AI Tool Bible

Which image and video model to use, for which job

CM Productions · Tested across close to 100 productions · October 2026

Why this guide exists

Every week there's a new model and another claim that it's the best. After close to 100 productions in the past year (short films, spec ads, music videos and social content, plus brand work ranging from UGC-style videos to explainer videos, internal training videos and cinematic ads), this is my honest map of which tools I actually reach for, what each one is good at, where it falls down, and what it costs. Every example image in this guide comes from a real project.

How to read it

- Image models chart: strengths, weaknesses, max resolution and price for every image tool I use.
- Video models tier list: from best quality down, with price falling alongside.
- Grok: why the subscription is the best-value tool on the market, and its trade-offs.
- Pipelines in practice: real projects showing how the tools chain together, from cheap exploration to production-ready references.
- Head-to-heads and known issues: the same prompt through different tools, plus common failures and their fixes.
- Cheat sheet: one page, task → my pick → alternative.

Reading the example images

Each image is labelled with the tool that made it and a catalogue number. Tags show GOOD (a strength), MIXED (usable with caveats), BETTER (improved, but not perfect yet) or BAD (a failure worth learning from).

> **Important:** These are guidelines, not hard rules. They're my opinions from hands-on production use. Every model can surprise you in either direction, and results vary from generation to generation. Pricing changes often, so always check the platform.

Image models

Pricing basis: Higgsfield Ultra plan, $129 USD/month for 3,000 credits, so 1 credit ≈ $0.043 USD before tax. Other plans and top-ups cost more or less per credit. Tools not on Higgsfield use their own platform pricing. To convert USD to CAD, multiply by ~1.40 (rate as of October 31, 2026; check the current rate).

Tool	Best for	Watch out for	Max res	Credits	USD
Early discovery (cheap exploration, needs refinement)					
Higgsfield Soul Cinema Higgsfield	Cheap original character design and casting; strong mood and film look	Soft, filmic output, not start-frame quality; refine in a higher-end tool	**4K**	0.125 (2K)	~$0.005 (2K) 4K: not listed
Higgsfield Cinema Locations Higgsfield	Cheap location exploration: layout, mood, composition	Soft textures and low fine detail; refine before production use	**2K**	0.125	~$0.005
Midjourney Standalone	Highly stylized characters; uniquely haunting locations; bending physics and reality. Very low cost per image for exploration	Output generally not usable as-is; refine in a higher-end tool to reach production quality	**2K (HD)**	n/a	~$0.005–0.02 / image*
Production / refinement models					
Nano Banana Pro Higgsfield	4K photorealism; creatures; prompt-only creature work; patterned wardrobe; refining Cinema Locations sets	Often pushes toward a slightly animated look, especially when refining a photoreal reference, which can land in the uncanny valley	**4K**	1K: 2 2K: 2 4K: 4	1K: $0.09 2K: $0.09 4K: $0.17
Nano Banana 2 Higgsfield	Strong creature and alien imagery; highly stylized photoreal looks (e.g. graphic noir) at good value	Older model; less suited to straight naturalism	**4K**	1K: 1.5 2K: 2 4K: 3	1K: $0.06 2K: $0.09 4K: $0.13
Nano Banana 2.1 Higgsfield	Photorealistic images generated from scratch (text-to-image)	Weak as a refinement tool; referencing other images often gives strange results	**4K**	1K: 1.5 2K: 2 4K: 3	1K: $0.06 2K: $0.09 4K: $0.13
GPT Image 2 Higgsfield	Top-tier photorealism and close-ups; accurate text; AI compositing; product sheets	Most expensive image model; noise and fractalization on fine textures; small text can break down	**4K**	1K: 4.5 2K: 8.5 4K: 14	1K: $0.19 2K: $0.37 4K: $0.60
GPT Image 2.5 Sunburst Higgsfield	Gritty, found-footage realism; strong text; complex character sheets at lower cost	Noise and fractalization; may need a refine pass	**4K**	1K: 2 2K: 3 4K: 4.5	1K: $0.09 2K: $0.13 4K: $0.19
GPT Image 2.5 Flare Higgsfield	Gritty handheld and night-op realism	Same noise risk as the GPT family	**4K**	1K: 2 2K: 3 4K: 4.5	1K: $0.09 2K: $0.13 4K: $0.19
Seedream 4.5 Higgsfield	Extreme value for native 4K (4K costs the same as 2K); cheapest route to 4K output	Older model with weaker quality than 5.0 Pro; not recommended for new images	**4K**	2K: 1 4K: 1	2K: $0.04 4K: $0.04
Seedream 5.0 Pro Higgsfield	Best cost vs quality; creature realism; hero shots and turnarounds; compositing; cleaning up GPT noise	2K max, so upscale for 4K video	**2K**	1K: 1.5 1.5K: 1.5 2K: 3	1K: $0.06 1.5K: $0.06 2K: $0.13
Subscription value					
Grok Imagine 2.0 (images) Higgsfield / Grok app	Least restrictive content moderation on the market (mature themes, violence, nudity) for images made inside Grok. See the Grok section	Referencing any image made outside Grok triggers much stricter moderation; pay-per-generation is poor value vs. the subscription	**2K**	1K: 2 2K: 2.5	1K: $0.09 2K: $0.11**

Soul Cinema / Cinema Locations: some Higgsfield plans may include a large number of free generations. Check the platform for current offers.

