

# Film Movie *Studio*

## User Manual

Covers version 3.3.3

2026-07-05 · [filmtovor.ai/docs](https://filmtovor.ai/docs)

# Contents

1. Getting Started
2. Launchpad & Projects
3. The Project Workspace
4. Canvas & Cards
5. Chat & the Agent
6. Skills & the Hub
7. Skills Reference
8. Timeline Editor
9. Cast & Asset Libraries
10. Settings & AI Providers
11. Advanced: MCP, Telegram & More
12. Troubleshooting & FAQ

# Getting Started

Install FilmMovieStudio, activate your license and connect an AI provider — three ways to run.

Welcome to FilmMovieStudio — a desktop studio for turning ideas into finished media with AI. You work on a freeform canvas, talk to an AI agent in the chat panel, and everything you generate — scripts, images, video clips, voiceovers and music — lands as cards you can arrange, remix and drag onto a timeline. When you're happy, the built-in timeline renders a real MP4 you can share anywhere.

This chapter walks you from download to your first render.

## System requirements & download

FilmMovieStudio runs on all three desktop platforms:

Platform	Installer	Notes
macOS	.dmg	Universal build — works on both Intel and Apple Silicon Macs
Windows 10/11 (64-bit)	.exe	Standard setup wizard
Linux	.deb / .rpm / .AppImage	Debian/Ubuntu, Fedora/RHEL, or portable AppImage

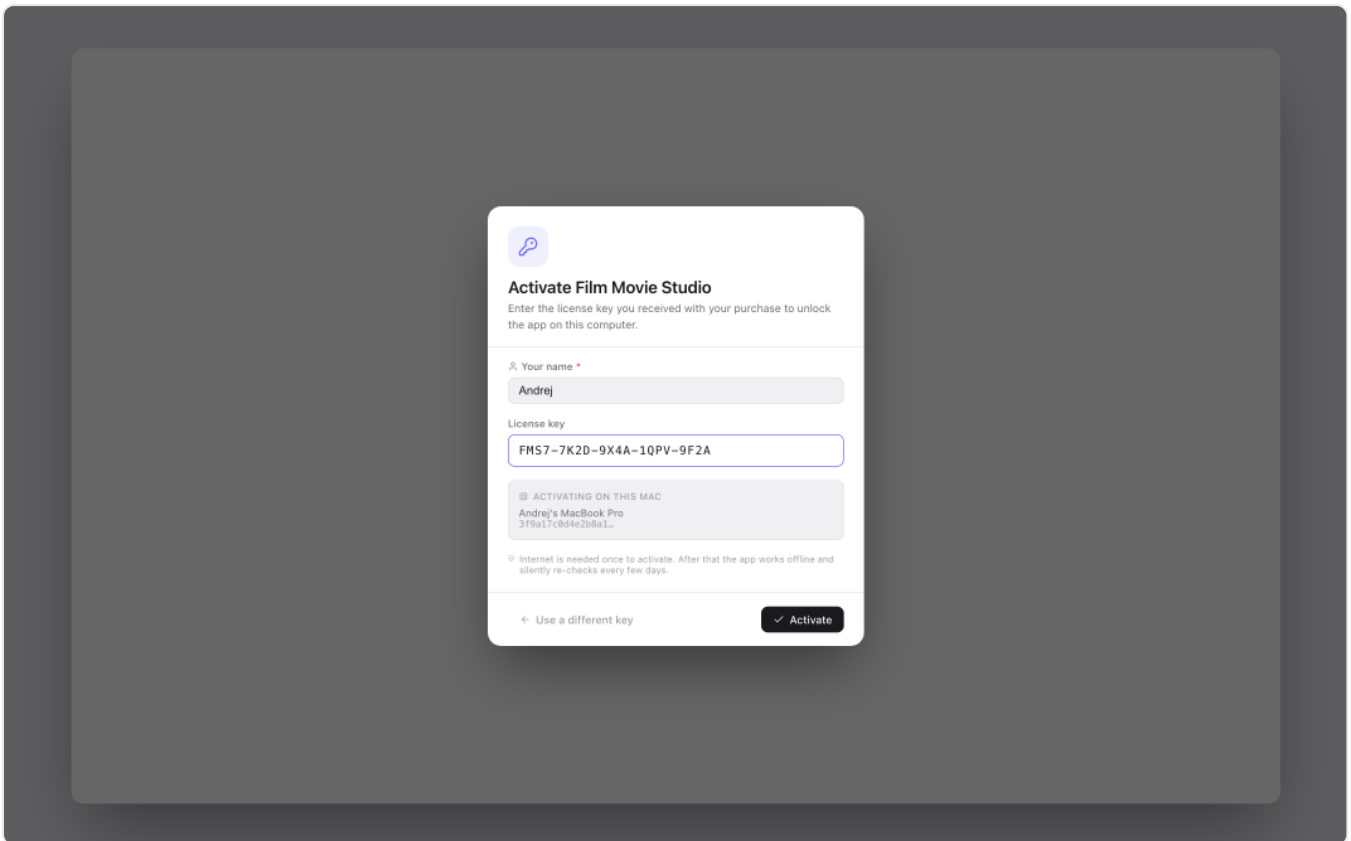
Download the app from [filmotvor.ai/studio](https://filmotvor.ai/studio). Sign in on the website first — the download links are available to logged-in users.

