ASSET·LAB User Manual

by Coastal Creative Lab

v2.5.0 · Linux

1 Getting started
2 The interface at a glance
3 Generating images
4 Generating video
5 Editing images
6 Audio — voice, music & sound effects
7 The Director — building a finished film
8 Prompt Lab & the Library
9 Writing prompts
10 Advanced controls
11 The Gallery & your media vault
12 Costs, balance & billing
13 Reliability
14 Keyboard shortcuts
15 Troubleshooting & FAQ

1 Getting started

Asset Lab is a generative studio for AI images, video and audio, powered by your own **Fal.ai** account. You pay Fal per generation — there is no subscription and no middleman. Full account setup (creating a Fal account, adding credits, creating an API key) is covered in the companion **Setup Guide** included alongside this manual. The short version:

1. Create a free account at **fal.ai** and add ~$10 of credits at **fal.ai/dashboard/usage-billing/credits**.
2. Create an API key — **the easy way:** on Asset Lab's welcome screen, click **"No key yet? Create or connect your Fal.ai account"**. fal.ai opens in your web browser; sign in, click **Add key**, and in the **Scope** dropdown switch from **API** (the default) to **ADMIN** — only an ADMIN key can show your live balance. Then click **Copy** — the key lands in Asset Lab automatically. Prefer to do it yourself? Create the key at **fal.ai/dashboard/keys** and paste it in.
3. Paste your **Gumroad license key** (from your purchase receipt), then click **Unlock the studio**. Both are verified on the spot and remembered forever — stored only on this computer. One license covers every computer you own.

Without keys, Asset Lab opens in **demo mode** — every screen is browsable, but Generate is locked until the license and API key are in.

✅ On Debian/Ubuntu the **.deb** is the recommended install — it registers the launcher and configures Chromium's sandbox for you. Elsewhere use the **AppImage** with the bundled `install-appimage.sh`.

2 The interface at a glance

ELEMENT	WHAT IT DOES
✦ Prompt Lab / Image / Video / Audio / Edit / 🎬 Director / Gallery tabs	The seven workspaces, in pipeline order: shape the prompt, generate the image, put it in motion, give it sound, refine it, cut it into a finished film, review it. Each create tab keeps its own prompt, model and settings — switching tabs never loses work. The app opens on Image.
Month (sidebar, under your balance)	Estimated Asset Lab spend this calendar month, tallied from per-generation estimates.

Balance (foot of the sidebar)	Your live Fal credit balance (requires an ADMIN-scoped key; shows "—" otherwise). Refreshes after every render, whenever you switch back to Asset Lab, and on click. Turns amber under $5 and red at zero.
Buy Credits	Opens Fal's billing page to top up.
Key pill (header)	Your masked API key. Click it — or **License & API key** in the sidebar — to manage your license & API key, or remove them.
"Demo mode · Unlock" (header)	Shown instead of the key pill when the app is locked — click it to enter your license and API key.
Clear (prompt toolbar)	Resets the current tab to a blank slate — prompt, attachments, options and result cards (after a confirmation). Finished renders always stay in the Gallery and vault.
Version tag	The app version (e.g. v1.4.0) — mention it when reporting an issue.

3 Generating images

1. Write a prompt (or use 🎲 **Surprise** / ✦ **Enhance** — see section 9).
2. Pick a model — click the **model selector** to open the catalog. Each entry states what it's best at:

MODEL	BEST FOR
Nano Banana 2	The all-round flagship — fast, smart, renders at 2K
Seedream 5.0 Pro	Logically consistent frames with visual reasoning
Seedream 5.0 Lite	Reasoning + web-grounded images at 3.5¢ each
GPT Image 2	Extreme detail and flawless text/typography in images
Qwen Image 3	**Best-in-class text rendering — posters, ads, typography that stays legible**
MAI Image 2.5 Pro	Microsoft's flagship image model — photoreal, up to 2:3 / 3:2 canvases
Ideogram V4	**Ideogram's flagship — crisp text, logos and poster layouts (≈ 1–2¢ per image)**
Ideogram V4 Fast / Instant	The same family in about a second or a fraction of one — for iterating on layouts
Ideogram V3	Previous Ideogram flagship — realistic or design looks with MagicPrompt
Ideogram Transparent	Logos, stickers and cut-outs on a real transparent background (alpha PNG)
Seedream 4.5	Razor-sharp output, strong at stylized work
Nano Banana 2 Lite	Sub-2-second renders, 14 aspect ratios
FLUX.2 [pro]	Pro-grade realism and style control
Nano Banana Pro	Deep-reasoning renders, realism + typography
Grok Imagine	Highly aesthetic, editorial-leaning frames

Kling Image O3	Kling's image model — consistent characters, ultrawide 21:9, up to 4K
Grok Imagine 2.0	**The current Grok image model — sharper, up to 2K, ultrawide ratios**
FLUX.2 [flash]	Near-instant, half-cent drafts — iterate here first
Recraft V4.1	Designer-grade composition and brand styles — includes an **Output** toggle that renders true vector SVG files (scalable, editable in Illustrator/Figma)
Recraft V4.1 Pro	Recraft's top design tier — up to 2048×2048 with the same raster/vector Output toggle (21¢ raster, 30¢ vector)

3. Choose an **aspect ratio** (the chip shows the actual shape) and how many **outputs** (1–4) you want.
4. Hit **Generate** — the button always shows the estimated cost first.

Social formats & resolution

Under the aspect chips, the **Social** row formats a render for a platform in one tap: **Instagram** and **Facebook** feed (4:5 portrait where the model supports it), **TikTok** (9:16) and **X** (16:9) — the tooltip shows the exact ratio being applied. Models that render multiple sizes also show a **Resolution** selector (1K/2K/4K on Nano Banana 2 and Pro, 1K/2K on Seedream 5.0 Pro and FLUX.2 [pro]); higher tiers cost more, and the price on the button updates instantly.

Reference images

On models that support it (Nano Banana 2, Seedream 5.0 Pro/Lite, GPT Image 2, Seedream 4.5), an optional **Reference images** slot appears under the model selector. Attach up to 4–10 images (click, drag-drop, or paste) and the model uses them to guide *style, subject, and composition* — great for keeping a character, product, or look consistent across renders. Leave it empty for pure text-to-image.

☆ **Cast a character** and ⌂ **Cast a location** sit above the slot. Each adds a saved entry's main image and its stitched sheet to whatever is already there, so a character and the place they're in can ride into the same render (see section 8).

Compare mode

Click **Compare** next to the Model label — the catalog stays open so you can select up to 3 models, and one Generate fires all of them with the same prompt. Results render side by side in a single strip so you can pick a winner. The button shows the combined estimate ("Generate ×3 ≈ $0.30").

