



NUVAR

Indie filmmaking toolkit · storyboard, pre-viz, AI studio.

VERSION 1.12.2 · USER GUIDE

Welcome

NUVAR turns a screenplay into a storyboard, a storyboard into an animatic, and an animatic into a shoot plan — all in one app. Your projects, your account and your files stay on your disk — there is no NUVAR server and no subscription. The AI steps run on your own provider account over the internet; see “What leaves your computer”.

NEW IN 1.12.2

NUVAR is its own MCP server. Settings → AI Agents (MCP): while the app is open, Claude Code, Cursor or the Codex CLI can operate your film through the app's own engine — one copied line, nothing to install. See [The MCP server](#). And the 3D editor on macOS works again: "Three.js could not be loaded" came from a space in the app's folder name.

NEW IN 1.12.1

Video on kie.ai. Seedance 2.0 — text, image and reference to video — and Kling V3 Turbo now run through a kie.ai key as well, next to the Atlas Cloud video models: 17 video models across the two providers. A film whose default clip model lives on kie.ai stays there even when a model is named loosely ("Seedance"), and naming a provider ("kie Seedance", "Atlas Seedance") is honoured. A video-only kie.ai setup no longer stops at "no Atlas key".

From a full code audit of both editions: the console (CREATE) reported "0 done" after every batch and announced a picture as done that could not be written into the film file — a second click paid twice; several controls could be pressed twice for one paid run; adopting a rebuilt picture, an AI merge, written lines and matched narration reported success even when the film could not be saved; a clip with its own sound could end a fraction early in the cut; copying a project left the location plates behind; deck and treatment pages of a storyboard PDF became shots; a pre-release was offered as an update. All fixed. A licence is now valid through the end

of its last day (local time), with two days of slack for a wrong clock; one that fails verification no longer counts as valid.

Video prices used to be in two different units across providers, so "the cheapest" compared unlike things. One unit now (per clip at the default length) — which is why the cheapest-first order of the video list changed.

NEW IN 1.10.0

Location sheets **STUDIO** — build a room once from several angles and tag the shots that play there; the matching plate then rides along with each shot, so a scene stops looking like a different room in every frame. **Beat anchoring** ties several angles of the same story moment to one image, which is the fix for "suddenly there are two tables". **Shot headings** — a card can carry a few human words of its own instead of the opening of its generation prompt, and *Tools* ▾ → *Label shots from images* writes them for a whole board by looking at the stills that are already there; nothing is re-rendered. **CREATE mode** puts script, image, motion and film on one console for when the full board is more than you need. **kie.ai** joins Atlas Cloud, fal.ai and Freepik as an image provider. PDF exports are a fraction of their old size.

All price figures were removed from the interface. They came from the model catalogue, and those numbers are not what the providers actually bill. See [What things cost](#).

Two silent data-loss bugs were fixed: opening NUVAR twice let both windows write the same project file, and a freshly generated image could be overwritten moments later by a stale autosave. Both are closed.

Earlier: 1.7 brought the character-sheet generator, camera coverage and non-destructive colour grading. The full history is in the [CHANGELOG](#).

STANDARD OR STUDIO?

This guide covers both. The dividing line is: *STANDARD makes the assets, STUDIO makes the film*.

STANDARD — scripts, boards, characters and character sheets, images, the 3D previz, mood board, animatic, PDF and XML export. AI video generation is included too, but through CREATE mode rather than on the board.

STUDIO adds — the timeline editor and final-cut assembly, the in-app assistant, AI voices, lip-sync, compositing, per-shot video generation on the board itself, the  Context panel, location sheets and the hand-drawn storyboard-PDF import.

Anything marked **STUDIO** in this guide is simply not present in a STANDARD install — it is hidden, not disabled. The [MCP server](#) is neither: it lives next to the app, not inside it.

What's in the box



Storyboard manager

Acts, layers, shots — drag-reorderable (cards across layers, whole layers & acts too), with per-shot type, duration, audio, notes.



AI script analyzer

Paste a screenplay (txt / Fountain / PDF / DOCX)
— Claude breaks it into shots.



Image generation

up to 30 image models across Atlas Cloud, fal.ai, Freepik and kie.ai — 16 with an Atlas key alone — and 17 video models across Atlas Cloud and kie.ai. Reference images make a character much more likely to look the same from shot to shot — it is not a guarantee.



Character sheets

Build a person once from many angles and reuse them. Both editions, from the Characters Library.



Location sheets **STUDIO**

The same for a room. Tagged shots then get the plate that matches their camera angle. Only reachable by asking the assistant.



3D scene preview

Drop GLB models, block the camera and the actors, then either take the viewport as the frame or have a model repaint it photoreal.



Colour grading

Non-destructive, per shot and project-wide, with LUTs. Baked into the film and the graded-PNG export — not into the PDF.



Mood Board

Pinterest-style reference grid per project. Drag from Explorer.



Animatic → WebM / MP4 / ProRes



XML export

Render shots as a video with mixed audio. Plays in Premiere & Resolve.

Direct import into Premiere Pro and After Effects timelines.



CREATE mode

Script, image, motion, film on one console — for when the full board is more than the job needs.



Timeline & final cut **STUDIO**

Assemble the clips into a film, with voices, lip-sync and compositing.



In-app assistant **STUDIO**

Takes a scene from script to board on its own — and asks before anything that spends.

There is no production-schedule or budget-report view. Earlier versions had them; they were removed rather than left half-working, and this guide no longer describes them.

Install & First Run

First launch creates your local admin account and starts a 14-day trial. Everything you need is built in — no extra installs required.

System requirements

OS	MIN. VERSION	NOTES
Windows	10 (64-bit)	SmartScreen warning on first launch — click "More info" → "Run anyway"
macOS	11 Big Sur	Apple Silicon (arm64) and Intel (x64). Right-click → Open on first launch.
RAM	4 GB	8 GB+ recommended for large projects (500+ shots)
Disk	500 MB	Plus space for your project images

First-run wizard

1

Splash screen — click anywhere or wait 5 seconds to continue.

- 2 **Create admin account** — choose a username (≥ 3 chars), password (≥ 8 chars), and optional recovery email. These are stored encrypted on your machine only.
- 3 **Accept the EULA** — required by law. Click the link to read the full text.
- 4 **Trial starts automatically** — 14 days of all features unlocked. After the trial the app keeps working, but AI generation and PDF/XML export are refused until you enter a license key. Your projects, the board, the animatic and the JSON export keep working.
- 5 **Pick a project** — install the sample project (recommended for the first time), create a blank project, or link an existing folder.

 TIP — TRY THE SAMPLE PROJECT FIRST

Click "TRY SAMPLE PROJECT" on the projects page. You'll get an 8-shot 3-act demo that exercises every major feature, with notes on each card pointing to what to try next.

5-minute Quickstart

Go from script to storyboard with three AI-generated images in five minutes. Needs an Anthropic key (for the script breakdown) plus any one image provider — Atlas Cloud, fal.ai, Freepik or kie.ai (Settings → API Keys).

- 1 **New project** — Projects page → "+ NEW PROJECT" → name it, pick a color.
- 2 **Paste a screenplay** — AI ▾ → **Script** → **Storyboard** → paste 1-2 pages → click "ANALYZE". Claude returns a storyboard preview.
- 3 **Apply** — "APPLY TO PROJECT" (replace or append). The full storyboard appears.
- 4 **Tag a character** — Settings ▾ → **Characters Library** → "+ NEW CHARACTER" → describe (e.g. "Mid-30s, leather jacket, weathered") and add a photo or two. Back on the board, click a shot's chip-row to attach them.
- 5 **Generate images** — Click  on any card → pick a model → "GENERATE". Image arrives in ~10-30 seconds. For the character to carry across, pick a model tagged [REF](#).
- 6 **Switch to timeline** — the fourth view button in the middle of the header (or Ctrl 4). Press Space to play your animatic.
- 7 **Export** — Export ▾ → PDF, XML, animatic video or JSON.