GPT Image models are currently on sale on Higgsfield. Regular pricing is shown, so check the platform for current prices.

***Midjourney:** plans run $10–$120/month and buy Fast GPU time rather than individual images. Per-image cost is estimated from plan prices (4 images per prompt). Standard ($30) and higher plans include unlimited Relax-mode generations.

****Grok:** Higgsfield per-generation pricing is shown for comparison only. Paying per generation isn't recommended. The value is in a Grok subscription (see the Grok section).

Video models

Tiered from best quality down, with price falling alongside. Prices are per 15-second generation unless noted, using the same Higgsfield Ultra basis (1 credit ≈ $0.043 USD).

Model	15s price	30s price	Res / length	References	Verdict
Tier 1: Premium quality					
Seedance 2.5	480p: 45 cr · $1.93 720p: 105 cr · $4.51 1080p: 240 cr · $10.32	480p: 90 cr · $3.87 720p: 210 cr · $9.03 1080p: 480 cr · $20.64	**1080p*** 30s	Up to 50 (≈30 practical max)	Hands down the best video model for quality, and the most expensive. Strongest cinematic intelligence. Draft mode (480p, 45 cr/15s): one-click regenerate at 1080p. Not an upscale, so expect small changes in framing, camera or action. Reportedly unavailable to US users on Higgsfield.
Seedance 2.0	480p: 45 cr · $1.93 720p: 68 cr · $2.92 1080p: 180 cr · $7.74 4K: 330 cr · $14.19	—	**Native 4K** 15s	9 images + 3 videos + 3 audio (12 max)	Best option for native 4K. On 15s clips it's often a toss-up with 2.5, and frequently just as good at 1080p or lower for 25% less.
Tier 2: Budget					
Kling 4.0	720p: 41.25 cr · $1.77 1080p: 60 cr · $2.58	720p: 82.5 cr · $3.55 1080p: 120 cr · $5.16	**1080p** 30s	Up to 10 keyframes	Huge value vs Seedance but a budget option: almost always waxy, plasticky skin and far weaker cinematic knowledge.
Kling 3.0	720p: 30 cr · $1.29 1080p: 37.5 cr · $1.61 4K: 90 cr · $3.87	—	**4K** 15s	—	Many creators swear by its character consistency; mixed results in my testing. Cheapest route to 4K video.
Kling 3.0 Omni	720p: 18.75 cr · $0.81 1080p: 30 cr · $1.29 4K: 90 cr · $3.87	—	**4K** 15s	Multiple (handles better than 3.0)	Cheapest Kling. Handles multiple references better than Kling 3.0.
MiniMax H3	2K: 30 cr · $1.29	—	**2K** 15s	—	Quality roughly on par with Kling 3.0. One flat price on Higgsfield, native 2K: the best resolution per credit on this list.
Subscription value					
Grok Imagine 1.5 (Higgsfield)	480p: 37.5 cr · $1.61 720p: 67.5 cr · $2.90 1080p: 120 cr · $5.16	—	**1080p** 15s	—	Shown for comparison only: not worth paying per generation. The value is a Grok subscription (see Grok section).

***Seedance 2.5:** native 4K isn't on Higgsfield yet but is reportedly coming soon. Draft mode is reportedly unavailable to US users on Higgsfield but still works in Canada; availability may vary by region.

Seedance 2.5 references: up to 50 on Higgsfield (some platforms cap it lower). There's almost never a reason to use all 50; around 30 is the practical maximum.

Kling: kept for completeness. Many creators rely on it, but I've had very mixed results and have retired Kling from my own work.

> **Resolution tip:** Match your start frames to your video resolution. For 720p or 1080p video, 2K images are the sweet spot. If you're generating video in 4K, you need 4K start frames, either native 4K (Nano Banana Pro, GPT Image, Seedream 4.5) or a 2K image upscaled to 4K.

Grok: the subscription play

Grok is the one tool I recommend using through its own subscription instead of paying per generation. For value per second of video, nothing else on the market comes close. While Grok provides incredible value, it's not a dedicated image and video model and should be treated as such. I wouldn't rely on it for professional finished work, but for social media projects, prompt testing and building reference videos for Seedance, it's packed full of value.

Where it shines

- Speed: one of the fastest video models I use day to day. A 15-second generation rarely takes more than about 90 seconds, far faster than most hosted models.
- Animation: very high-quality 720p generations, especially for animated styles.
- Cheap exploration: test ideas, blocking and timing without burning premium credits.
- Video references for Seedance: build a rough version in Grok, then use those clips as motion and camera references for a final Seedance pass.
- 720p social content: often good enough to publish directly.
- Native 1080p video is available, but only on the Super or Ultra Grok plans. Plus and Premium are limited to 720p.

Content moderation

- Grok has the least restrictive content moderation on the market, but only for content created inside Grok.
- As soon as you reference an image made outside Grok (including anything in that image's reference chain), the generation falls under a much stricter set of rules. For mature work, keep the whole chain Grok-native.
- The only real alternatives for mature content are open models run locally with no moderation, such as MiniMax H3 or LTX 2.5.

Trade-offs

- Grok isn't a dedicated image or video model; it's an LLM generating images and video, so expect some trade-offs in control and consistency.
- Grok has shown it can produce 20-second clips, but Grok Imagine 2.0 is currently capped at 15 seconds. Hopefully that's upgraded soon.
- Speed caveat: LTX 2.5 is faster on paper. Its vendor benchmark is a 10-second clip in 6.8 seconds, self-hosted on two NVIDIA GB200 data-centre GPUs, or about 24 seconds through its API at 1080p. Local speeds on consumer GPUs will be slower than the headline, and I haven't tested LTX myself.
- Subscription pricing varies by region, so check the Grok app for current plans.