4 Generating video

MODEL	BEST FOR	AUDIO	MAX LENGTH
Seedance 2.5	**Newest flagship — up to 30 s in one take, native audio**	Yes	30 s
Seedance 2.5 References	**Cast reference images, clips and audio into one shot — bound by @Image1 / @Video1 / @Audio1**	Yes	30 s
Seedance 2.0	Flagship cinematic shots, real physics, up to 4K	Yes	15 s
Kling O3 Pro	Newest-gen Kling — top-tier motion and visuals	Yes	15 s
Kling O3 Standard	O3 motion and audio at a friendlier rate	Yes	15 s
Kling O3 4K	O3's cinema tier — native 4K output	Yes	15 s

Kling 3.0 Pro	**Kling's current flagship line (distinct from O3) — top of the text-to-video arena; character-driven, multi-shot, native audio**	Yes	15 s
Kling 3.0 Standard	Kling 3.0 motion and audio at the friendlier rate	Yes	15 s
Kling 3.0 4K	Kling 3.0's cinema tier — native 4K output	Yes	15 s
Kling 3.0 Turbo Pro	**Kling 3.0 speed tier — 1080p, native audio, stronger lipsync**	Always on	15 s
Kling 3.0 Turbo Standard	Fast 720p Kling 3.0 with audio — for high-volume iteration	Always on	15 s
Veo 3.1	Google-grade realism, up to 4K	Yes	8 s
Veo 3.1 Fast	**Veo realism at speed — same look, a fraction of the flagship rate**	Yes	8 s
Veo 3.1 Lite	Veo realism at a friendlier price (1080p)	Yes	8 s
FLUX 3	**Black Forest Labs' frontier video model — up to 20 s with native audio**	Yes	20 s
FLUX 3 Draft	Fast, low-cost FLUX 3 previews — block the shot before you pay for it	Yes	20 s
Grok Imagine	Stylish clips with audio	Yes	15 s
Grok Imagine 1.5	Sharpest Grok motion — animates a start frame (image-to-video only)	Yes	15 s
Wan 3.0	**Coherent, story-first shots — up to 30 s in a single take, 480p–1080p**	Yes	30 s
Wan 3.0 Prime	Wan 3.0's premium tier — same controls, stronger fidelity	Yes	30 s
Wan 2.7	Smooth motion, high scene fidelity (1080p)	—	15 s
HappyHorse 1.1	Alibaba's flagship — the widest ratio choice (even 4:5 and 21:9)	—	15 s
Gemini Omni Flash 1.1	Fast, grounded in real-world knowledge — now 360p to 4K, start + end frames	Yes	10 s
Gemini Omni References 1.1	Cast up to 4 reference images into the shot — bind them in the prompt ("the product from image 1…")	Yes	10 s
MiniMax H3	Frontier motion and real-world physics — up to 4K	—	15 s
MiniMax H3 Max	H3's bigger sibling — stronger motion and physics (480P / 768P)	—	15 s
Seedance 2.0 Fast	Budget/turbo — iterate here before going flagship	Yes	15 s
Seedance 2.0 Mini	Seedance control at the lowest cost (720p)	Yes	15 s

Drag the **duration slider** in 1-second steps up to each model's true maximum (30 s on the Seedance 2.5 pair and Wan 3.0; 15 s on Seedance 2.0/Kling/Wan 2.7/Grok/HappyHorse; Veo uses fixed 4/6/8 s options). Set **resolution** where offered (720P / 1080P / 4K; the turbo tiers offer 480P/720P). The price on the Generate button recalculates as you change duration or resolution, so you always see the exact cost of the clip you've dialed in.

The Seedance 2.0 family also offers a **Bitrate** toggle: **Standard** (more compression, smaller file) or **High** (less compression, larger file) — pick High when the clip is headed for editing or big screens. Bitrate does not change the price.

Audio On/Off: models that generate native audio (Seedance 2.0 family, Kling O3 family, Veo 3.1 & Lite) show an Audio toggle. Off delivers a silent clip for your own sound design — and on Kling O3 Pro/Standard the price on the button drops to the cheaper audio-off rate. The setting sticks as you switch models.

Director drawer (cinematography): structured controls — pick a genre, color palette, lighting, camera move and lens, and Asset Lab appends precise direction to your prompt (shown verbatim before you spend). Save any combination as a named **look** ("Dusty Western") and re-apply it with one click to keep a whole project's cinematography consistent. The drawer is free — it's prompt engineering, not an extra model call. (It styles a single clip; to build a whole multi-shot film, use the 🎬 **Director tab** — section 7.)

Picking a model by task

The model picker opens with **task chips** across the top — *Cinematic*, *Character continuity*, *Native audio*, *Start → end frames* and *Cheap drafts* on the Video tab (*Product accuracy*, *Readable text*, *Character continuity* and *Cheap drafts* on Image). Click one to narrow the list to the models that are genuinely good at that job; click it again to see everything.

Prompt limits & the character counter

Some models cap how long a prompt can be — fal refuses anything over the limit. For those models a counter appears under the prompt once you're past 70% of the limit (for example **1,840 / 2,500**). It counts what is actually sent, including any Director-drawer or brand-kit notes Asset Lab appends, and turns red past the limit — Generate then stops you *before* the request goes out, so an over-long prompt never costs anything. The Kling 3.0 and O3 lines allow 2,500 characters; the Kling 3.0 Turbo tiers allow 3,072.

Multi-shot (Kling 3.0 & O3)

On the Kling 3.0 and O3 models the prompt box has a **Single prompt / Multi-shot** switch. In Multi-shot you write up to six short shots, each with its own length (1–15 s), and Kling renders them as **one continuous take** that cuts between them — same characters, same world. The shots must add up to 3–15 seconds; the total and every shot's character count are shown as you type, and the price follows the total length.

> 💡 Multi-shot is one clip from one model call. For a longer piece built from separate, individually approved shots — with narration, music, transitions and fades — use the 🎬 Director tab.

Animating an image (start frame)

Attach a **start frame** (click, drag-drop, or paste) and the model animates *your image*: the scene comes from the picture, the *prompt directs the motion and camera only*. The button changes to **"Animate image"** to confirm, and the aspect chips are replaced by **Auto** — the video automatically keeps your image's exact aspect ratio, whatever shape it is. (Gemini Omni Flash is the one exception: its image-to-video mode always renders 16:9 or 9:16, so the aspect chips stay visible.) **Grok Imagine 1.5** works exclusively this way — it always animates a start frame, so Generate stays disabled until you attach one.

Start → end frame

On Seedance, Kling and Wan 2.7 you can also attach an **end frame**. The model renders a transition from your first image to your second — the prompt describes how the shot evolves between them.