Projects

Every project is a folder on your disk. You decide where it lives — usually `Documents/NUVAR/` on Windows or `~/Documents/NUVAR/` on Mac.

Project anatomy on disk

```
Documents/NUVAR/my-film/ ├── cron-project.json ← the only mandatory file ├── ai-generated/ ← AI stills | ├── ESTABLISHING-001.jpg | └ ... ├── ai-video/ ← AI clips ├── ai-voice/ ← voice-over (STUDIO) ├── characters/ ← character reference images ├── refs/ ← per-shot reference images ├── export/ ← the assembled film (STUDIO) ├── audio/ ← your own imported audio ├── scenes/ ← 3D scene JSON snapshots ├── mood-board/ ← reference images ├── drawings/ ← hand-drawn frames └ .nuvar-backups/ ← the 5 most recent autosaves, plus one per hour for 12 hours and one per day for 7 days
```

Project actions

- **New project** — picks a fresh folder, optionally pre-fills from a template.
- **Link existing folder** — bring back a project after machine change or restore.
- **Scan drives** — searches all your drives for `cron-project.json` files. Useful when you've forgotten where you put something.
- **Edit meta** — right-click a card to rename, recolor, or add a subtitle without opening it.
- **Unlink** — removes the project from your list but never touches the disk folder.

The Storyboard

Four views, one model. Switch views with the buttons in the header or `Ctrl 1-4`. Every view edits the same project — moving between them never loses work.

The four views

KEY	VIEW	BEST FOR
Ctrl 1	Grid	The default. See all shots in their acts & layers.

KEY	VIEW	BEST FOR
Ctrl 2	Center	Focus on one shot at a time, like a Lightbox.
Ctrl 3	Tree	Outline view — quick re-ordering at the act/layer level.
Ctrl 4	Timeline	Animatic — proportional widths reflect shot duration.

Aspect ratios

Eight presets in the ratio dropdown in the middle of the header (it shows the current one, e.g. 16:9): 16:9 Landscape · 9:16 Portrait · 2.35:1 Cinemascope · 2.39:1 Anamorphic · 1.85:1 Theatrical · 1.43:1 IMAX · 1:1 Square · 4:3 Classic. Changes all card thumbnails instantly without re-rendering, and the choice is remembered per project. The animatic video export and the PDF thumbnails both follow it.

Shots, Layers, Acts

Three-level structure mirrors real productions. Most projects use one layer per act — but you can have parallel layers (alt cuts, reference videos, narration) when needed.

Shot anatomy

Open a card's  for the full form.

Code / Name Your shot label — "ESTABLISHING", "INT_CAFE_03". Free text.	Heading A few words in plain language, shown on the card and in the PDF. Empty falls back to the description.
Type MAIN · GLITCH · HELMETCAM · FLASH. Visual distinction on the board and in the timeline.	Description This <i>is</i> the image prompt. Two helpers sit above it: DESCRIBE writes one from the attached image, REFINE rewrites the one you have.
Camera Shot Type ECU · CU · MCU · MS · MWS · WS · EWS · OTS · POV · AERIAL · MACRO · DUTCH · TOP · LOW ·	Status WIP · LOCKED · OLD · MISSING. "Generate All Missing" works off this.

HIGH, plus your own from Settings. SUGGEST picks one from the description.

Duration

Seconds — feeds the animatic, the video export and the PDF. "Lock length" keeps a shot at exactly this length when the final cut re-times the rest.

Colour grade

Non-destructive, per shot, applied in the final cut. Appears once the shot has an image.

Notes

Internal director notes. They print in the PDF; they are never sent to an image model.

Insert after

For a cutaway on a non-main layer: which main-line shot it cuts in behind. Only appears where it applies.

Audio

Per-shot audio file for animatic playback, with an inline player.

VO / Speaker

STUDIO

The spoken line and who says it — the input for AI voices and subtitles.

Characters

Tag chips. Their references are sent with image generation.

Path

The image file, relative to the project folder or absolute.

Per-shot actions

Each card has a row of buttons along its bottom edge:

-  — Generate AI image for this shot
-  — Generate AI video for this shot STUDIO
-  — Draw / paint over the image (Wacom-friendly canvas)
-  — Open the 3D scene editor
-  — Duplicate this shot into another layer
-  — Edit shot (modal with all fields)
-  — Delete shot (undoable)
-  — More actions

In narrow aspect ratios (9:16) the card is too slim for the whole row, so the middle buttons collapse into  . Nothing is lost —  opens the same menu you get by right-clicking a card.

The shot menu (or right-click)

This is where the less common actions live. Which entries appear depends on the shot: most of them need an image first.

- Edit shot... · Open in Lightbox (or double-click) · Generate AI Image
- **Rebuild image (edit)...** — give an instruction and let the model change the existing frame instead of inventing a new one. The right tool when text generation keeps reinventing the scene.
- **Add coverage angles...** — alternate angles locked to the frame you already have
- **First / last frame** — generate a video-start frame, generate a matching end frame, pick either from a file, or promote the current image to the end frame
- **Image versions...** / **Video versions...** — flip back through earlier generations and reuse one
- **Colour grade...** — plus copy/paste a grade between shots and export the graded PNG
- Draw / paint over · 3D Scene Editor · Duplicate to layer... · Export image as PNG... · Delete shot
- STUDIO Generate AI Video · Composite / BG replace...

Multi-select & bulk operations

shift-click a card to extend a range; ctrl-click to toggle one. The bottom action bar appears once anything is selected:

-  **CHARACTER...** — tag or untag a character across the whole selection
- → **MOVE TO...** /  **COPY TO...** — pick a target layer
- • **STATUS...** — WIP · LOCKED · OLD · MISSING
- ↓ **IMAGES** / ↓ **VIDEOS** — save the media of every selected shot into one folder
- × **DELETE**

There is no bulk *generate* button here. Batch generation runs from **AI** ▾ → **Generate All Missing**, which works off shot status rather than the selection.

Characters

Reference images and a written description are sent with every generation, which makes a character much more likely to look the same shot after shot — it is not a guarantee. Define them once, tag them on shots, and reference images get sent with every generation.

How character consistency works

- 1 Header → **Settings** ▾ → **Characters Library** (also under **AI** ▾) → "+ NEW CHARACTER"
- 2 Give them a name, color (for chip), and a written description (gets prepended to prompts)
- 3 Upload reference images — front, profile, full body if possible
- 4 On any shot, click the character-chip row to tag them
- 5 Next time you  generate, those references are uploaded to the image model

⚠ REFERENCE IMAGE RULES

NUVAR downscales references locally to 1280–1536 px JPEG before sending. How many go out depends on the model — from 3 (Qwen Image Edit Plus) up to 16 (GPT Image 2 · Edit on kie.ai). Models that take no reference at all, such as FLUX or Imagen 4, ignore them entirely; the picker marks the ones that do with  REF .

When **two or more** characters are tagged on the same shot, each one contributes at most three photos, so no single face dominates and gets painted onto the wrong person. A shot with only one character is exempt — more angles of one subject only sharpen the likeness. Pick the strongest, clearest shots; squinting, profile and motion blur all hurt consistency.

AI Script Analysis

Paste a screenplay or upload a file — Claude breaks it into shots structured as acts and layers. Edit before applying.

Supported formats

- Plain text (`.txt` , `.md`) — or just paste into the box
- Fountain (`.fountain` , `.spmd`) — the screenwriter-friendly markup
- PDF — text-only PDFs from Final Draft / Highland / WriterDuet. A *scanned* PDF has no text layer and is rejected with a clear message rather than silently producing nothing.
- Word (`.docx`) — works for un-formatted scripts too. Old `.doc` is not supported; save as `.docx` .