> **Responsible use:** Follow each platform's terms of service and the law where you live. Mature content is for adults only, and never depict a real person without their consent. Just because a tool allows something doesn't mean it's okay to make.

Pipelines in practice: characters

The pattern behind almost every character I build: explore cheaply, rebuild the winner in a production model, then break it into reference sheets for start frames and video.

Human character, start to finish

Soul Cinema to find the look, then everything downstream in GPT Image 2.5 Sunburst (The Back 40).



1 · Explore (Soul Cinema)
Higgsfield Soul Cinema · IMG-013



2 · Refined hero shot
GPT Image 2.5 Sunburst · IMG-014



3 · Approval sheet
GPT Image 2.5 Sunburst · IMG-015



4 · Basic turnaround
GPT Image 2.5 Sunburst · IMG-016



5 · Chemistry-test start frame
GPT Image 2.5 Sunburst · IMG-017

Watch out: Complex sheets can contradict themselves. The approval sheet's notes say 'bare ring finger' while its close-up shows a ring. Proofread every text callout.

Tip: The basic turnaround drops the head from the body views, so video tools can't grab a tiny face crop as the identity reference.

Creature, all in Seedream 5.0 Pro

Soul Cinema exploration, then a Seedream hero shot and turnaround (Behold Earth).



1 · Explore (Soul Cinema)
Higgsfield Soul Cinema · IMG-004



2 · Hero shot
Seedream 5.0 Pro · IMG-005



3 · Turnaround sheet
Seedream 5.0 Pro · IMG-006

More from the same episode, same method: Soul Cinema original → Seedream hero → Seedream turnaround.



Crested stalker · Soul Cinema original
Higgsfield Soul Cinema · IMG-111



Crested stalker · Seedream hero
Seedream 5.0 Pro · IMG-112



Crested stalker · turnaround
Seedream 5.0 Pro · IMG-113



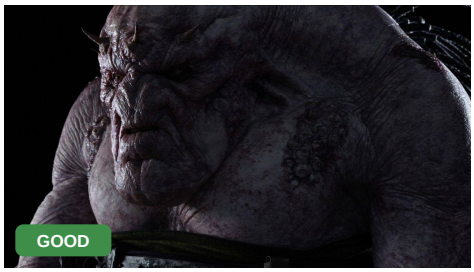
Ram-horned clawbeast · Soul Cinema original
Higgsfield Soul Cinema · IMG-114



Ram-horned clawbeast · Seedream hero
Seedream 5.0 Pro · IMG-115



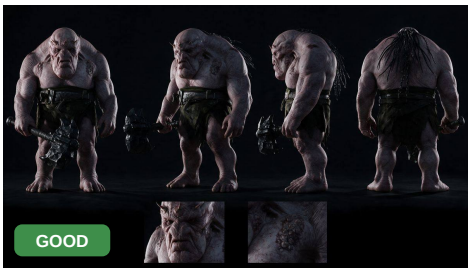
Ram-horned clawbeast · turnaround
Seedream 5.0 Pro · IMG-116



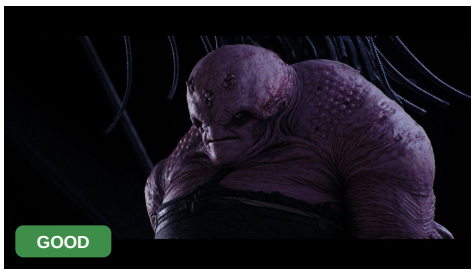
Hammer brute · Soul Cinema original
Higgsfield Soul Cinema · IMG-117



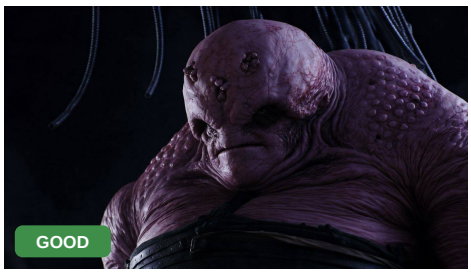
Hammer brute · Seedream hero
Seedream 5.0 Pro · IMG-118



Hammer brute · turnaround
Seedream 5.0 Pro · IMG-119



Announcer · Soul Cinema original
Higgsfield Soul Cinema · IMG-120



Announcer · Seedream hero
Seedream 5.0 Pro · IMG-121



Announcer · turnaround
Seedream 5.0 Pro · IMG-002

Three tools, plus a scale plate

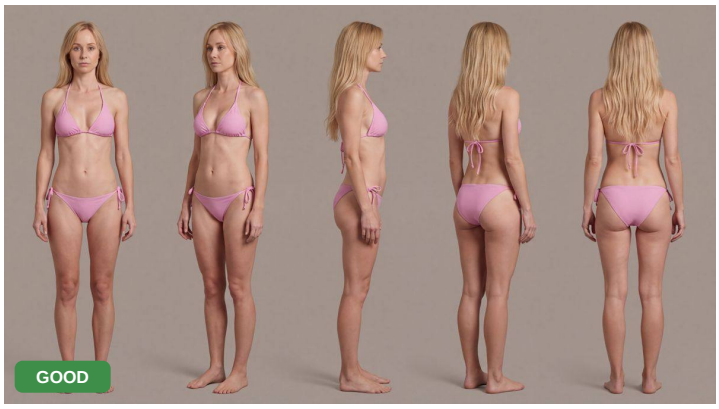
Soul Cinema → Seedream 5.0 Pro refine and scale check → Nano Banana Pro final turnaround (Trails End).