Reference images on video

Sixteen video models accept reference images as well as a start frame. Add images to the **Reference images** slot (or cast a character or location into it) and Asset Lab switches the render to that model's reference version automatically — the price on the Generate button follows.

MODEL	REFERENCE IMAGES
Wan 3.0 · Wan 3.0 Prime	up to 10
Seedance 2.0 · 2.0 Fast · 2.0 Mini · MiniMax H3 · H3 Max · HappyHorse 1.1	up to 9

Kling O3 4K · Grok Imagine · Grok Imagine 1.5	up to 7
Kling O3 Pro · Kling O3 Standard · Wan 2.7	up to 4
Veo 3.1 · Veo 3.1 Fast	up to 3

Reference versions have their own rules — Veo renders 8-second takes with references, and some can't combine references with a start frame. Asset Lab checks these before you spend and tells you what to change. MiniMax and Grok bill reference images on top of the clip (MiniMax H3 only past the first five); the price on the button already includes it.

Casting with references (Seedance 2.5 References)

Seedance 2.5 References goes further than any other model here: alongside reference **images** it accepts reference **clips** and reference **audio**, and blends them into a single new shot of up to 30 seconds with sound.

REFERENCE	HOW MANY	LIMITS	REFER TO IT AS
Images	up to 30	30 MB each	`@Image1`, `@Image2`…
Clips	up to 10	1.8–30.2 s each, 30.2 s combined, 200 MB each	`@Video1`, `@Video2`…
Audio	up to 10	MP3/WAV, 1.8–30.2 s each, 15 MB each	`@Audio1`, `@Audio2`…

⚠️ **The @ tokens are what bind a reference.** Writing "the product from image 1" does nothing on this model — you must write `@Image1`. Example: *"@Image1 walks through the alley from @Video1, scored like @Audio1, slow push-in, night."*

Each thumbnail is labelled with its own token, so the name you need to type is sitting on the thing you're pointing at. Asset Lab checks every limit above *before* you spend anything — a clip that's too short, an oversized file, or audio attached with no image or clip beside it is refused instantly and for free, rather than failing after a render is queued.

Reference audio needs at least one reference image or clip alongside it — that's a rule of the model, not of Asset Lab.

Casting with references (Gemini Omni References)

This model doesn't animate a picture — it **composes a new shot that contains your pictures**. Attach up to 4 reference images (a product, a face, a location) in the References slot and bind them by number in the prompt: *"the bottle from image 1 standing on the cliff from image 2, slow orbit, golden hour."*

- **Spoken lines:** put dialogue in quotes — *the woman from image 1 looks into the camera and says: "Own your tools."* — and the character speaks it with synced lips and native audio.
- **Limits to know:** 3–10 second takes (roughly 20–25 spoken words); the model chooses the voice, and it can differ between takes — for longer scripted reads with one consistent voice, generate the voiceover separately and use a lipsync workflow instead.
- **Pricing:** this endpoint is token-billed; the button shows fal's estimate (≈ $0.13 per second at 720p).

Extending a video

Open any video in the Gallery and click **Extend** →. Asset Lab grabs the clip's final frame and loads it as the next shot's start frame — describe the next motion and generate. Chain this to build sequences.

5 Editing images

The **Edit** tab transforms existing images with plain-English instructions.

1. Add 1–6 **source images** (click, drag-drop, or paste — or send any gallery image here via **"Edit this image"**).
2. Write the instruction: *"Make it golden hour", "Remove the car", "Put the product from image 1 into the scene from image 2".*

3. Pick an editor and Apply. Ten to choose from — the full **Nano Banana** family (2 / Pro / Lite), the **Seedream** family (5.0 Pro's region-precise edits, 5.0 Lite, 4.5), **GPT Image 2** (fine-grained detail), both **FLUX.2** editors, and **Grok Imagine**.

When you send a render here from the Gallery, Asset Lab preselects the editor from the **same model family that made it** — a Seedream render opens in Seedream Edit, keeping the look consistent. Renders from models without their own editor open in Nano Banana 2 Edit, and you can always pick any editor from the list: an edit works on any image, whatever made it.

Marking up regions (✎)

For surgical edits, click a source image (or its pencil badge). Drag to draw numbered, colored boxes around the parts you want changed, then **type what should change in each box** right in the panel — when you hit **Apply marks**, those notes are written into your instruction automatically ("In the red box: make the bottle emerald green…"). Prefer to write it yourself? Leave the note fields empty and reference the boxes in your own words. The boxes are baked into the copy sent to the model along with a note telling it to edit those regions and leave the rectangles out of the result. A **marked** badge shows which sources carry marks.

> 💡 Edits keep what you don't mention. "Make it daytime, keep everything else identical" changes the sky and lighting while preserving the scene.

Editing videos

Videos edit the same way images do. Open any clip in the Gallery and click **Edit this video** — it lands in the Edit tab as a playable source (hover to preview; its length shows on the thumb). Write the instruction and Apply:

1. **Kling O3 Edit / Edit Pro** — flagship restyles and reworks, 3–15s clips, with an Audio toggle that carries the clip's own sound through the edit ($0.126/s standard, $0.168/s pro).
2. **Wan 2.7 Edit** — style transfer and scene edits at the friendliest rate ($0.10/s).
3. **Gemini Omni Flash Edit** — simple instructions, fast turnarounds ($0.13/s).
4. **Grok Imagine Edit** — the cheapest way to try an edit on for size ($0.06–0.08/s by resolution).
5. **HappyHorse Video Edit** — reference-guided edits up to 1080p at a premium rate ($0.28–0.56/s).

Edits are priced per second of the source clip, and the estimate on the Apply button uses the clip's real length. Output duration matches the source. Expect video edits to render for several minutes on the flagship models — the job keeps running if you switch tabs.

New shot from a clip (same subjects, new camera)

Editing restyles a clip. **New shot from this clip** (in every video's Gallery lightbox) does something different: it generates a *brand-new shot* that keeps the clip's subjects, scene and style — a different angle of the same moment. It uses Kling O3's reference engine ($0.126/s standard, $0.168/s pro; 3–15s clips; several minutes to render).

> 💡 Refer to your clip as **@Video1** and ask for *one single continuous shot*: "A tight close-up of the driver from @Video1 seen through the windshield, driver on the same side as in @Video1. Do not replay footage from @Video1." If you skip the @Video1 reference, Asset Lab adds a single-shot anchor to your instruction automatically — and ✦ Enhance rewrites instructions into this grammar for you.