Director presets

Pick a director preset to bias Claude toward a particular sensibility. Twelve are built in: *Horror Template* · *Musical Template* · *Classic Drama* · *Action / Combat* · *True Crime* · *Anime / Stylized* · *Documentary / Verité* · *Music Video* · *Car Commercial* · *Advertising* — *Luxury* · *Advertising* — *Commercial* · *Neutral (Structure only)*. Each one carries a full cinematographer's brief — palette, lighting, lens, camera behaviour and how the layers are used — not just a genre label. *Neutral* is the one to pick when you want the script's own structure and no house style.

You can edit any of them or write your own in **Settings** → **Director Prompts**.

Applying the result

You always see a preview first, and can go ← **BACK** from it without changing anything. Three ways forward:

- **+ ADD TO CURRENT PROJECT** — appends the new acts after your existing ones. Use for revisions and for feeding in a script scene by scene.
- **+ NEW PROJECT** — creates a fresh folder and leaves the project you have open untouched. The continuity bible travels with it.
- **⚠ REPLACE CURRENT** — wipes the board and starts over. Undoable with `Ctrl Z`, but the red button is red for a reason.

Image Generation

up to 30 image models across Atlas Cloud, fal.ai, Freepik and kie.ai — 16 with an Atlas key alone — pick the right model for the shot, batch-generate the whole project, or refine one shot at a time.

Choosing a model

The model dropdown groups options by family (FLUX, Seedream, Nano Banana, GPT Image, Imagen, Luma, ...) and lists them cheapest-first inside each group. No price figure is shown — see [What things cost](#). Two tags matter:

- [REF](#) — a reference model. It **needs** at least one image: your character photos, a location angle, or the shot's own still. This is the family to use for continuity.
- [LORA](#) — needs a trained LoRA.
- Everything untagged is text-to-image: it reads your description and nothing else. Attached references are ignored, not blended in.

Some concrete picks from the current catalogue:

- **Seedream v5.0 Lite** — fast text-to-image, fine for establishing shots
- **FLUX Dev** — strong photorealism, text-to-image only
- **Seedream v5 Edit (Reference Mode)** — up to 14 references; the workhorse for keeping a face across a sequence
- **GPT Image 2** — the best of the batch at plates with legible text on them; its Edit variant takes up to 16 references
- **Nano Banana Pro** — available both on Atlas Cloud and kie.ai; the Edit Ultra variant takes up to 10 references

The catalogue moves as providers add and retire models. What is in the dropdown on your machine is the truth; this list is a snapshot of 1.12.1.

Video models

Clips are made from a shot's still (image-to-video), in CREATE's **03 MOTION** step — and, in STUDIO, straight from a card on the board. 17 video models across two providers: **Atlas Cloud** (Seedance 2.0 and 2.0 Fast, Kling 2.6 Pro, Veo 3.1, MiniMax H3, Hailuo 2.3 Pro) and **kie.ai** (Seedance 2.0 — text, image and reference to video — and Kling V3 Turbo). Atlas video runs on your Atlas image key; kie.ai video on your kie.ai key. A provider without a key is not offered.

- The list is **cheapest-first by clip** — every model's rate at its default length, in one unit — so "the cheapest" really is the top entry, whichever provider it belongs to.
- Each model brings its own lengths, resolutions and aspect list; the sheet only offers what that model accepts, and a still's own aspect is kept where the model allows it.
- The same family may exist on both providers (Seedance 2.0 on Atlas Cloud *and* kie.ai). Your film's default clip model decides which one a loose name ("Seedance") resolves to; naming the provider ("kie Seedance", "Atlas Seedance") always wins.

Batch generation

AI ▾ → Generate All Missing generates images for every shot whose status is "missing". A progress modal shows the current shot and a cancel button.

TIP — TEST FIRST

Generation runs on your own provider account (you bring your own API key). NUVAR no longer shows price estimates — provider prices change and can be wrong — so always start with a small test batch and check the real charges in your provider's own dashboard.

Shot Refinement (✨)

The ✨ button next to shot description rewrites it via Claude. Optionally give a direction like "more dramatic" or "tighter".

- Vague descriptions like "Bob enters" become evocative: "Mid-shot — Bob steps through the doorway, autumn light catching the dust"
- The direction field accepts anything — "more dramatic", "shorter", "in style of Kubrick", "show character motivation"
- The original description is replaced; use `Ctrl Z` to undo

Character Sheets

A character with three snapshots from a phone is a weak reference. A *sheet* is the same person rendered from every angle you need, all locked to one canonical front view — which is what a reference model can actually work from. Available in both editions.

Building one

- 1 **Settings** ▾ → **Characters Library** → ✨ **Generate Sheet**
- 2 Write the **character concept** in plain language. Optionally anchor it to a picture you already have: 📋 **Use storyboard shot...** takes a frame from the board, 📁 **Browse file...** takes a file from disk.
- 3 Tick the **views**. Six are on by default — front, 3/4 angle, side profile, back, full body, face close-up — plus two optional expressions, angry and smile. *Front* is always generated; it is the anchor everything else is locked to. Ten views is the ceiling.

4

Pick the **front model** (text-to-image) and the **reference model** that locks the other views. Then ✨ **Generate**.

The front is generated first — or, if you supplied an anchor image, taken from it unchanged. Every other view is then generated *against* that front, so they are versions of one person rather than several people who happen to match the same description.

What you get

- One image per view in [ai-generated/](#), saved as a new character with all of them as references.
- A **composite cover sheet** — every view laid out on one large image with labels, the way a production model sheet looks. It is deliberately *not* sent as a generation reference by default: a grid of pictures makes a weaker model paint the grid.

SENDING THE WHOLE SHEET INSTEAD OF SEPARATE PHOTOS

The per-shot generation dialog has a 🧑🏻 **Komplett-Sheet als Referenz** switch. With it on, each tagged character contributes exactly one image — their cover sheet — instead of several individual views. That eases the crowding on a multi-character shot, and premium reference models handle it well.

It is **off by default**, and the switch only appears once at least one tagged character actually has a cover sheet. If you don't see it, that is why.

Building a sheet needs an Anthropic key as well as an image key: Claude writes the locked look and the per-view prompts before any picture is made.

Colour Grading

A non-destructive look, per shot and project-wide. Nothing is written back into your image files — the grade is stored as numbers and applied when it matters.

Getting there

- Right-click a shot → **Colour grade...** (appears once the shot has an image)
- Or 🎨 Edit shot → **Open grading...**

- Right-click also offers **Paste grade here** and **Export graded PNG...** once there is something to paste or export

The panel

SECTION	CONTROLS
CORRECTION	Exposure · Contrast · Saturation · Temperature · Tint
LUT	A LUT picker, + .cube to add your own, and a <i>LUT amount</i> mix. Six LUTs ship with the app: bleach bypass, cold, noir, teal/orange, vintage, warm. Your imported ones are marked ★.
OVERLAYS	Grain · Vignette · Blur · Dirt · a VHS switch
LOOK TRANSFER	✨ Auto · 🎯 Match... · copy · paste · reset

✨ **Auto** measures the image's own brightness and colour balance and moves the sliders to even it out. It is arithmetic, not a model — it costs nothing and sends nothing. 🎯 **Match...** takes a reference picture and transfers its colour character onto the shot by building a small LUT from the two. A LUT you picked by hand always wins over a matched one.

Set as project look at the bottom applies the current settings underneath every shot in the project. Per-shot corrections then add on top; effects take whichever is stronger. This is the only place a project-wide look can be set.

⚠ WHERE THE GRADE ACTUALLY LANDS

It is baked into: the assembled film from the timeline **STUDIO**, and **Export graded PNG...** for a single frame.

It is shown but not baked into: the board, the lightbox and the animatic *preview* — these approximate it live in the app.