### First launch notes:

- **macOS:** the app is currently unsigned, so the very first time you open it, **right-click the app icon and choose "Open"**, then confirm. macOS remembers your choice and launches normally after that.
- **Windows:** you may see a SmartScreen prompt on first run — click "More info" → "Run anyway".

## Activate your license

When the app opens for the first time you'll see the **Activate Film Movie Studio** screen. It asks for the license key you received with your purchase — a code in the format `FMS7-XXXX-XXXX-XXXX-XXXX`, delivered by email after you buy on [filmotvor.ai/studio](https://filmotvor.ai/studio).



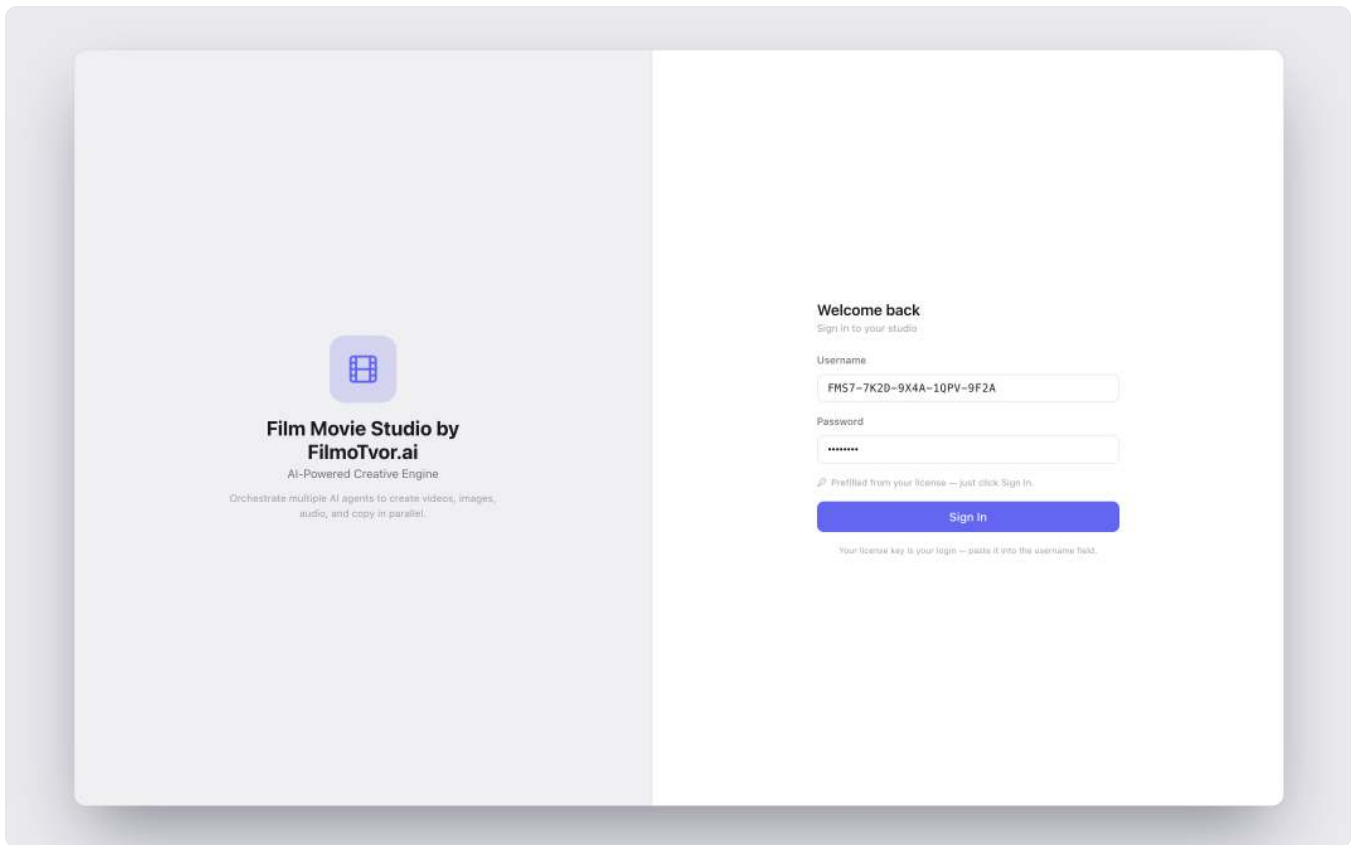
1. Type **Your name** — this becomes your display name inside the app.
2. Paste your key into the **License key** field (it auto-formats as you type).
3. Check the machine preview box — it shows which computer you're activating on.
4. Click **Activate**.