Shot Set — coverage from one prompt

On the Image tab, the **Shot set** button (or **Build shot set** on any gallery image) turns one scene into film-set coverage: pick from six angles — Wide, Medium, Close-up, Over-shoulder, Low angle, Aerial — and Asset Lab renders each as a reference-anchored still of the *same* moment. Start from a gallery image (it becomes the hero) or from a bare prompt (a hero frame generates first). Results land in the Gallery grouped as one shot set; the total price shows before you run. From there, **Animate this image** on your keepers and you've storyboarded, shot and animated a scene from a single prompt.

6 Audio — voice, music & sound effects

The **Audio** tab is a fourth workspace with four lanes across the top: **Voice**, **Music**, **Sound FX** and **Score**. Each lane keeps its own prompt, model and settings, and each one lays out exactly like the Image and Video tabs — results on the left, controls in the inspector on the right, the exact price on the Generate button. Finished audio lands in the Gallery with a waveform tile, and on disk in **~/Music/Asset Lab** — not Pictures, where your stills and clips live.

Voice — choosing one

Click the voice selector to open the picker. It holds **172 curated English voices** across six engines. Cartoon rosters and low-quality open models were left out on purpose — every voice on the list is one you could put in a client's video.

- **Preview is free.** Every voice ships with a recording of the same line, bundled with the app. Press ▶ on as many as you like — it costs nothing, needs no API key, and works before you've even unlocked the app. The samples are loudness-matched so no voice wins simply by being louder.
- **Filters:** Gender (male / female), Pitch (Deep / Mid / Bright) and Model. Or type in the search box to match a name.
- **About the gender labels:** most vendors don't publish this, so where it isn't documented Asset Lab measures the voice's pitch and infers it. An inferred label is shown in *orange* and can be wrong — trust your ears over the tag.
- **Pitch, not age or accent.** Vendors don't publish age or accent either, and guessing them would be worse than useless, so the app offers the one thing it can actually measure: the voice's fundamental frequency, banded into Deep / Mid / Bright.

ENGINE	RATE	NOTES
Inworld TTS 1.5	$0.01 / 1,000 characters	By far the cheapest, and the largest roster. Start here.
xAI TTS	$0.015 / 1,000 characters	Small roster, very cheap.
Seed Speech 2	$0.03 / 1,000 characters	ByteDance; natural conversational delivery.
Gemini 3.1 Flash	$0.05 / 1,000 characters	Google; wide roster.
MiniMax Speech 2.8 HD	$0.10 / 1,000 characters	The highest-fidelity option, *and the only engine that can speak a cloned voice.*
ElevenLabs Multilingual v2	$0.10 / 1,000 characters	The familiar benchmark for narration.

For scale: a 150-word script is about 900 characters, so it costs well under a penny on Inworld and about 9¢ on MiniMax HD. The exact figure is on the Generate button before you click.

Pauses — timing a read

Write `(pause)` anywhere in a script for a beat of silence, or `(pause 1.5)` for a timed hold of that many seconds (up to 5). *"OCHRE (pause 1.5) Own the moment."* reads as two deliberate phrases instead of one rushed sentence. It works in the Voice lane and in the Director's narration lines.

- **Exact silence** on MiniMax Speech 2.8 HD (and therefore every cloned or designed voice) and on ElevenLabs.
- **A natural breath** on Inworld, xAI, Seed Speech and Gemini, which have no pause control — the read slows at that point, but the length isn't guaranteed.

Voice — cloning your own

On the Voice lane, **Clone from a recording · $1.50** turns an audio file of someone speaking into a reusable voice. **Design from a description · $3.00** invents one from words alone ("warm, gravelly, late-night radio"). Both are *one-off charges* — after that the voice costs the same as any other to speak with.

1. Open **Your voices** (the same panel handles cloning, designing and your library).
2. Choose an audio file — **at least 10 seconds** of clean speech, one speaker, no music behind it. About a minute is the sweet spot.
3. Expand **How to record a sample that clones well** for a ready-made script and the recording technique. The script is written to cover every English sound, three kinds of sentence and the spoken numbers clones most often fumble, and runs about a minute at a normal pace.
4. Tick the consent box — you must own the recording or have the speaker's permission.
5. Name it and click **Clone voice**. It appears at the top of the voice picker under **Your voices**, ready for any future project.

The seven-day rule, and why Asset Lab handles it for you. Cloned voices live on your own fal account, and MiniMax deletes one that goes **seven days without being used**. That's real money evaporating over a long weekend, and fal has no way to list the voices you own — *this app is the only registry of them*. So Asset Lab speaks one word with a new clone the instant it's created, tracks when each was last used, warns you from day five, and gives you a one-click renew. Use **Back up** to write your voice list to a file and **Restore** to bring it back on another machine or after a reinstall — do this, because losing the list loses the voices.

Recording tips

- **Somewhere soft.** A room with a rug, curtains or a wardrobe. Hard surfaces add echo, and the clone copies the echo along with the voice.
- **Kill the hum.** Fans, air conditioning, anything with a compressor. Steady background noise is the most common reason a clone comes back thin.
- **Your phone is fine.** Voice Memos, held a hand-span away and slightly off to one side so your breath misses the mic, beats a cheap USB microphone in a bad room.
- **Read the way you want it to sound.** The clone copies your delivery, not just your timbre. Don't perform — an exaggerated read makes an exaggerated voice.
- **Don't clean it up.** No noise reduction, EQ, compression or de-essing. Processing strips out the detail the cloner needs.

Music

The Music lane opens on **Composer**, which turns a brief into chips: genre, mood, instrumentation, tempo, key, arrangement and vocals. Expand **What gets sent** to read the exact caption before it goes — nothing is hidden. Switch to **Write it myself** if you'd rather type the description yourself.

MODEL	RATE	BEST FOR
MiniMax Music 3	$0.002 / second	The only model that genuinely parses the composer's BPM and key. Also takes lyrics with section tags.
ElevenLabs Music	$0.60 / minute, *rounded up*	The most controllable. Note the rounding: a 20-second sting costs a full minute.
Lyria 3	$0.04 flat	The cheapest way to fish for ideas. Fixed 30 seconds.
Lyria 3 Pro	$0.08 flat	Up to three minutes — the model picks the length, so there's no slider.
Sonilo v1.1	$0.0025 / second	The vendor markets it as licensed and commercial-use-safe. Reach for it on client work.
CassetteAI	$0.02 / minute	Near-free background beds, very fast.

Why some tracks come back short. On MiniMax Music 3 the length slider is a *maximum*, not a target — the model stops when the piece feels finished, so a 60-second request can return 33 seconds. The app labels that slider "Max length", says so underneath, and quotes the price as "up to $X" rather than pretending. Every other music model here honours the length you set.