It is NOT in: the storyboard PDF, the plain *Export image as PNG* and bulk image export, or the exported animatic video. Those use the original files. If you need a graded still, use *Export graded PNG....*

Temperature and Tint are also the two sliders where the live preview and the final bake differ slightly — the preview approximates them with a colour wash, the export uses real colour-temperature filters.

3D Scene Preview

Lay out a shot in 3D space. Use the built-in GLB starter pack or drop your own models from Sketchfab / Mixamo / RenderHub.

Workflow

- 1 Click  on a shot card (or right-click → "3D Scene Editor") to open it
- 2 Drag a GLB from the library panel into the viewport
- 3 Click to select, use the gizmo to move/rotate/scale
- 4 Adjust camera (pan/orbit/zoom with mouse); the camera preset and focal length are in the top bar
- 5  **RENDER** → **SHOT** — writes the viewport as a PNG into the shot's image slot
- 6  → **REALISTIC** — keeps the 3D camera and composition as the lock, and has an AI model rebuild it as a photoreal frame. This is what the 3D editor is really for: block the shot yourself, then let the model paint it.
- 7  **SCENE** — saves the 3D scene only, so you can come back and refine it without touching the image

Mood Board

A Pinterest-style reference grid per project. Open with **Ctrl M** or from **Settings** ▾ → **Mood Board**.

- **Drag from Explorer / Finder** — drop straight onto the grid
- **Per-item tags** — "lighting reference", "wardrobe", "color palette"
- **Per-item notes** — "use as primary mood for act 1"
- **Search** — filters by tag or note text
- Items stay project-local — never shared between projects

CREATE mode

A deliberately small console for when the full board is more than the job needs: little text, big keys, four steps. It works on ordinary NUVAR projects — the same files, the same folders — so you can start here and open the same film in the full storyboard afterwards.

Getting in and out

Since 1.11.2 there is one switch for this, and it sits in the header of every view: **CREATE · WORKSPACE**. The half you are on is the raised plate; the other half is a button. Click it and you are on the other surface — same film, same files, nothing to import or export. It looks the same and sits in the same place in the console, on the board and on the Projects page, because it is one program.

Before 1.11.2 the way out of CREATE was a grey **STORYBOARD** key at the top right — a confusing name, since CREATE has a storyboard of its own — and the way in was a **■ CREATE** button on the Projects page. Both are gone; the switch replaces them.

ENGLISH OR GERMAN

CREATE opens in **English**. The **EN / DE** switch in the header changes every label, tooltip and message; the choice is remembered. Switching reloads the view, so finish or cancel a running generation first — the switch says so if you try.

The four steps

A dial on the left selects the step; a filmstrip along the bottom selects the shot. Three keys at the bottom change with the step.

STEP	KEYS	WHAT HAPPENS
01 SCRIPT	SCRIPT · CAST	Pick a script file; it is read and broken into shots. It creates a new project — this step cannot re-split the film you already have open, and it asks for your provider key first if none is stored. CAST opens the cast.
02 IMAGES	GENERATE · PLAY · CAST	GENERATE opens one sheet for the still — its text, its film look, the model, the aspect and the size — with the key that starts it at the bottom. CAST is where you add cast members, describe them and build their sheets.

STEP	KEYS	WHAT HAPPENS
03 MOTION	ANIMATE · PLAY · SHOT	ANIMATE opens one sheet for the clip — the motion, the camera, the model (17 video models across Atlas Cloud and kie.ai, cheapest-first — see Video models), the length, the resolution, the aspect and sound. SHOT holds the shot's own settings: caption, framing, length and who is in it.
04 EXPORT	EXPORT / BILDER · ORDNER · ANSEHEN	The big key does what your edition can do: in STUDIO it assembles the film, in STANDARD it writes all the stills out to a folder. The other two open the project folder and play the selected shot.

Opening a project that already has shots drops you straight into step 02.

The sheet behind step 02 — the one that carries the shot's text, its film look and the model — also has an **ASPECT** (BILDFORMAT) picker of its own, next to **SIZE** (GRÖSSE). Choose the format first: it filters the size list down to the sizes that really have that shape, so a size cannot quietly undo the format you just picked. What is on offer belongs to the model — one that knows only its own format, or only one size in that format, says so instead of opening an empty list.

The picture on either side of the dial opens full-screen on click. In full screen, ⌵ and ⌶ — or the arrow keys — page through the film, and a counter shows where you are. Esc closes it. The **STORYBOARD** button in the middle of the dial lays every shot out as one scrolling board.

The storyboard view

Two controls sit in its header. The **card size** slider steps through LARGE · MEDIUM · SMALL — + and - do the same. Beside it stands the **aspect ratio**: 16:9 · 21:9 · 2.35:1 · 2.39:1 · 1.85:1 · 4:3 · 1.43:1 · 1:1 · 9:16. It crops the *cards*, not the pictures — the files on disk are not touched, and going back to 16:9 shows the full frame again.

It is the same setting as the aspect-ratio picker in the header of the full board, not a second one: a film keeps its format across both views, so 9:16 set here is 9:16 there. It is a way of looking, not a change to the film, and it does not land in the undo history.

GENERATE (ERZEUGEN) and **ANIMATE** (BEWEGEN) sit in the foot of this view too, and they work on the selection — this is the place where you can see what is still missing, so it is the place to fill it. Tick the cards you want, or take **ALL**, then:

- **One shot selected** — exactly the sheet you would get on the console opens, prompt and all.
- **Several** — a batch sheet: the same dials once, a list of the shots that are about to run, and one shared text field. Its *emptiness* means something — leave it empty and every shot keeps its own text; whatever you type there is written into all of the listed shots.

For clips the length stands on **PER SHOT** (JE SHOT), so each shot keeps the duration it already has, and a camera plate reads **UNCHANGED** until you pick one. Selected shots that have no still are left

out of the run and the sheet says how many — a clip is made *from* the still.

Nothing starts before you have seen that sheet. Its heading and its button name the **number of shots** about to run — a count is something NUVAR can state truthfully, an amount is not. See [What things cost](#).

Cards are selected with a click, extended with **Shift**-click, and the rest of the foot moves, re-anchors or deletes what you have picked. **Esc** goes back to the console.

While something is running

A bar stays on screen for as long as CREATE is working: what is running, how far it has got, and **STOP** (STOPP). A batch counts itself off — 3 / 12 — and fills the bar. A single image has no total to count against, so the bar runs as a stripe instead of claiming a progress nobody knows.

⚠ STOP ONLY APPEARS WHERE SOMETHING CAN REALLY BE STOPPED

Not every job can be called back once a provider has it. The button is therefore shown only for the kinds that can actually be cancelled — image and clip runs, single or batch, character sheets and film assembly. If the bar carries no STOP, the job has to finish; a button that only pretends to cancel would be worse than none, because you would believe it and the run would carry on. Stopping a batch part-way keeps everything that was already finished.

Leaving CREATE while a run is going is refused rather than done quietly: the bar and its STOP would go with you while the job kept running.

Export targets

EINZELBILDER writes the stills out as separate files, and **STORYBOARD · PDF** builds the print-ready board — the same generator the full storyboard uses. Both work in every edition.

FILM · 16:9 and **HOCHFORMAT · 9:16** assemble the clips into one video file. **STUDIO** In STANDARD these two do not appear at all: a tile whose only job is to show what you are missing is not an offer.

What CREATE deliberately leaves out

No look presets, no prompt editing outside the sheets, no acts or layers, no drawing, 3D, mood board, grading, subtitles or timeline. All of that lives in the full storyboard, on the same project.

A project look set in the grading panel is **not** applied when a film is assembled from CREATE — only per-shot grades are. Assemble from the timeline if the project look matters.

Animatic Playback

The timeline view (**Ctrl 4**) plays your shots in sequence, respecting each shot's duration. Audio plays in sync if attached.