1 · Explore (Soul Cinema)
Higgsfield Soul Cinema · IMG-022



2 · Refine + scale plate
Seedream 5.0 Pro · IMG-023



3 · Final turnaround
Nano Banana Pro · IMG-024

> **Tip:** Build a scale plate once, then attach it to every later generation featuring both characters, labelled as a scale-only reference. It's far easier to prevent size drift than to fix it in video.

Handing off one step to another tool

Kevin the alien: Flare draft → Sunburst refine → Nano Banana Pro to add the cow-print shorts Sunburst couldn't land.



1 · Flare draft
GPT Image 2.5 Flare · IMG-058



2 · Sunburst refine (base sheet)
GPT Image 2.5 Sunburst · IMG-059



3 · Nano Banana Pro adds the shorts
Nano Banana Pro · IMG-060

> **Tip:** When one tool can't land a detail, don't fight it. Hand that single step to a different model.

One method, a whole cast

Every Under My Bed creature built the same way: Soul Cinema original → Nano Banana Pro hero → Nano Banana Pro turnaround. Names are internal working labels only.



Jailer ogre · Soul Cinema original
Higgsfield Soul Cinema · IMG-073



Jailer ogre · NB Pro hero
Nano Banana Pro · IMG-074



Jailer ogre · NB Pro turnaround
Nano Banana Pro · IMG-075



Big-head goblin · Soul Cinema original
Higgsfield Soul Cinema · IMG-076



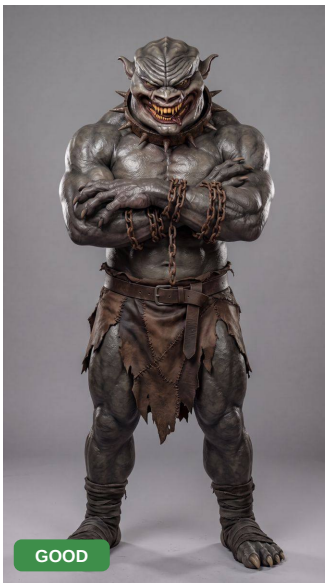
Big-head goblin · NB Pro hero
Nano Banana Pro · IMG-077



Big-head goblin · NB Pro turnaround
Nano Banana Pro · IMG-078



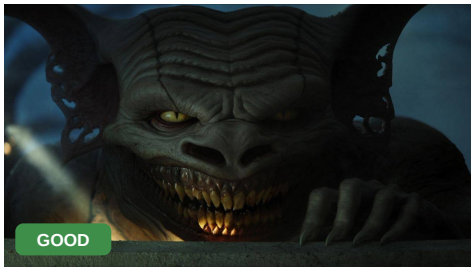
Chained brute · Soul Cinema original
Higgsfield Soul Cinema · IMG-079



Chained brute · NB Pro hero
Nano Banana Pro · IMG-080



Chained brute · NB Pro turnaround
Nano Banana Pro · IMG-081



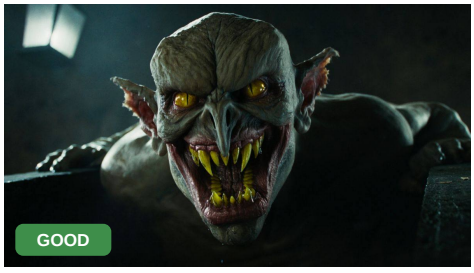
Bat-eared lady · Soul Cinema original
Higgsfield Soul Cinema · IMG-082



Bat-eared lady · NB Pro hero
Nano Banana Pro · IMG-083



Bat-eared lady · NB Pro turnaround
Nano Banana Pro · IMG-084



Vampire · Soul Cinema original
Higgsfield Soul Cinema · IMG-085



Vampire · NB Pro hero
Nano Banana Pro · IMG-086



Vampire · NB Pro turnaround
Nano Banana Pro · IMG-087



Knit-poncho goblin · Soul Cinema original
Higgsfield Soul Cinema · IMG-088



Knit-poncho goblin · NB Pro hero
Nano Banana Pro · IMG-089



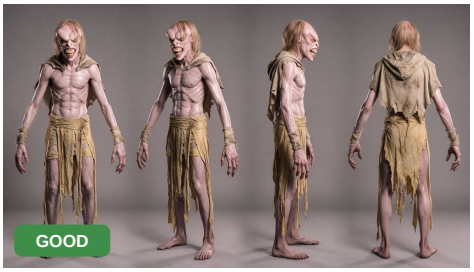
Knit-poncho goblin · NB Pro turnaround
Nano Banana Pro · IMG-090



Dangle · Soul Cinema original
Higgsfield Soul Cinema · IMG-091



Dangle · NB Pro hero
Nano Banana Pro · IMG-092



Dangle · NB Pro turnaround
Nano Banana Pro · IMG-093

Skipping exploration: Seedream hero → Nano Banana Pro turnaround

The rest of the Under My Bed cast, including one character in two looks.



