the figure's left wrist -. The cuff's oxidised silver chain wraps sit among the metalwork and cloth, the labradorite teardrop resting against the dark wood. The lighting is museum-grade - even, slightly warm, the kind of diffused overhead illumination used in ethnographic collections.

But the focus pull tells a different story: the cuff is tack-sharp, the figure's face and body are slightly soft, going just out of focus. The jewellery is the subject; the figure is the context. Shallow depth, 85mm, f/2.8. The background is museum-matte, off-white amber-beige- The image sits exactly at the intersection of commodity and ceremony, and refuses to tell you which one it is."

Model:

Nano Banana 2

Add a separate reference for each product

Each product may need its own reference image.

A combined shot or packshot collage can sometimes work, but it runs the risk of spreading the model's attention too thin.

It can work, but if not, try uploading each product separately as a dedicated packshot reference image.

You can then refer to them individually, and explicitly describe the placement for each product in your prompt.

This method eliminates a lot of surface area for ambiguity, and gives the model the best chance to correctly interpret your 'vision' for the ad.



Two approaches to adding multiple packshots

All at once, or one at a time. Each has pros and cons.

All at once:

Attach all product references in a single prompt and generate the full scene in one go.

- This way is quicker, and composition feels more natural and unified.
- But the model's attention is split, so per-product fidelity may drop.
- Works best for simpler compositions: flat-lays, white backgrounds, evenly weighted products.




All-at-once product placement example prompt:

The jewelry pieces attached presented on luxurious rich deep cyan satin or velvet bed. Studio shot; key lighting - even and clear.

Generate at 2K, 19:6 AR."

Model:

Nano Banana
Pro




Two approaches to adding multiple packshots

All at once, or one at a time. Each has pros and cons.

One at a time:

Place the first product, then feed that output back in and add the next. This method, while slower, offers more fine-grained control.

- More control over placement and fidelity at each stage.
- But each generation re-renders the full image, which can degrade earlier products.
- Re-attach all previous product references with each new generation.
- Works best for complex scenes with specific depth, lighting, or layering.



One-at-a-time product placement example prompt:


1: "She holds out her hand, palm facing inward, back of hand facing outward toward the viewer, holding up her fingers to show the rings she is wearing (reference shot attached). She wears the three rings stacked all on one finger."

2: "Now she raises her hand to show off the bracelet (reference attached) she is modeling and wearing around her wrist, in addition to the trio of stack rings."

3: "Zoom out slightly to reveal that she also wears the labradorite pendant necklace - the pendant of which is quite small - around 2 - 3 cm long (reference shot attached)."

Model:

Nano Banana 2



Describe Physics, not Vibes

Today's models can parse material specificity.

The more precisely you describe a material, the more convincingly the model can render it.



For high-stakes client campaigns - product shots, packshots, and beauty editorial work - a physics-driven prompt structure gives you fine-grained control and repeatable results. This methodology offers a robust, modular framework for achieving a granular level of creative control.

When you write "soft foam" or "nice lighting," you're describing an outcome, not a behaviour. The model fills in the gaps with its own interpretation, and that interpretation drifts from your product. Instead, describe what materials do. ***“Foam compressing around a bottle.” “Surface tension creating collapse pockets.” “Light falling off directionally across a label.”*** Behaviour language constrains the model. Aesthetic language should come only after the fundamentals of your shot are lockin in.

Prompting is still an iterative, experimental process, and elsewhere in this guide we encourage you to try different workflow approaches to see what lands. Physics-first language doesn't remove that flexibility; The idea here is to use precise prompting to give each generation a stronger starting position.

Describe Physics, not Vibes

Today's models can parse material specificity.

The more precisely you describe a material, the more convincingly the model can render it.