- **Space** — play / pause. This is the only keyboard shortcut in playback; everything else is on the bar at the bottom of the screen.
- **⏮** **⏭** — reset to the first shot · skip to the next one
- A scope selector limits playback to a single act instead of the whole board
- Click any shot on the timeline, or anywhere on the progress bar, to jump there
- Centered preview shows the current shot at full size during playback

Export Animatic Video (WebM / MP4 / ProRes)

Renders your animatic to a real video file with per-shot audio mixed in. Frames are assembled directly by the built-in encoder — much faster than real time, and you always choose where to save.

Formats

FORMAT	CODEC	USE FOR
WebM	VP9 + Opus	Smallest file · browsers, VLC, quick sharing
MP4	H.264 + AAC	Universal — every player, YouTube, NLEs
MOV	ProRes 422 HQ + PCM	Editing master — Premiere, Resolve, Final Cut

The **Quality** selector applies to MP4 & WebM; ProRes is always full quality.

Layers to include

A per-act checklist lets you render exactly the layers you want — with **All / Main only / None** shortcuts and a live shot + runtime counter. Default: the main (first) layer of every act. If you include an act's main layer plus other layers, anchored inserts/cutaways are woven into the main line just like in playback.

Resolution presets

PRESET	LONG EDGE	USE FOR
720p	1280	Quick previews · sharing
1080p	1920	Client review · social
1440p	2560	Conform pre-edit · presentations

The aspect ratio is read from the current storyboard view — change it via the top-right buttons before exporting. The preset is the *long* edge, so a 9:16 board at 1080p exports as 1080×1920.

Export PDF

Export ▾ → Export as PDF. Produces an A4 landscape document: a title cover with the project name and stats, then four shots per page with the image, its heading, act and layer, duration, status, description and notes. A page footer carries an AI-generated-content notice.

The thumbnails follow the board's aspect ratio — a 9:16 or 2.39:1 board is laid out to match rather than cropped to landscape.

The one option is a watermark. The export dialog offers a single free-text field ("DRAFT v3.2", "CONFIDENTIAL", ...) which is printed diagonally across every page at low opacity. There is no paper-size or orientation choice, and descriptions are always included.

During the trial the PDF additionally carries a "TRIAL · NUVAR" stamp. Once the trial has ended, PDF and XML export are *refused* rather than watermarked — as is AI generation. A licensed build exports clean.

Export to Premiere / After Effects

Generates a Final Cut Pro 7 XML — accepted as import by Premiere Pro and After Effects. Shots come in as a sequence at the right durations.

- **Premiere** — File → Import → pick the .xml. Each shot becomes a clip with the right offset.
- **After Effects** — File → Import → pick the .xml. Creates a comp with layers.
- Images are referenced by absolute path — keep the project folder intact when sharing the XML.

What things cost

NUVAR shows no money amounts anywhere. Before anything is generated it tells you exactly how many images, clips or spoken lines it is about to make — and nothing happens until you confirm.

- Models are listed cheapest-first, so "the cheapest" is simply the top of the list. No figure is shown. Video models are ranked by their rate per clip at the default length — one unit across providers, so the order compares like with like (before 1.12.1 it did not).
- The real charges live in your own provider's dashboard — Atlas Cloud, fal.ai, Freepik, kie.ai, Anthropic, OpenAI or ElevenLabs. You bring your own keys, so NUVAR never sees a bill.

Earlier versions displayed a per-image estimate. It came from the providers' published rates, and those are not what they actually charge — one advertised a hundredth of its real price. An estimate you plan against and that turns out wrong is worse than no estimate, so it was removed rather than patched.

The Assistant STUDIO

A chat panel that can actually operate the app: read your project, write prompts, tag characters, generate images and clips, speak the lines and assemble the film. It runs on your own Anthropic or OpenAI key, and it never spends money at a provider without asking you first.

Opening it

The  **Agent** button, available on three pages: the storyboard, the Projects page and Settings. Esc closes it. Pick the model in the panel header — Claude Sonnet, Opus or Haiku, or GPT-4.1 / 4o. The starter chips (*What now?* · *Images only* · *Video only* · *Read the script (then stop)* · *Whole film* · *Character reference*) are just pre-written messages; you can type anything instead.

What it can do without spending anything

Reading your project, and editing it. It can rewrite a shot's prompt or heading, write the motion description for a clip, tag characters, set the camera type, group shots into a beat, set the director look, edit the story bible, assign a voice to a speaker, attach a screenplay, delete or import characters, and assemble a film locally with the bundled encoder. It can also *look at* your images and tell you what is wrong with them — that needs a Claude model, since the GPT path here has no vision.

Added since the 1.11.2 release. Four things it could not do before, all free:

- **Write the spoken lines.** It could already *speak* a line, but not write one — you had to type every line into the shot editor yourself. Now you can say "write the narration for act one" and it fills the lines in. If a shot already has a recording, changing its line does not change the recording: the film keeps the old words until the voice-over for that shot is generated again, and the assistant tells you which shots that affects.
- **Create and edit characters without paying.** Until now a character could only come into existence through the paid sheet generator, so correcting a name cost money. Describing someone in words is now enough.
- **Turn a picture that is already in the film into a character reference** — a shot's still, a picture you dropped into the chat, or a file in the project folder. Pay for a rendered sheet only when no usable likeness exists yet.
- **Measure what it produced.** After generating clips or voice-overs it can check what actually landed on disk: is the file there, does the clip have a picture and a sound track, how long is it, and — the important one — is the recording longer than the time the cut plans for that line. If it is, the sentence is cut off in the finished film, and nothing else in the program would have told you.

A camera move drawn into a storyboard PDF (the "Camera move" column) now reaches the clip as well. Before, it was read in and shown on the card, but thrown away when the shot was animated.

Since 1.12.1. A shot code is accepted in any case ("s05" finds S05) on every tool, not just some; a list of named shots always wins over a scope word, so "tag S05" can no longer re-tag the whole film; the film look sent along with a prompt is no longer dropped when the code case differs; a talking-head clip does not speak a line that already has its own recording; and pressing Cancel on a running character sheet can no longer be undone by a second run. When you name a video provider — "make it on kie", "use Atlas Seedance" — that provider is used.

What costs money — and the gate in front of it

Exactly eight things spend at a provider: generating images, videos, voice-overs, a character sheet, a location sheet, lip-synced clips, reading a script, and reading a storyboard PDF. Every one of them stops and asks first, in the chat, with this wording:

The agent wants to generate 12 *images*. This costs money at your provider. Nothing is generated unless you confirm. [ Confirm] [ Skip]

The number is the exact count, and it is the honest one — it is what will actually be made. Script and PDF imports say instead how many AI calls or pages they will read, and which file. Nothing happens until you click  **Confirm**.

ABOUT PRICES

NUVAR itself displays no money figure anywhere, and the assistant is instructed never to name one — including from its own memory of provider pricing. But it is a language model writing free text: that instruction is a rule it follows, not a filter that removes numbers. If a figure ever appears in its replies, do not budget against it. The only true number is in your provider's account.

Showing it pictures

Drag images onto the panel, paste them from the clipboard, or use the  button — up to eight per message. They are downscaled before sending. Attached pictures are also written into the project so that a character or location sheet can be anchored to those exact pixels rather than to a description of them. There is a  button for dictating instead of typing.

Memory

The conversation is stored inside the project folder and survives closing the app, so you can pick up where you left off. The last 80 messages are kept. Attached *images* are not kept — only the conversation about them. Switching to another project switches to that project's own conversation.  starts a fresh one.

The panel calls itself a preview, and that is fair: it is the newest part of the app and the one that changes most between releases.