Little beast · Seedream hero
Seedream 5.0 Pro · IMG-094



Little beast · NB Pro turnaround
Nano Banana Pro · IMG-095



Human · Seedream hero
Seedream 5.0 Pro · IMG-096



Human · NB Pro turnaround
Nano Banana Pro · IMG-097



Human, deformed · hero
Seedream 5.0 Pro · IMG-098



Human, deformed · turnaround
Nano Banana Pro · IMG-099

> **Tip:** For a transformation or before/after state, build each look as its own matched hero and turnaround, with identical wardrobe and body.

Putting a real person in the frame

Real photo → Nano Banana Pro hero → Seedream turnaround and creatures → Seedream backplate → GPT Image 2 composite (CM Productions promo). Several tools were tried at each stage; the tools shown are the ones that made the final.



1 · Real photo → AI hero
Nano Banana Pro (from real photo) · IMG-034



2 · 8-view turnaround
Seedream 5.0 Pro · IMG-035



3a · Creature: fairy
Seedream 5.0 Pro · IMG-036



3b · Creature: critter
Seedream 5.0 Pro · IMG-037



4 · Backplate
Seedream 5.0 Pro (several iterations) · IMG-038



5 · Final start frame
GPT Image 2 · IMG-039

> **Tip:** The little dragon in the backplate was a happy accident. Stay open to what the model adds.

Pipelines in practice: locations

The same idea for sets: lock layout and mood cheaply, then rebuild for texture, light and detail.

Cinema Locations → Nano Banana Pro

Five horror sets explored in Cinema Locations (left), then refined in Nano Banana Pro (right).



Lava canyon · before
Higgsfield Cinema Locations · IMG-100



Lava canyon · after
Nano Banana Pro · IMG-101



Haunted cabin · before
Higgsfield Cinema Locations · IMG-102



Haunted cabin · after
Nano Banana Pro · IMG-103



Flesh cavern · before
Higgsfield Cinema Locations · IMG-104



Flesh cavern · after
Nano Banana Pro · IMG-105



Abandoned campsite · before
Higgsfield Cinema Locations · IMG-106



Abandoned campsite · after
Nano Banana Pro · IMG-107



Glowing tent trail · before
Higgsfield Cinema Locations · IMG-108



Glowing tent trail · after
Nano Banana Pro · IMG-109

Interior refinement

Cinema Locations exploration → GPT Image 2 final after many iterations on light, textures, the view through the window, log walls and the ceiling paneling (The Back 40).



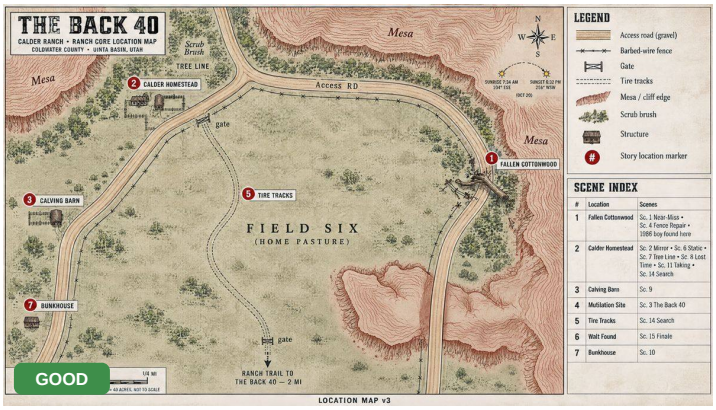
1 · Cinema Locations exploration
Higgsfield Cinema Locations · IMG-009



2 · GPT Image 2 final
GPT Image 2 · IMG-010

Master map → photoreal aerial

A Seedream 5.0 Pro location map anchors every location on one site, then GPT Image 2 converts it to a photoreal overhead for a drone flyover (The Back 40).



1 · Master location map
Seedream 5.0 Pro · IMG-007



2 · Photoreal overhead
GPT Image 2 · IMG-008

Establishing shot → coverage

One Sunburst establishing shot drives a 9-panel coverage storyboard; the best panels are regenerated full-scale.