Sound FX

Describe a single sound — the source and material, the action, the space it happens in, and how it decays. Two models:

- **ElevenLabs SFX v2** — $0.002 per second, up to 22 seconds, with a **Seamless loop** toggle for ambience that has to run under a scene and a **Prompt influence** slider. A five-second effect costs about a penny, which is why the lane offers 1 / 2 / 4 variations — rolling four and keeping the best beats re-rolling one.
- **Stable Audio 2.5** — $0.20 flat, for anything longer or stranger than 22 seconds.

Score — music written to a video

The Score lane starts from footage instead of a description. Pick a clip from your Gallery, or **Or upload a clip…** to use any video on your drive, and Sonilo reads the clip's pacing and cuts and writes music that lands with them. You can also start from the Gallery: open any video in the lightbox and click **Score this clip**.

- **$0.009 per second, per take.** A 30-second cue is about 27¢ for one take, 81¢ for three.
- **Takes (1 / 2 / 3)** give you different scores for the same clip — the price on the button already accounts for them.
- **Leave the description empty and Asset Lab watches the clip first:** for ≈ $0.001 it samples frames, writes a period-accurate scoring brief from what's actually on screen (a 1940s scene gets 1940s music), and drops the brief into the description box — so you can refine it for the next take.
- ✦ **Enhance** in this lane speaks film-composer — era, genre, instrumentation, how the music evolves — never sound-effect language.
- Clips already in your Gallery are used directly; a clip you upload is sent to fal first.

Give the voice a face — lip sync

A voiceover is only half a piece. Open any take in the Gallery and click **Make a talking video**: Asset Lab hands it to the **Video** tab, you attach a face, and a character speaks your words — in your voice, or your cloned one. Starting from the other end works too: open any image and click **Make this talk**, then attach a voice. You can also press **Add a voiceover** on the Video tab and pick an audio file from disk.

What you attach decides which models you get. Drop in a *picture* and you get the avatar models, which animate a still. Drop in a *clip* and you get the dubbing models, which keep the footage and change only the mouth and the soundtrack.

STILL + VOICE → A TALKING CHARACTER	RATE	NOTES
OmniHuman 1.5	$0.16 / second	ByteDance. Emotion and gesture follow the voice, not just the mouth, and an optional prompt directs the performance. **720p takes up to 60s of audio, 1080p up to 30s** — the app checks before it spends.
sync-3 Avatar	$0.133 / second	Works on illustrations and animation frames, not only photographs.
Flashtalk	$0.02 / second	768×448, so rough one out for pennies and commit on a better model.

CLIP + VOICE → THE CLIP RE-VOICED	RATE	NOTES
HeyGen Lipsync — Speed	$0.05 / second	Dub a clip you already made. Optional burnt-in captions and speech clean-up.

HeyGen Lipsync — Precision	$0.10 / second	The same with tighter mouth tracking.
sync-3 Lipsync	$0.133 / second	The one for when the voice and the clip are different lengths — choose cut off, loop, bounce, hold, or retime.

The price is exact, not an estimate. Every one of these bills per second of the audio you hand it, and Asset Lab already measured your take's length when it drew the waveform — so the figure on the Generate button is the real number. A 24-second read costs about 48¢ on Flashtalk and $3.84 on OmniHuman.

Two things worth knowing. The output is exactly as long as your audio, so there is no length slider — trim the script, not the video. And **a reference model is not a lip-sync model**: Seedance 2.5 References and Gemini Omni References both accept audio, but they use it for *timing* — cutting the action to a beat — and generate their own speech. Only the models above put *your* recording in a character's mouth.

About the files you attach. Every face, clip and voice is uploaded to fal's storage first and checked to be actually there before the model is asked to fetch it, so a stale link or an unreadable file fails instantly with a message that names the input — "the face", "the voice" — instead of a render that dies with "failed to download the file". **OmniHuman refuses images over about 5 MB** (the other models don't care), so Asset Lab re-encodes a large photo to a JPEG under 4 MB at the same resolution before sending it — a 20 MB desktop PNG just works. Clips for the dubbing models can be up to 200 MB.

Prompt tools

🎲 **Surprise**, **History**, **Clear** and ✦ **Enhance** work in every audio lane, and they know which lane they're in. Enhance rewrites a voiceover script *for reading aloud* — short sentences, spoken rhythm, no em dashes (some engines read them out) — while on Music it turns a rough idea into a proper brief with genre, instrumentation and tempo, and on Sound FX into a concrete single-sound description.

Licensing

Every audio model in Asset Lab is marked for commercial use on fal, so output is yours to use in client work. Two caveats worth knowing: fal warrants nothing about originality or non-infringement, and Lyria 3 embeds Google's **SynthID** provenance watermark in what it generates. **Sonilo** is the only vendor here that actively markets itself as licensed and commercial-safe, which is why it carries a tag in the picker. For voice cloning, the consent is yours to obtain — clone your own voice, or one you have written permission to use.

7 The Director — building a finished film

The 🎬 **Director** tab turns single clips into a finished piece: a commercial, a trailer, a short. You plan it shot by shot, approve the best take of each, add music and a narrator, choose transitions, and export one video — straight into your Gallery.

1 · Start a sequence

Click **+ New**. Each sequence appears as a chip on the Director bar; the bar also holds **Rename**, **Save to gallery**, **Close** and **Delete**. Sequences are remembered between sessions.

2 · Sequence setup — what every shot shares

The orange-framed panel at the top holds everything that must stay the same across the cut, in three cards:

CARD	CONTROLS
✦ The look	**Model** (every shot uses it, so the look holds), **Aspect** (the ratios that model offers), **Resolution** (every shot renders at it — change it and every shot's price re-quotes), **Model audio** on/off where offered, and **Open & close**: Fade in / Fade out, each with its own length (0.5–3 s).
◉ The world	**Anchor frame** — one image every shot can open from, so the whole cut shares one world. **References** (models that take reference images) — images that ride into every take: pick from your Gallery, upload, or cast a saved ☆

	character, ⌂ location or ⬒ brand kit from the clearly labelled sections of the picker.
♪ The sound	**Soundtrack** — any track from your Audio gallery. **Narrator** — one voice for the whole cut, chosen from the full voice picker (your cloned voices included).

💡 **Keeping characters consistent.** On models that take reference images (section 4) put the character — and the location — in *References* and name them in each shot's prompt. On every other model use the *Anchor frame* — or chain shots together (below). Sending an image from the Gallery with *Use as reference in Director* switches a sequence to a reference model automatically.