The MCP server — driving NUVAR from outside

NUVAR is its own server for the *Model Context Protocol*: the open standard through which AI assistants such as Claude Code, Claude Desktop, Cursor and the Codex CLI operate other programs. Switch it on in **Settings** → **AI Agents (MCP)**, connect your assistant with one copied line, and it can open your film, read every shot, write prompts, tag the cast, set the look — and, behind a confirmation and a spend cap, generate the pictures, clips and voices through NUVAR's own engine. Same references, same prompt, same files as a click in the app.

WHAT IT IS — AND WHAT IT IS NOT

While NUVAR is open, it listens on this computer only (`127.0.0.1`), behind a token that is yours alone. It is **not** an edition and not a cloud service — nothing leaves your machine except what the paid tools send to your own provider, exactly as the app does. Close NUVAR and the server is gone with it.

There is nothing to install for Claude Code, Cursor or the Codex CLI. Claude Desktop only launches local programs, so it needs Node.js for the small *mcp-remote* bridge; the copied configuration includes it.

What the assistant can do through it

63 tools, in three groups. **Free, read**: the project list, status, every shot with its text, camera, cast and location, the characters and locations, the model catalogue, versions of a shot, what is selected on the board, and an analysis that names the single best next step. **Free, write**: the same field-setters the in-app assistant has — prompt, motion, title, beat, camera, location, cast, duration, blocking, spoken line, references, the film look, cut template, colour grade, director preset, image model, story bible — plus adding, deleting and reordering shots, creating and editing characters, casting voices, importing a shot list, a folder of images, or the cast and locations of another film, and exporting the character sheets to files. **Paid**: reading a script or a drawn storyboard PDF into shots, generating images, clips, voice-overs, character and location sheets, lip-synced clips, editing a shot's picture by instruction, and assembling the film with the bundled encoder.

Three guards, all on the server side

- **Every paid tool is a dry run by default.** Without an explicit `confirm` it returns the plan — how many shots, which model — and spends nothing. The assistant has to ask you before it sets that flag.
- **A session spend cap**, set when you register the server (`NUVAR_MCP_SPEND_CAP` , default 5, 0 = off). It stops the loop the moment the *next* shot would cross the line, and says what it skipped. Image, clip, voice, edit, sheet and lip-sync generation count against it; script and PDF reading spend tokens on your own Anthropic/OpenAI account and are disclosed but not metered.
- **The app's own gates run first** — the real licence state for anything paid, the real edition for video and voice. The server cannot do more than your NUVAR can.

Scope defaults to *missing* everywhere: "generate the images" means the ones that are still missing, never a re-bill of work that exists. The film you have open on screen is exactly the one the agent can work on — and it reloads by itself when the agent changes it, so what you see is what the agent did.

Your API keys never reach the assistant. The app calls the providers itself and hands back results — the assistant only ever sees those.

Connecting it

Open **Settings** → **AI Agents (MCP)**, make sure *Serve MCP while NUVAR is open* is ticked, and press the button for your client. The ready-made configuration lands on your clipboard — with the address and your token already in it. Claude Code takes a one-liner in a terminal:

```
claude mcp add --transport http nuvar "http://127.0.0.1:41777/mcp" --header "Authorization: Bearer <your token>"
```

Cursor and Claude Desktop take a `mcpServers` block in their settings file, the Codex CLI a `[mcp_servers.nuvar]` block in its `config.toml` — each button copies the right shape. **New token** invalidates every connected client at once; copy a fresh configuration afterwards. Then talk to it the way you talk to the in-app assistant:

```
"Open BOSCH – what is missing?" → it opens the project and reports the gaps "What would the missing images cost?" → a dry run: 12 images, which model, nothing spent "Do it." → the same 12 images, through the app's own engine
```

One honest gap: the server reaches NUVAR's *engine*, not its *windows*. It cannot press a button in a dialog that is open on your screen, and the film it assembles is the app's hard-cut export — no crossfades or burned subtitles; those stay in the Editing window. A tool your edition does not have is simply not listed.

Location Sheets

STUDIO

What a character sheet does for a person, a location sheet does for a room. Build the space once from several angles, tag the shots that play there, and each shot gets the plate that fits its camera — so a scene stops looking like a different room in every frame.

⚠ THERE ARE NO BUTTONS FOR THIS

Location sheets have no panel, no library and no field in the shot dialog. The only way to build one, or to tag a shot with one, is to **ask the assistant** — "build a location sheet for the kitchen", "put shots 4 to 9 in the kitchen". Everything below happens through that conversation.

What gets built

Seven angles by default: *establishing* · *reverse* · *corner-left* · *corner-right* · *high-angle* · *entrance* · *detail*. As with a character sheet, the establishing view comes first — or from a photo you supply — and every other angle is generated locked to it. The room is described and rendered **empty**; people come from the shot itself.

Which angle a shot gets

Exactly one room plate travels with a shot, picked from the shot's camera type:

CAMERA TYPE	PLATE SENT
EWS · WS · WIDE · MASTER	establishing
AERIAL	high angle
MWS · MS	corner-left
OTS · REVERSE	reverse
POV	reverse, then the entrance
CU · ECU · MACRO · INSERT	detail
DUTCH	corner-left

If the shot has no camera type, or the location has fewer angles than that, it falls back down the list starting at the establishing plate. It is capped to **one** slot on purpose: the room must always ride along, but it must never crowd out the faces on a multi-character shot.

The reference order that goes out with a shot is: the beat anchor, then the shot's own references, then the room plate, then the character photos — cut off at whatever the chosen model accepts.

Beats — several angles of one moment

A location keeps the room identical; it does not keep the furniture and the people where they were. That is what a **beat** is for: two to four shots that cover the same story moment from different angles. When one of them has an image, the others are generated locked to it — same seat, same table, only the camera moves.

Beats are created for you by the script import: its **Angle coverage** setting offers *Single angle*, *Varied* (two angles on long beats) and *Full coverage* (shot/reverse, 2–4 angles). The assistant can also group existing shots into a beat, and the timeline can merge a beat's separate voice lines back into one continuous take.

Timeline & Final Cut STUDIO

A real editor for the clips you generated. Open it with the **Editing** button in the header — not to be confused with the fourth *view* button, which is the animatic playback on the board.

If the project has more than one filled layer you are asked **WHICH LAYERS GO INTO THE CUT?** first, with *all* / *none* shortcuts. One layer, and it opens straight through.

Getting around

Transport: ⏮ ⏪ ⏩ ⏭ ⏮ ⏭ , with Space to play, ⏪ / ⏩ for the previous and next clip, Home for the start. Tools: 🖱 selection (v), ✂ blade (c) which cuts a video or audio clip where you click, and 📏 snapping. ⏴ and ⏵ set In and Out — the export renders only what is between them. Zoom with - + or 🔍 Fit.

The four menus

-  **Timeline** ▾ — a cut *Style* and ⚡ **Auto Cut**, plus *Close gaps* — *keep my cuts*, *Reset to storyboard*, *Bring back removed shots* and *Unlock all clip lengths*.
-  **Tracks** ▾ — add a video track (V2 and up), reorder or delete tracks, edit only the selected one, and per-track  hide /  lock /  solo in the gutter.
-  **Audio** ▾ —  *Narrate all*, *Rebuild audio lane*,  *Merge beat VO*, *Fit all to VO*, and  *Load audio file...* for your own music or atmos. Dragging an audio file onto the timeline works too.
-  **Subs** ▾ — auto-generate subtitles, add one by hand, and choose whether to burn them in or mux them as a soft track. English and German.

The  ▾ menu holds two switches for the render: generate any missing voice-over on export, and normalise loudness (on by default).

Editing a clip

Drag a clip to move it, drag its right edge to trim. Select one and the panel on the side gives you **Length**, **Speed** ($\frac{1}{2}\times$ $1\times$ $2\times$ $4\times$ or a free number), a **Lock length** switch, *Fit to VO*, and an **opacity envelope** with fade-in and fade-out.  **Time-Remap** is the After-Effects-style curve: drag a point sideways to change *when*, right-drag to change *which source frame* — that is how you build a ramp or a freeze.