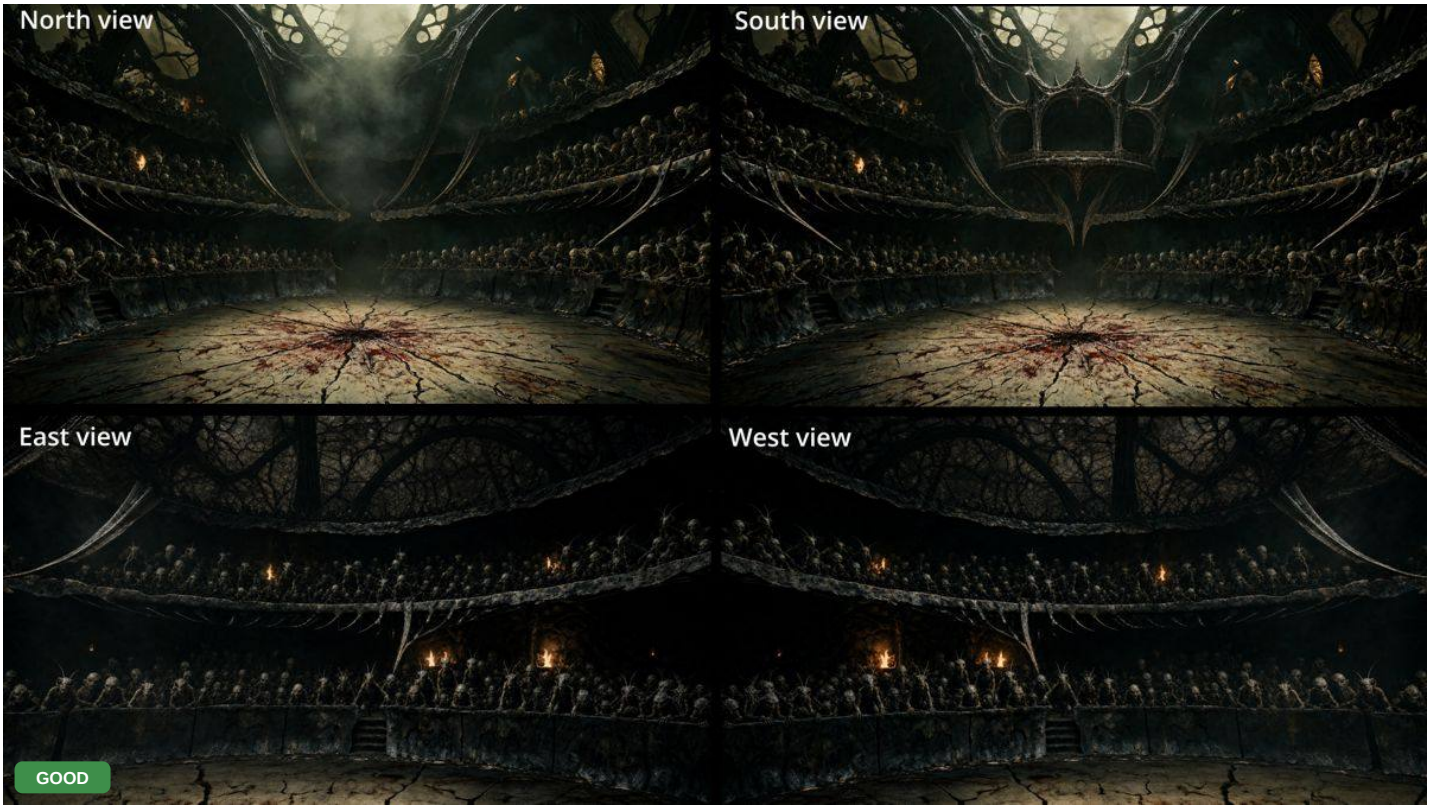
1 · Wide establishing shot
GPT Image 2.5 Sunburst · IMG-018



2 · 9-panel coverage storyboard
GPT Image 2.5 Sunburst · IMG-019

Set plate: four directions

Four compass-direction views generated in Seedream 5.0 Pro, hand-composited onto one sheet as the environment reference (Behold Earth, arena episode).



N/S/E/W set plate
Seedream 5.0 Pro + hand composite · IMG-003

Props, products and text

Text and hard-surface detail are where model choice matters most.

Vehicle: Nano Banana Pro at 4K



4K hero shot
Nano Banana Pro (4K) · IMG-011



Vehicle reference sheet
Nano Banana Pro (4K) · IMG-012

> **Watch out:** Even great sheets slip. The badge and tire lettering are garbled, and two views labelled FRONT and REAR are actually side views.

Text-heavy product: GPT Image 2

Hero shot with clean packaging text, then a full product turnaround (Cheese Yums).



1 · Product hero
GPT Image 2 · IMG-020



2 · Product turnaround
GPT Image 2 · IMG-021

Watch out: Trust the big text, proofread the small text. The turnaround's small panel text is garbled and two layout labels have typos.

Story assets in Sunburst

Cast lineup → wardrobe swap with a reusable scale reference → postcard front and back, plus a storyboard sheet used to find frames (Alien Interrogation).



Cast lineup
GPT Image 2.5 Sunburst · IMG-054



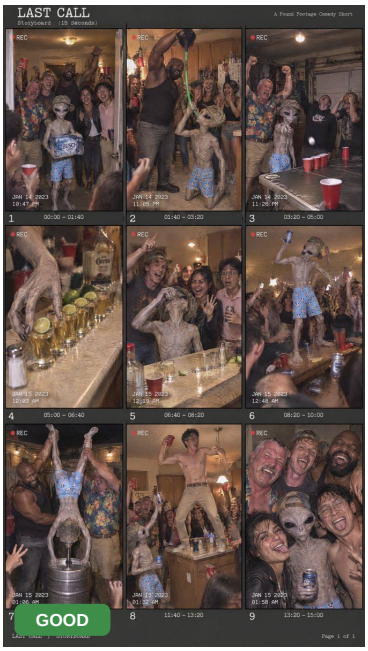
Wardrobe swap + scale reference
GPT Image 2.5 Sunburst · IMG-055



Postcard front
GPT Image 2.5 Sunburst · IMG-056



Postcard back (text-heavy)
GPT Image 2.5 Sunburst · IMG-057



9-panel storyboard → pulled frames
GPT Image 2.5 Sunburst · IMG-061

> **Tip:** Generate a storyboard sheet of many story beats at once, then pull the best panels into standalone full-scale frames.

Compositing and key art

Two approaches: one AI pass that combines separately built elements, or hand compositing in an editor.

Single-pass composite

Four characters composited onto a backplate in one Seedream 5.0 Pro generation (Behold Earth).