3 · Write and shoot each shot

Each shot card has a prompt, a length, ↑ ↓ to reorder and × to remove — up to six shots. The length menu offers exactly what the sequence's model accepts (4, 6 or 8 s on Veo; up to 30 s on Seedance 2.5 and Wan 3.0), and switching models moves every shot to the nearest length the new model allows. Where the model can open from an image, a **Start** control sets how the shot begins:

- **anchor** — from the sequence's anchor frame;
- **chain prev** — from the last frame of the previous shot's approved take, for a seamless continuation;
- **fresh** — from the prompt alone.

✦ **Generate take** shows its ≈ price and renders a take; generate as many as you like and click the one you want to **approve** it (✓ take). Already made the clip on the Video tab? **Use a gallery clip…** drops it in as a take at no cost.

4 · Narration

With a narrator chosen, every shot gets a narration line. Type it and press **Read** (≈ $0.002–0.01 a line — the exact price is on the button). The finished read appears as a chip with ▶ to listen, its length, and × to discard.

- **at *n* s** — when the voice comes in, in seconds after the shot starts. A 10-second shot with a 3-second line set to *at 5* is heard from 5 s to 8 s.
- **(pause)** or **(pause 1.5)** inside the line adds a beat of silence (see *Pauses* in section 6).
- If a line would run past the end of its shot, the chip turns red and **fit shot** lengthens the shot to fit.

5 · Transitions

Between every pair of shot cards sits a transition picker: ✂ **Cut** (the default), ⇄ **Cross-fade**, → ← ↑ ↓ **Swipe** in any direction, ◐ **Dip to black** and ◑ **Dip to white**. Each transition overlaps the two shots by half a second; narration timing is adjusted automatically so every line still lands exactly where you put it.

6 · Preview, mix, export

- ▶ **Preview cut** plays the approved takes in order with the soundtrack and narration. (The preview uses straight cuts — transitions and fades appear in the export.)
- **Mix** — the ♪ Music and 👤 Voice sliders set the levels, heard live in the preview and baked into the export.
- **Render to Gallery** assembles the film, keeps it in your Gallery and vault, and opens it. **Render to Disk** assembles the same film but asks where to put it and writes exactly one copy there, leaving the Gallery and the vault untouched — the choice for a deliverable that belongs in a client folder rather than in your library. Either way, a cut with any transition or fade is rendered **on your computer — free, and usually in seconds**, and a plain-cut sequence is stitched on fal for a few cents. If you dismiss the save sheet on a paid stitch, the finished cut is kept in your Gallery rather than thrown away.

Saving and reopening

Save to gallery stores the whole sequence — settings, shots, takes, narration, offsets and mix — as a 🎬 card in the Gallery (filter: 🎬 Sequences). Saving again updates the same card. **Close** saves first, then clears the sequence off the Director bar; clicking its card any time reopens it exactly as you left it. Saved sequences are included in Back up / Restore.

8 Prompt Lab & the Library

The ✦ **Prompt Lab** tab (first in the bar) is a guided prompt builder plus a library of everything worth keeping — saved prompts, reusable blueprints, your characters and your locations. If a render ever comes out generic, this is where you learn why.

Build — the Blueprint

Pick **An image** or **A video**, then answer one question per step: **Subject** (who or what, with specifics), **Setting** (where and when), **Camera** (framing and angle), **Light** (name the source and direction — never just "cinematic"), **Mood** (one or two words), **Finish** (an optional look — skip it for realism), and for video, **Motion** (one camera move plus what happens in frame). Tap the suggestion chips or write your own; the prompt assembles live in the side rail, and the rail shows which blocks are still blank — blank blocks are where generic renders come from.

- ✦ **Draft it (≈ $0.001):** type a rough idea and the AI fills every block for you — review and edit each one.
- ✦ **Polish (≈ $0.001):** rewrites the assembled prompt into fluent, model-ready prose. Undo restores it.
- **Send to Image / Video** → drops the finished prompt straight into that tab.
- On the Image and Video tabs, the **"Need a stronger prompt?"** link under the prompt box jumps here with your draft carried along.

The Library

Five kinds of entries, all renameable, searchable by eye, and one click from use:

- **Prompts** — saved from the lightbox ("Save prompt to Library") or from a finished Blueprint. Each carries its model and seed, so **Use** → restores the whole recipe. Saved prompts also appear, starred, at the top of History on the Image and Video tabs.
- **Blueprints** — a saved set of blocks. **Open in Build** to swap the Subject and keep the world: same setting, light, mood and finish across a whole campaign.
- **Characters** — see below.
- **Locations** — see below.
- **Brand kits** — see below.

Characters — the same face in every scene

When a render nails a person, open it in the lightbox → **Save as Character…** and name them. (Works on videos too — the character anchors to the clip's last frame.) A character is a reusable identity: its reference images ride into any reference-capable model so the same face appears in new scenes, new outfits, new worlds.

- **Casting:** the ☆ **Cast a character** button on the Image and Video tabs' reference strip adds the character's main image and its stitched sheet — two images that carry the whole identity, leaving room beside them for a location or your own references. Or use **New scene: Image → / Video →** on the character's Library card, or start the Prompt Lab Subject step from a saved character.
- **Build character sheet:** renders a full 360° turnaround — six stops around the character (front, ¾ left, left profile, right profile, ¾ right, back), each as a full-length figure and as a head shot, on one seamless studio backdrop — 12 views by default at ≈ $0.12 per view (≈ $1.44 for the set). Untick any angle or framing, or add Smiling, Serious or a custom shot you describe. The finished sheet shows on the character's card as a lookbook grid (full-length row on top, head shots beneath); **View sheet** opens it full-size. The stitched sheet is saved to the Gallery for you (**Save as one image** re-stitches it any time), and every view plus the sheet is filed in a folder named after the character and grouped as one project in the Gallery. Every view joins the character's references automatically, so each sheet makes the identity lock stronger in every future scene.
- **Details:** type what must stay the same — ripped jeans, leather shoes, a silver nose ring — in the card's **Details** box. They are written into every sheet view and carried into new scenes started from the card.

Locations — the same place in every scene

When a render nails a place — a diner, a harbor street, an office — open it in the lightbox → **Save as Location…** and name it. Locations sit in the Library under their own filter and work the way characters do, for places.

- **Details:** what must never change — rain-slick cobblestones, a red neon bar sign, overcast dusk — carried into every sheet view.

- **Build location sheet:** renders the place from seven angles — Wide, Reverse, Left side, Right side, High angle, Low angle and a close Detail — at ≈ $0.12 per view (≈ $0.84 for the set). Untick any view, or add Night, Golden hour or a custom view you describe. No people, text or captions are added. The views and the stitched sheet are saved to the Gallery, filed in a folder named after the location, and grouped as one project.
- **Casting:** ⌂ **Cast a location** on the Image and Video tabs adds the location's main image and sheet to whatever is already in the reference strip — cast a character, then a location, and the character appears in that place. In the 🎬 Director, pick it from the Locations section of a sequence's References picker.