Trimming ripples automatically. There is no ripple switch to find: changing a clip's length moves every later clip on the main video track and the base audio lane by the same amount, so you never leave a hole. Overlay tracks are deliberately left alone.

An overlay clip (V2 and up) additionally gets **Track**, **Size** (100 % = full frame, more = punch in, less = picture-in-picture) and **Pos X / Pos Y** — or just drag it in the preview.

THERE ARE NO PER-CUT TRANSITIONS

The only thing that sets a transition is ⚡ **Auto Cut**, and it sets *one* for the whole film from the chosen style. You cannot put a crossfade between two particular clips. The fade controls on a clip are an opacity envelope on that clip alone — useful for coming up from black, not for dissolving into the next shot.

The audio lanes

LANE	WHAT IS ON IT
AUDIO	The main voice lane: one clip per shot's spoken line, plus narration blocks and any audio file you imported. Fully editable — move, cut, fade, set gain.
AUDIO 2	The same, for lines belonging to overlay shots. Only appears once you have overlays.
CLIP	Not a real track — a read-out of the sound baked into each video clip, ducked under the voice. Clicking mutes that shot's clip audio.
SUB	Subtitle cues.

A clip you have time-remapped keeps its own sound in the preview but not in the exported film. If that sound matters, put it on the audio lane.

Auto Cut

Sends every clip on the main track — its length, its spoken line, its speaker — to Claude with the chosen style, then rewrites the durations, locks them, and re-flows the track gaplessly from zero. It does not touch your audio, your overlays or your images, and `ctrl z` undoes it in one step. Treat it as a first pass to argue with, not a verdict.

Exporting the film

↓ **Export...** asks where to save and writes an MP4. The size comes from the project's aspect ratio (16:9 → 1920×1080, 9:16 → 1080×1920, 1:1 → 1080×1080, 4:3 → 1440×1080) at 24 fps — there is no resolution picker here. Subtitles can be burned in and/or muxed as a soft track, with `.srt` files written alongside. Progress runs on the button itself; the toast at the end names the full path, and says so explicitly if the film came out without sound. If a track is soloed you are warned before rendering, because solo affects the export and not just the preview.

Voices & Lip-sync STUDIO

What a shot needs first

Two fields in the shot editor, both STUDIO-only: **VO / Subtitle** — the line that is spoken — and **Speaker**. A shot with an empty VO field is never a voice target. You can also type the line straight onto an audio clip in the timeline, or have `↵ Sync from script` write it from your screenplay.

Which voice a line gets

- 1 An empty speaker, or one that reads like NARRATOR / ERZÄHLER / V.O. / OFF, gets the **narrator voice**.
- 2 A speaker whose name matches a character gets **that character's voice**. The match tolerates the usual script annotations, so "KAYA (V.O.)" still finds Kaya.
- 3 An unrecognised name gets a consistent voice picked for it, so at least it stays the same person throughout.

Set a character's voice in the **Characters Library** → edit the character → **ElevenLabs Voice**.

Override a single line from the timeline's audio panel. The narrator voice is set by asking the assistant to assign a voice to "NARRATOR".

Generating them

-  **Audio** ▾ →  **Narrate all** — makes every missing line and lays the narration out at its real spoken length
- Select an audio clip →  **Regenerate voice** — re-records that one line and re-fits the clip
-  ▾ → **Generate missing VO on export** — catch the stragglers at render time
- Or ask the assistant

Needs an **ElevenLabs** key. Finished audio is cached against the text, so re-generating an unchanged line costs nothing — but editing a single word means a new recording. Numbers and clock times are written out into speakable words first.

Lip-sync

ONLY THROUGH THE ASSISTANT

There is no button or menu entry for lip-sync anywhere in the app. Ask the assistant — "lip-sync shots 4 to 9" — and confirm when it asks.

A shot needs all three of: a spoken line, its **already generated** voice-over file, and a still image. The still is the face; the voice drives the mouth. Output is 720p by default, 480p optionally, and you can add a short note about expression or posture.

The clip that was on the shot before is **kept as a version**, and so is the new one, so you can always flip back. The still and the voice file are untouched. It runs on **Atlas Cloud** — not ElevenLabs — so

that is the key it needs. It is a talking-head model: it earns its keep on close and medium shots of one person speaking, and struggles with anything wider.

Compositing / BG Replace STUDIO

Put real footage over a generated background. Right-click a shot → **Composite / BG replace...** — that is the only way in.

The core of it costs nothing. The chroma key runs locally with the bundled encoder; no request leaves your machine. Three optional buttons in the same dialog *are* paid, and are marked as such below.

- 1 Foreground clip** — your footage. Pick a file, use the shot's own generated clip, or supply a clip that already carries an alpha channel (ProRes 4444 `.mov` or VP9 `.webm`). ☐ **AI Matte** cuts a subject out without a greenscreen — *this one is paid and needs a `fal.ai` key*.  `Matte` shows the cutout in black and white so you can judge it.
- 2 Key colour** — green, blue, or eyedropper a colour straight off the frame.
- 3 Key strength** — *Similarity* and *Edge softness*, plus spill reduction to kill the green tint along the edges. This is where a mediocre key becomes a good one; move them slowly.
- 4 Background** — the shot's own AI video, any file, a still from a board layer named `BACKGROUNDS`, or generate one right here from a prompt (*paid*). Foreground audio is kept by default.

Render composite writes an MP4 into the project's `comp/` folder and attaches it to the shot, taking the size and length from your foreground clip. If the shot had no still, one is pulled from the new clip. The background you chose is remembered for next time.

A note in the dialog says AI matting is "coming next". It is already there — the note is out of date, and the button above it works.

Storyboard-PDF Import STUDIO

A storyboard someone drew by hand, as a PDF, becomes a NUVAR project — every panel cropped out and kept both as the shot's picture *and* as the composition reference the AI must follow.

Running it

- 1 **AI** ▾ → **Script** → **Storyboard**, then the **STORYBOARD-PDF (Skizzen)** → button in that dialog
- 2 Pick the PDF
- 3 Pick an **empty folder** for the new project — this always creates a new project, it never appends to the one you have open

Each page is rendered at high resolution and read by Claude vision, one call per page. It looks for the layout a drawn board almost always has — one shot per row, top to bottom: the drawing, the shot number, the action, and often a camera-move note. Production columns (VFX, SFX, backgrounds, animals, doubles) are ignored on purpose. A page it cannot read is skipped with a message rather than aborting the whole import.

What each shot ends up with

- **The cropped drawing as its image** — so the board is readable at a glance immediately
- **The same drawing as reference #1**, with a fixed instruction attached: follow this composition, framing, camera angle and the placement of everyone in it — but redraw it photorealistically in the film's look, do not copy the pencil style
- The shot label as its **code**, the caption as its **description**, the camera framing if the page named one, and the camera move in the notes
- A **scene number** read off the label — `73_01` becomes scene 73

A panel whose crop fails still becomes a shot, with its caption text and no picture; the closing message tells you how many.

Afterwards

When it finishes you land on the Projects page — open the new project from there. Because every card has a picture, they show as *WIP*, not *missing*. Generate on a reference model and the AI follows the drawing.

You can then **attach the screenplay**, which pairs the scene numbers on the drawings with the scene numbers in the script, so prompts can be written from the actual dialogue instead of a caption. Like lip-sync, this has no button — ask the assistant to attach the screenplay. It will tell you if some drawings carry scene numbers the script does not have, which usually means the wrong file or an excerpt.

There is no page limit — every page costs one vision call, so a long board costs proportionally more. The confirmation tells you the page count before it starts.

Settings

Settings live on a separate page accessible via the gear icon. The left sidebar lets you jump between sections; the search box (`Ctrl K`) filters across all sections.