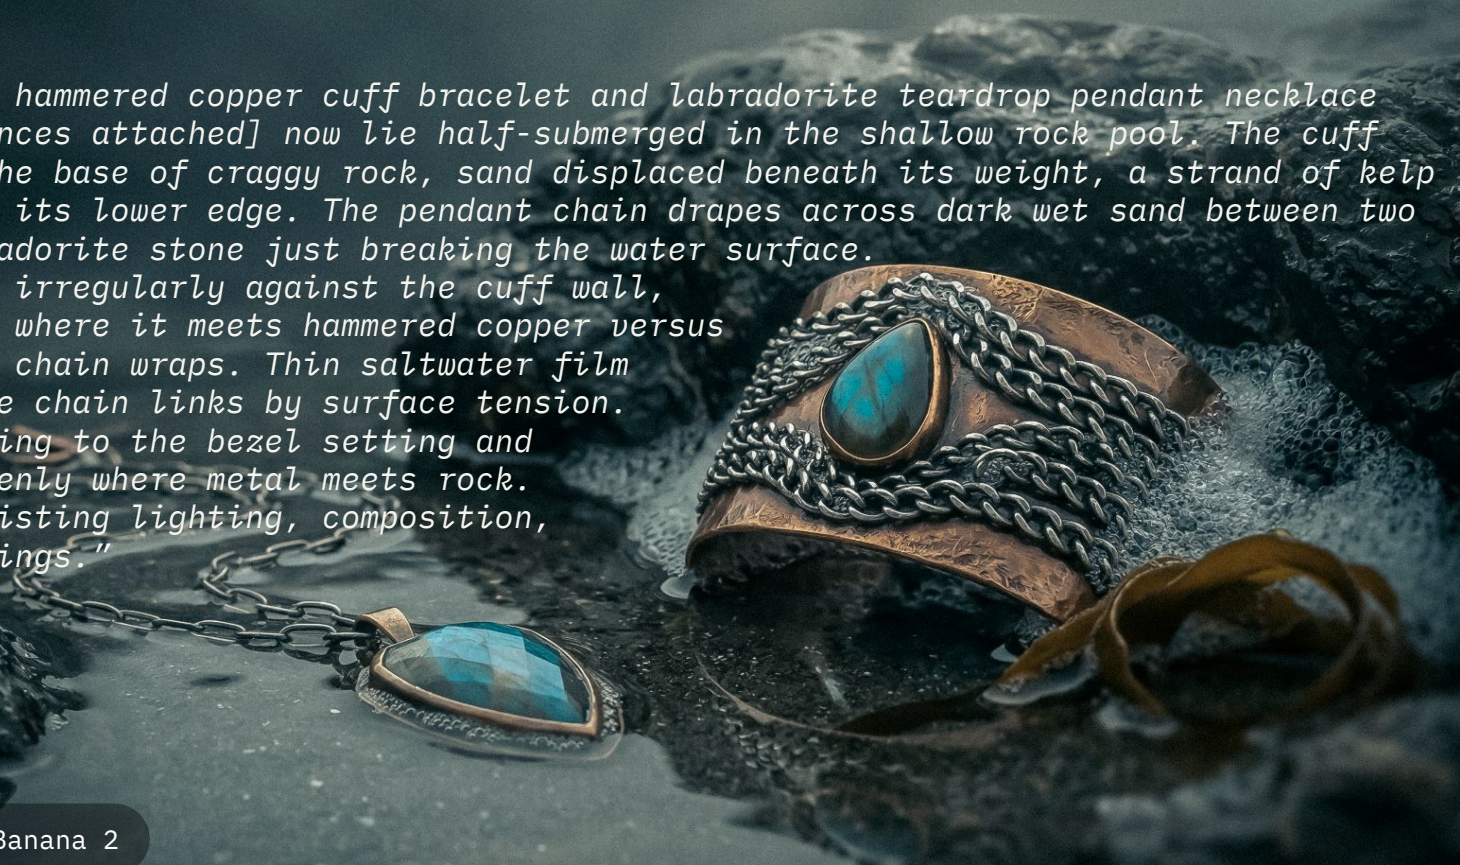


- 1 **Camera + geometry first:** Lock your framing, orientation, and lens geometry up front.
- 2 **Material physics + object behaviour:** Describe what materials *do*, not just what they look like. Physics language - such as ***tension variation, collapse pockets, microbubble fields, thin-film interference***, works well here. ***"Bubbles with Burst membrane traces and uneven collapse" renders more convincingly than "delicate bubbles."***
- 3 **Lighting:** Specify the rig, not just the vibes. Directional key, cast shadows, highlight falloff, controlled fill, these translate into dimensional, commercially credible output.
- 4 **Aesthetics comes last:** Once structure, physics, and lighting are locked down, aesthetic language has a resilient context to land in.



Physics-based product placement example prompt:

"The Mythopoeic hammered copper cuff bracelet and labradorite teardrop pendant necklace [product references attached] now lie half-submerged in the shallow rock pool. The cuff rests against the base of craggy rock, sand displaced beneath its weight, a strand of kelp curling against its lower edge. The pendant chain drapes across dark wet sand between two rocks, the labradorite stone just breaking the water surface. Foam compresses irregularly against the cuff wall, density varying where it meets hammered copper versus oxidised silver chain wraps. Thin saltwater film pulls across the chain links by surface tension. Microbubbles cling to the bezel setting and accumulate unevenly where metal meets rock. Maintain all existing lighting, composition, and camera settings."





Where you can do this in Pencil

These methods can be applied in various places across the platform.

The methods described in this guide are workflow-agnostic. We're always working on cutting edge new ways to create with AI, so this list of places where you can apply product fidelity tips will grow.

- Universal Chat
- Custom Agents
- Workflows
- Image Editor
- Ads Editor
- Sheets

As image models improve across the board, common product placement and fidelity issues will be ironed out at the model provider level. This means, over time, it'll get easier to achieve product fidelity in your ad creatives, without needing any fancy tricks or hacks, in whichever paradigm you prefer to work.





MYTHOPOEIC
WORN MYTHOLOGY.

* **Made with AI.**

No reptiles were harmed during the making of this playbook.

Pencil & the
brandtech
group