✦ Analyze & fix

The first button beside a render in the lightbox. Describe what's wrong with a render (or leave it blank for a straight critique) and for ≈ $0.001 the AI looks at the actual image or clip next to the prompt that made it, then hands back a plain-language diagnosis plus two exits: a surgical **Fix in Edit** → instruction applied to this render, and a rewritten **Re-generate** → prompt. The analysis is saved with the item, so the prompt you didn't use is never lost — reopen any time, refine your notes, and ask again.

9 Writing prompts

✦ Enhance

Rewrites your rough idea into a detailed, model-optimized prompt (composition, lighting, palette, lens for images; motion + camera for video). It knows when a start frame is attached and will direct *motion only*, never re-invent your scene. **Undo** restores your original. Costs a fraction of a cent (it runs an LLM through your Fal key).

🎲 Surprise me

Generates one wild, fully-formed prompt for the current tab. Great for exploring a model's range or demoing the app.

History

Your last 50 prompts across all tabs — click any to restore it. Prompts you've saved to the Library appear ★-starred at the top of the panel, above the rolling history.

Tips for good prompts

- **Images:** subject → setting → light → palette → style/medium. Specific beats long.
- **Video:** describe *one* continuous shot with one camera move. Multiple scene changes confuse models.
- **Edits:** say what changes *and* what must stay ("keep the face identical").
- Iterate on the Turbo models (FLUX.2 [flash] / Seedance Fast), then re-run the winner on a flagship.

10 Advanced controls

Click **+ Advanced** (shown when the selected model supports extras):

CONTROL	EFFECT	AVAILABLE ON
Seed	Same seed + same prompt = same image. Use to reproduce or make controlled variations.	FLUX family, Seedream, Nano Banana, Veo
Negative prompt	What to avoid ("blur, text, watermark").	Veo, Wan

11 The Gallery & your media vault

Everything you generate lands in the **Gallery** automatically — stills, clips and audio alike. Audio tiles show a waveform and play in the lightbox. The desktop app also saves every file to **~/Pictures/Asset Lab** (the "vault"). That means your work survives even after Fal's temporary download links expire.

The gallery never forgets. There is no cap on how many renders it keeps — every generation stays reachable forever. Big collections are paginated 50 to a page, with pager controls above and below the grid; filters and search always span the whole collection.

Browsing

- Filter by **images / videos / Audio / 🎬 Sequences / ★ Starred**, or search prompts and model names. A 🎬 Sequence card is a saved Director cut-in-progress — click it to reopen it in the Director (section 7).
- Every card shows its prompt and cost; the corner ☆ adds it to your favorites.
- **Select** mode enables bulk download, bulk delete — and **Project…**, which groups the selected items into a named project. A project filter appears in the gallery bar, tiles show a project tag, and on the desktop app the files themselves move into a matching subfolder of `~/Pictures/Asset Lab` — so your file manager mirrors your gallery. "Remove from project" moves everything back to the vault root.

Delivery packs

In Select mode, pick the renders a client should receive and click **Pack…**. Name the pack and Asset Lab builds a folder in `~/Pictures/Asset Lab/Exports` with the files numbered in order plus a `manifest.json` recording the prompt, model and seed of every asset — then reveals it in your file manager, ready to zip and send. Files come straight from your vault (or are re-downloaded if a local copy is missing). Free.

Back up & Restore

Back up (in the gallery bar) writes one portable folder to `~/Pictures/Asset Lab/Backups`: your media plus the gallery records, prompts, characters, brand kits, saved looks and settings. With a project filter active it backs up just that project. On another computer, click **Restore** and choose the backup's `assetlab-backup.json` file — everything merges into that computer's gallery and vault, and nothing you already have is duplicated. Saved Director sequences travel along too.

The item view (lightbox)

Click any item to open it. Use ← → (or the on-screen arrows) to flip through the gallery, `Esc` to close. Actions:

The bar keeps ☆ Favorite, Download, Show in file manager, Delete and Close as plain links, and shows every creative action as buttons beside the render — grouped by workflow: iterate first, finish second, then save.

BUTTON	WHAT HAPPENS
✦ Analyze & fix…	AI diagnosis + two ready-made fixes (see the Prompt Lab chapter)
Edit this image / video	Opens the Edit tab — in the same model family that made the render, when that family has an editor
Animate this image	Sends an image to the Video tab as the animation source
Use as reference — clean prompt	Sends an image to the Image tab as a reference with an empty prompt — a fresh start with the same subject. The aspect ratio follows the image
Use as reference in Director	Adds an image to the active Director sequence's references (starting a sequence if none is open)
Make this talk	Sends a face to the lip-sync lane, ready for a voice

Build shot set	Six film-grammar angles of the same scene, anchored to this render
Extend this video	Grabs the last frame and preloads the Video tab to continue the shot
New shot from this clip / Score this clip	(Videos) re-shoot the same scene, or write music to it
Save as Character… / Save as Location… / Save prompt to Library	Keeps the identity, the place, or the whole recipe (prompt + model + seed), in the Prompt Lab Library
Add to Brand Kit…	(Images, when you have a kit) adds the image to a kit's anchors
Reuse prompt	Restores that item's prompt and model in its original tab (audio takes open the matching Audio lane)

Tools

The last group of those buttons is **Tools**: finishing passes that work on the render you already have instead of generating a new one. Picking one opens a confirm bar with its options and the exact price for *this* item — measured from its real pixels and running time, not a generic rate — and nothing is billed until you press Run. The result lands in your gallery as a version of the original.

TOOL	WHAT IT DOES	TYPICAL COST
Upscale image	Topaz Precision — a faithful 2× or 4× enlargement that sharpens what is there rather than inventing detail	$0.08 per 24 megapixels of output
Remove background (image)	Ideogram — cuts the subject out onto transparency	$0.01
Upscale video	Topaz Precision (Proteus) — a resolution lift with no restyling	$0.40 / $0.80 / $3.10 per minute at 720p / 1080p / 4K
Smooth / slow motion	Topaz frame interpolation — 60 fps motion, or true 2×/4× slow motion without ghosting	from $0.30 per 10s clip
Remove background (video)	Bria — keeps the subject, drops the scene. Transparency is delivered as WebM, the only container here that carries an alpha channel	$0.05 per second
Crop to aspect	A straight crop, run by the app on your own computer. Pick 9:16, 4:5, 1:1, 4:3 or 16:9 and which edge to keep, and the confirm bar tells you the output size and how much of the frame survives. Nothing is invented, nothing is uploaded	Free
Clean plate — erase the subject	LTX — rebuilds the background as if nothing were in front of it. The model's control budget caps the plate near 6s at 720p or 15s at 480p, and the confirm bar tells you how much of the clip it will cover	≈ $0.32 for a 6s 720p plate

Tools need a cloud copy of the render, so they are hidden for Director cuts assembled on your own computer, and for audio, which has no upscaler of its own.