The sections, in sidebar order

Account

- **API Keys** — Anthropic, OpenAI, Atlas Cloud, fal.ai, Freepik, kie.ai and, in STUDIO, ElevenLabs. Encrypted by Windows DPAPI (macOS Keychain), tied to your OS user account — copying the files to another machine makes them unreadable.
- **License** — Activate / deactivate your key. Trial status.

Workflow

- **Default Settings** — Default image model and size, and which Claude model does the script analysis.
- **Director Prompts** — Edit the twelve built-in presets or write your own.
- **Storage Paths** — Move the projects folder or the 3D/GLB library to a different drive.
- **Camera Shot Types** — Add your own shot abbreviations alongside CU / MS / WS / OTS / POV.

App

- **Keyboard Shortcuts** — the same list the `?` overlay shows.
- **About** — Version, bug report, EULA, Impressum, AGB.
- **Maintenance** — Self-registration toggle, crash logs, danger zone (full reset).

What leaves your computer

NUVAR runs on your machine and your projects stay on your disk. Generation, however, is done by AI providers over the internet — with **your own API keys**. This page says exactly what is sent, to whom, and when, so you can answer that question for a client before you sign an NDA rather than after.

The short version

- **There is no NUVAR server.** We do not host your projects, your scripts or your images, and we never receive them. There is nothing for us to hand over, lose or be breached.
- **You are the customer of the AI provider, not us.** Requests go out under your key, are billed to your account, and are governed by your agreement with that provider.
- **Nothing you make is public.** No other NUVAR user can see any part of your work.
- **Nothing is sent until you ask for it.** Writing, importing, cutting and exporting are entirely local. Only the four actions in the table below reach the internet.

Exactly what is sent

WHEN YOU...	WHAT IS SENT	WHERE IT GOES
Run <i>Script</i> → <i>Storyboard</i>	The full script text you loaded, plus your director preset and context notes	Anthropic or OpenAI — whichever the model you picked belongs to (the dialog names it before you send)
Generate an image	The shot description, style prefix and any reference images attached to that shot	Your image provider — Atlas Cloud, fal.ai, Freepik or kie.ai
Generate a video clip	The shot prompt and its still / reference frames	Your video provider — Atlas Cloud or kie.ai
Generate a voice-over	The spoken line only	ElevenLabs

What is never sent

- Your project file, folder structure, file names or media that you did not explicitly generate from.
- Your timeline, cut, notes, schedule or budget.
- Any analytics or telemetry. NUVAR does not phone home. The only other outbound connection is the update check, which asks GitHub for a version number and sends nothing about you.

Training and retention

Anthropic and OpenAI both state that **API** input is not used to train their models — this is the important difference from their consumer chat products. Both retain requests for a limited period for abuse monitoring, and both offer zero-retention arrangements for business customers. Because you use your own account, these terms are between you and them: check the current version of your provider's terms, and ask them for a data-processing agreement if your work requires one.

If you work under an NDA

Sending a script to an AI provider means engaging a subprocessor. Many NDAs allow this with a data-processing agreement in place; some do not allow it at all. Two practical options:

- Ask your provider for a data-processing agreement and, if the work is EU-based, confirm the transfer mechanism for processing in the US.
- Or use NUVAR without the generative steps: importing, boarding, cutting, the animatic, PDF and video export all work fully offline. Only the four rows in the table above need the internet.

This page describes how the software behaves — it is not legal advice. If a client's requirements are strict, have the agreement checked by someone qualified.

License & Trial

NUVAR is commercial software. The 14-day trial gives you the full feature set, with a stamp on the PDF. When it runs out, the app keeps opening and your work stays fully editable — but everything that costs money or leaves the building stops until you activate a key.

States

STATE	WHAT IT MEANS	AI GENERATION	PDF / XML EXPORT
trial	First 14 days, full features	Works	Works, stamped "TRIAL · NUVAR"
trial-expired	14+ days, no license activated	Refused	Refused
licensed	Valid key activated	Works	Clean
expired	License had an expiry date that passed	Refused	Refused

"Refused" means a plain message — *"Your NUVAR trial has ended — activate a license in Settings → License to keep generating."* — not a silent failure and not a degraded result. Opening projects, boarding, re-ordering, the mood board, the animatic and the JSON export are never gated.

Activate a key

Settings → License → paste the key starting with `NUVAR-...` → click ACTIVATE. The signature is verified offline against a built-in public key — no internet round-trip.

PRIVACY NOTE

License activation is fully offline. We never see your key. Even our updater (when enabled) only checks for new versions; nothing about your usage is sent anywhere.

Keyboard Shortcuts

Press  anywhere in the app for the full cheatsheet overlay.

Universal

SHORTCUT	ACTION
 + 	Save immediately
 +  / 	Undo / Redo
 + 	Command palette
 + 	Find / Replace
 + 	Mood Board
 + 	Zoom UI
	Cheatsheet

Storyboard

SHORTCUT	ACTION
 + 	Switch view (grid / center / tree / timeline)

SHORTCUT	ACTION
<code>Ctrl</code> + <code>A</code>	Select all shots
<code>Space</code>	Play / pause animatic (timeline view only)
<code>Shift</code> + <code>Click</code>	Range select
<code>Ctrl</code> + <code>Click</code>	Toggle single shot
<code>N</code>	New shot in the selected shot's layer — needs a shot selected first
<code>Del</code>	Delete selected shot(s) — <code>Backspace</code> works too, for Mac keyboards
<code>Alt</code> + <code>↔</code>	Move shot left / right <i>within</i> its layer
<code>Alt</code> + <code>↑↓</code>	Move shot to the layer above / below
<code>Esc</code>	Clear selection / close modal

Alt-arrow reordering acts on exactly one selected shot. During animatic playback `Space` is the only key — use the `⏮` / `⏭` buttons on the playback bar for the rest.

Lightbox

SHORTCUT	ACTION
Double-click <code>a card</code>	Open the lightbox (and reset zoom once inside)
<code>←</code> / <code>→</code>	Previous / next shot
<code>Scroll</code>	Zoom under the mouse pointer
<code>Drag</code>	Pan when zoomed in

Troubleshooting

First step for anything weird: Settings → About → "Report a Bug". Attaches the latest crash log automatically.

IMAGE GENERATION FAILS

- Check the key for the provider that model belongs to — Settings → API Keys → "TEST KEY". A Seedream model needs the Atlas Cloud key, a `fal:` model the fal.ai key, and so on.
- Check the credit balance in that provider's own dashboard
- If the message mentions the trial having ended, activate a key: after expiry all generation is refused
- Check the crash logs: Settings → Maintenance → Crash Logs

"PROJECT IS ALREADY OPEN IN ANOTHER WINDOW"

- You have a second NUVAR instance with the same project open — close it
- Or a crash left a stale lock — wait 90 seconds, or use "FORCE-OPEN" from the warning

LICENSE KEY WON'T ACTIVATE

- Make sure you copied the entire string including the `NUVAR-` prefix
- Check the key wasn't expired or revoked — contact support
- On Linux without libsecret, encrypted storage isn't available — NUVAR refuses to persist plaintext for security

APP WON'T START ON MAC ("DAMAGED")

- Right-click the app → Open → Open (one-time Gatekeeper approval)
- Or terminal: `xattr -cr /Applications/NUVAR.app`

LOST A PROJECT

- Projects page → "SCAN DRIVES" — searches all drives for `cron-project.json`
- The `.nuvar-backups/` folder inside each project holds far more than the "last 5" the Tools menu offers: the 5 most recent saves, plus one per hour for the last 12 hours, plus one per day for the last 7 days. A mistake you only notice the next morning is still recoverable.
- If `cron-project.json` itself is unreadable — a quarantined file, a half-finished cloud sync — NUVAR restores the newest backup that still parses, rather than opening an empty project.