Four characters on one backplate
Seedream 5.0 Pro · IMG-001

Hand composite: group thumbnail

Character lineups and a backplate in Seedream 5.0 Pro, keyed and composited in DaVinci Resolve, then text, border and a 4K upscale. Extra characters from other projects were dropped in to fill gaps.



Lineup row 1
Seedream 5.0 Pro · IMG-027



Lineup row 2
Seedream 5.0 Pro · IMG-028



Lineup row 3
Seedream 5.0 Pro · IMG-029



Lineup row 4
Seedream 5.0 Pro · IMG-030



Lineup row 5
Seedream 5.0 Pro · IMG-031



Empty backplate
Seedream 5.0 Pro · IMG-032



Final composite thumbnail
Seedream 5.0 Pro + DaVinci Resolve + upscale · IMG-033

AI composite head-to-head

Elements built separately (Nano Banana Pro characters and space plate, Seedream title), then composited in two tools. GPT Image 2 narrowly won a vote across three school communities (Globs Ad 3).



Element: alien
Nano Banana Pro · IMG-040



Element: chocolate man
Nano Banana Pro · IMG-041



Element: scientist + bag
Nano Banana Pro · IMG-042



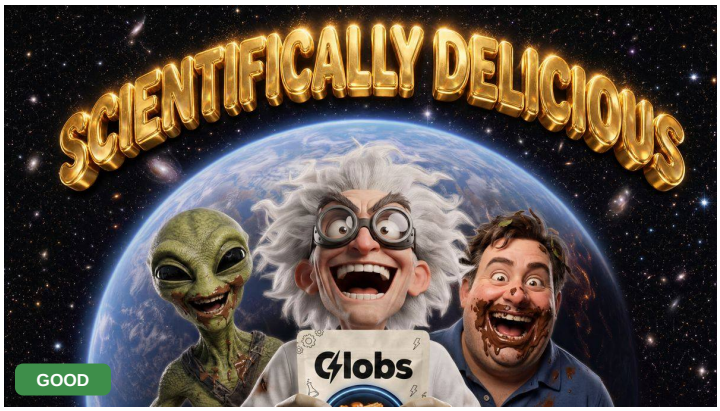
Element: 3D gold title
Seedream 5.0 Pro · IMG-043



Element: space backplate
Nano Banana Pro · IMG-044



Composite A · Seedream 5.0 Pro
Seedream 5.0 Pro · IMG-045



Composite B · GPT Image 2 (won vote)
GPT Image 2 · IMG-046

> **Tip:** Which tool did the best job is sometimes subjective. Both composites were excellent, and the vote was close.

Head-to-heads and known issues

The same references through different tools, plus the failure modes worth knowing before you spend credits.

Four tools, one brief

Same references through four tools for a 'leaked handheld footage' look (Alien Interrogation).



Flux 3 · not recommended
Flux 3 · IMG-047



Nano Banana 2.1 · painted skin
Nano Banana 2.1 · IMG-048



Seedream 5.0 Pro · too polished
Seedream 5.0 Pro · IMG-049



GPT 2.5 Sunburst · chosen
GPT Image 2.5 Sunburst · IMG-050

> **Tip:** The 'best' image depends on what the shot needs. Here the polished result lost to the imperfect one on purpose.

Photoreal vs animated drift

Same backplate and character references. GPT Image 2 holds photorealism; Nano Banana Pro drifts toward a stylized, game-cinematic look (Legendary Freshtness).



GPT Image 2 · stays photoreal
GPT Image 2 · IMG-069



Nano Banana Pro · drifts animated
Nano Banana Pro · IMG-070



GPT Image 2 · clean close-up
GPT Image 2 · IMG-067



GPT Image 2 · clean close-up
GPT Image 2 · IMG-068

Same prompt, both usable

Seedream 5.0 Pro vs Nano Banana Pro (Globs). There isn't always one best tool, so run key shots through more than one and cherry-pick.



Seedream 5.0 Pro (chosen)
Seedream 5.0 Pro · IMG-025



Nano Banana Pro
Nano Banana Pro · IMG-026

Older model vs current

Seedream 4.5 at native 4K vs GPT Image 2, with the same references and on-screen text.



Seedream 4.5 · native 4K
Seedream 4.5 (native 4K) · IMG-071



GPT Image 2 · clear winner
GPT Image 2 · IMG-072

Known issue: GPT noise and fractalization

GPT Image 2 and 2.5 (Sunburst and Flare) can add heavy noise to fine textures like skin, fabric, foliage and sky. They're still among the best tools available: prompt against it, refine, or remake the frame elsewhere.



Problem · Sunburst noise
GPT Image 2.5 Sunburst · IMG-062



Fix · Seedream 5.0 Pro remake
Seedream 5.0 Pro · IMG-063

Tool showcases

Single images that show a model at its best.

Gritty realism: GPT Image 2.5 Flare and Sunburst

Final start frames, cherry-picked from several iterations (Alien Interrogation).