The item's caption line also shows the **seed** the model used — same seed, same prompt, same settings reproduces the render.

Versions. Every render remembers what it was made from — an edit, an extension, an upscale, a reference re-shoot. When an item has relatives, the lightbox shows a **Versions · N** strip of the whole family in the order they were made; click any thumbnail to jump to it.

Director cuts with transitions are rendered on your computer, so they have no cloud copy: the lightbox offers Analyze, Extend, Save as Character and Show in file manager for them, but not the ones that need to send the clip to fal (Edit, New shot, Score, and the whole Tools section).

12 Costs, balance & billing

- **Before you spend:** the Generate button always shows an estimate (e.g. "≈ $0.12") computed from Fal's published rates for the exact duration, resolution, aspect and output count you've chosen. A few models without published rate cards fall back to Fal's historical average — the ≈ covers both cases.
- **One-off charges:** Voice cloning ($1.50) and voice design ($3.00) are the only click-once-money-gone actions in the app, so both ask for confirmation first and state the price on the button. Everything else is billed per render.
- **After you spend:** each gallery item records its estimated cost, and the **Month** line under your balance keeps a running total.
- **Balance:** live Fal credits at the foot of the sidebar (ADMIN-scope key required), refreshed after every job. **Buy Credits** opens Fal's top-up page.
- **Running low:** when your Fal balance drops to **$5 or less**, Asset Lab warns you before a render fails, and again if the account reaches zero. Both warnings link straight to Fal's top-up page. Dismiss it and the balance in the sidebar stays coloured as a reminder; top up and it resets. **You do not need to restart Asset Lab after topping up** — the balance refreshes the moment you switch back to the app, and you can click the balance in the sidebar to refresh on demand. This warning needs an ADMIN-scoped key, since Fal only reports balances to those.

ROUGH PRICE TIERS	EXAMPLES
Under a penny	FLUX.2 [flash] drafts, Enhance/Surprise prompts, a 5-second sound effect, a short script on Inworld
1¢ – 20¢	Music beds, a 150-word script on any voice, a 30-second scored cue
$0.05 – $0.20	Flagship images, image edits, image upscales
$0.20 – $1.50	Budget video clips, video upscales
$2 – $6	Flagship video (Seedance 1080p+, Veo 3.1, Kling 3.0 / O3)

13 Reliability

- **Parallel queue:** fire as many generations as you like — each job card tracks its own queue position and render status, across all tabs at once.
- **Cancel:** every running job has a Cancel button.
- **Crash recovery:** if the app quits mid-render, the job reconnects automatically on next launch and the result still lands in your gallery and vault.
- **Notifications:** when Asset Lab is in the background, a soft chime, a flashing taskbar entry, and a desktop notification announce finished renders.

14 Keyboard shortcuts

KEYS	ACTION
`Ctrl` `Enter`	Generate (from the prompt box)
`←` `→`	Previous / next item in the lightbox

`Esc`	Close the lightbox
`Ctrl` `V`	Paste an image as a start frame / edit source (Video & Edit tabs)

15 Troubleshooting & FAQ

How long does my license last? What about updates?
Your license **never expires** — the app keeps working forever, on every computer you own. Your purchase includes **12 months of updates** (new models and features) counted from your purchase date, which Asset Lab reads automatically from your Gumroad receipt when you unlock. After the first year, the app keeps working exactly as-is; if you'd like another year of new-model updates, a **discounted renewal** will be available — entirely optional, and never a subscription.

"Invalid API key" when generating
The key was mistyped or deleted on fal's side. Open **License & API key** in the sidebar, re-paste, save — the app validates it immediately.

The guided Fal connect didn't pick up my key (Wayland)
Wayland blocks background apps from reading the clipboard. Copy the key, then **switch back to Asset Lab** — it is read the moment the window is in front. Pasting into the key field by hand always works as well.

The app won't start / sandbox error
Ubuntu 23.10+ restricts unprivileged user namespaces. Install the **.deb** (it configures this for you), or run `install-appimage.sh` next to the AppImage and accept the AppArmor profile.

Balance shows "—"
Your key is API-scoped. Create a new key with the **ADMIN** scope at fal.ai/dashboard/keys and swap it in. Generation works fine either way.

Generations fail instantly / "exhausted balance"
Out of credits — click **Buy Credits**. With an ADMIN-scoped key Asset Lab warns you at $5 and again at zero, so you should see this coming.

I topped up but the Balance is unchanged
Switch to another app and back — the balance refetches on focus. You can also click the balance in the sidebar to refresh it immediately. No restart is needed.

A start-frame video ignored my image
Make sure the image thumbnail shows under "Start frame" before generating — the note under it confirms i2v mode. Remember the prompt should describe *motion*, not re-describe the scene.

"Failed to download the file" on a lip-sync render
The model's own fetcher rejected one of the two inputs. Asset Lab now checks both before spending and names the culprit — if you still see this wording, the message lists the exact voice and face links it sent; open each in a browser. OmniHuman in particular refuses images over ~5 MB, which the app shrinks automatically; Flashtalk and sync-3 accept the same face unchanged.

One model keeps failing
Providers have occasional outages. Try another model — your prompt carries over.

Where exactly are my files?
`~/Pictures/Asset Lab` — every render, named with its date and model; items you've grouped into projects live in a subfolder per project. The Gallery is a view over this plus your generation history.

Is my API key safe?
It's stored only in the app's local storage on this computer and sent exclusively to fal.ai over HTTPS. Nobody else — including Coastal Creative Lab — ever sees it.

How do I update the app?

You'll receive a new .deb or AppImage. Close Asset Lab, then install the .deb over the old one (or replace the AppImage file). Your key, gallery, and vault carry over automatically.

My Director export has no transitions in the preview.
That's expected — ▶ Preview cut plays straight cuts so it starts instantly. Transitions and fades are rendered into the exported file; open it from the Gallery to see them.

Where do Director exports go?
Into the Gallery and your vault like every other render. Cuts with transitions or fades are rendered on your computer (free), so they exist only in the vault — use Show in file manager to reach the file.

A narration line runs past the end of its shot.
The narration chip turns red. Click **fit shot** to lengthen the shot, shorten the line, or lower its **at** offset. If you leave it, the line carries on over the next shot in the export — except on the last shot, where it stops when the film ends.