Flare · crash-site approach
GPT Image 2.5 Flare · IMG-051

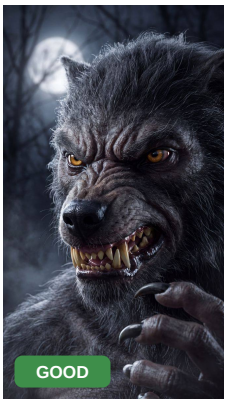


Flare · POV inside wreck
GPT Image 2.5 Flare · IMG-052



Sunburst · bar party
GPT Image 2.5 Sunburst · IMG-053

Nano Banana family: creatures and stylized looks



NB Pro · prompt-only werewolf, 4K
Nano Banana Pro (4K, prompt only) · IMG-066



NB 2 · alien stasis tank, 4K
Nano Banana 2 · IMG-065



NB 2 · stylized noir key art
Nano Banana 2 · IMG-110

GPT Image 2: clean photorealism



Earth from low orbit
GPT Image 2 · IMG-064

Working with Claude

Claude has been my production partner on almost every project in this guide. It doesn't generate AI images or video itself, but it handles a huge share of the work around them, often saving credits and hours.

Fixing images without spending credits

- Mirror a flipped profile on a turnaround sheet, or crop, rotate and re-frame an image, with no regeneration needed.
- Match the colour grade of one image to another, and build a LUT for DaVinci Resolve.
- Pull the last frame of a clip to use as the next start frame, for continuity.
- Split a storyboard grid into individual panels, or assemble images into reference and contact sheets.
- Resize and reformat for every aspect ratio (16:9 to 9:16, letterboxing, safe-area checks).

Video work in code

- Motion graphics rendered onto finished footage: title cards, lower thirds, animated text, end cards and HUD overlays.
- Subtitles and burned-in captions.
- Trimming, joining, reversing, flipping and speed changes, to get more out of footage you've already paid for.
- Stripping or swapping audio tracks, and exporting for each platform.

Audio

- Room tones and ambient beds such as wind, hum, air and rumble.
- Designed sound effects like whooshes, risers, impacts, beeps and glitches.
- Loudness matching across clips.

Prompts and planning

- Image and video prompts tailored to each tool and its quirks.
- Character, product and location reference sheets, using reusable skills you build once.
- Scripts, beat sheets, shot lists and storyboards.
- Invented languages and dialogue for characters.

Organization and business

- Production books and prompt logs, so nothing gets lost.
- Cataloguing and captioning hundreds of images. This guide's 110-image catalogue was built that way.
- Researching current prices, specs and model limits, with sources.
- Budgets, credit-cost calculators, client contracts and pitch decks.

Limits: Claude doesn't generate AI images or video itself. Its procedural sound effects are best for tones, ambience and designed sounds; realistic Foley like footsteps or cloth is still better from a library.

Going further: Claude can also connect directly to creative tools. Higgsfield's official MCP connector lets Claude generate images and video in Higgsfield from the chat, using your existing credits, and community-built connectors let it work inside Blender and DaVinci Resolve Studio. It's worth knowing that's possible, but I don't automate everything. I like staying in the creative loop: working with Claude, not having Claude do it all for me.

Cheat sheet

My default pick for each job, and what to try next. Guidelines, not hard rules.

Task	My pick	Also try
Images		
Cheap character or location exploration	**Soul Cinema / Cinema Locations**	Midjourney for stylized, surreal or physics-bending ideas
Hero shots and basic turnarounds	**Seedream 5.0 Pro**	GPT Image 2.5 Sunburst
Complex sheets and text-heavy images	**GPT Image 2.5 Sunburst**	GPT Image 2 for top fidelity (highest cost)
Gritty, handheld, found-footage realism	**GPT Image 2.5 Sunburst / Flare**	Seedream 5.0 Pro when you want polish
Creatures and aliens	**Nano Banana Pro / Seedream 5.0 Pro**	Nano Banana 2 for stylized looks
Photoreal from a text prompt only	**Nano Banana 2.1**	Nano Banana Pro
Refining a photoreal reference	**Seedream 5.0 Pro / GPT Image 2**	Avoid NB Pro (animated drift) and NB 2.1
Combining several elements in one frame	**GPT Image 2**	Seedream 5.0 Pro, or hand-composite
Fixing GPT noise	**Remake in Seedream 5.0 Pro**	Prompt against noise and regenerate
Cheapest native 4K stills	**Seedream 4.5 (1 credit)**	Nano Banana Pro 4K
Video		
Best possible quality	**Seedance 2.5**	Seedance 2.0
Native 4K video	**Seedance 2.0**	Kling 3.0 / Omni (cheapest 4K)
Testing a prompt cheaply	**Seedance 2.5 Draft (480p)**	Grok subscription
Budget quality	**MiniMax H3 (2K, 30 credits)**	Kling 3.0 Omni
Exploration, rough cuts, video refs for Seedance	**Grok subscription**	Seedance 2.5 Draft
Mature content (legal, adults only)	**Grok, kept Grok-native**	Open models run locally (MiniMax H3, LTX 2.5)

Habits that save the most credits

- Explore in the cheapest tool that shows the idea, and only pay premium prices for the winner.
- Run key shots through two models and cherry-pick.
- Build scale plates and basic turnarounds once, and reuse them everywhere.
- Proofread every piece of generated text, especially the small text.
- Simplify designs: every extra detail is one more thing to keep consistent shot to shot.

A final word

This guide's contents are my personal opinion after creating tens of thousands of images and half a dozen hours of video for projects over the last year. Treat it as rough guidelines for the best applications and practices for each tool.

The world of AI image and video generation is moving so fast that the best tool today might not be the best tool tomorrow. Those of us who make a living in this space keep up with these developments every day.

I hope you found this document useful, and I wish you the best of luck with all your future projects.

Chris Murray, CM Productions

CM PRODUCTIONS cm-promusic.ca | IG @cm_promusic | TikTok @cmpromusic.ca