You need an internet connection once to activate. After that the app works offline and silently re-checks the license every few days. Activation also signs you in automatically — your license key doubles as your login.

**Tip:** Each license runs on a limited number of machines. If you ever see "License is registered to another machine", contact support to transfer it.

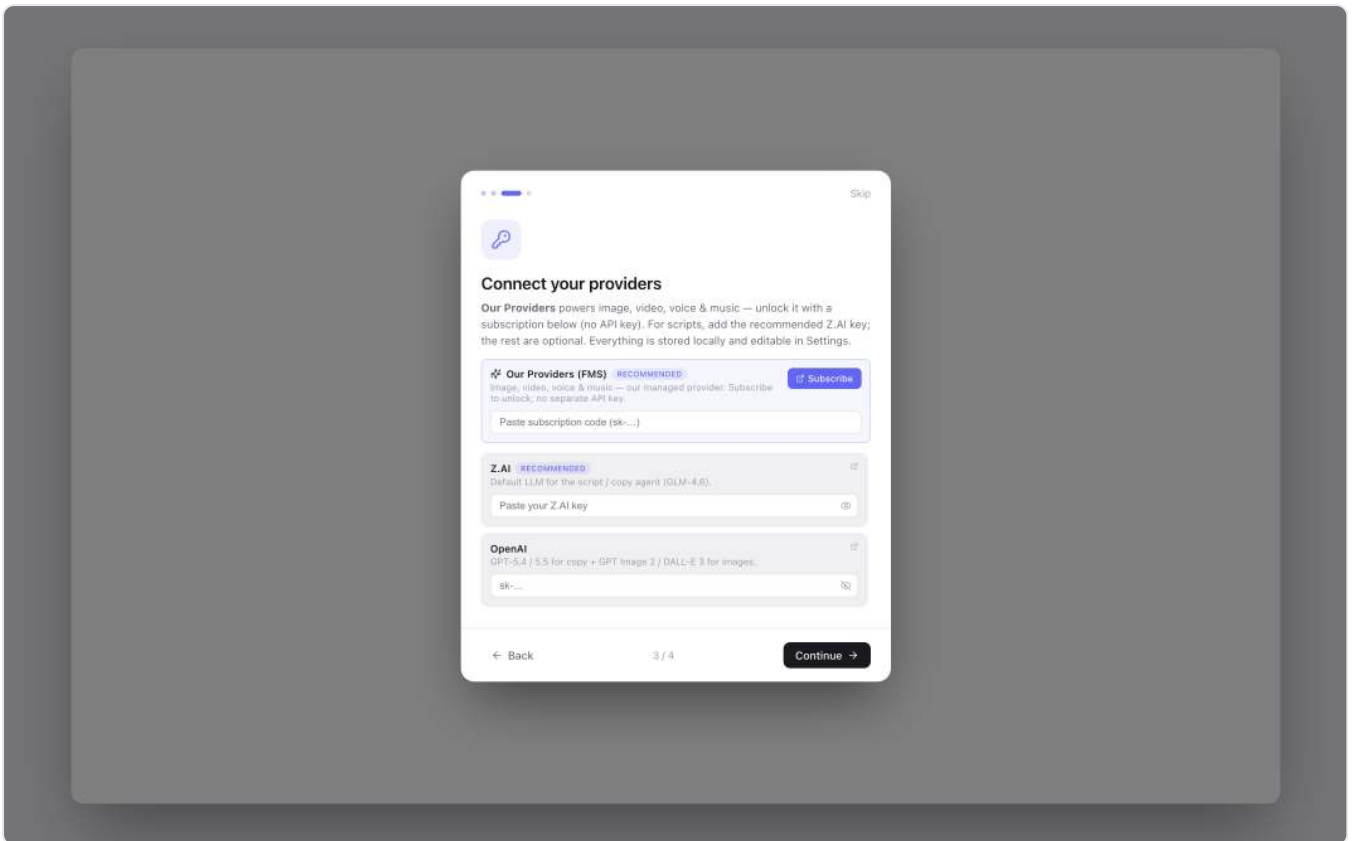
## Signing in later

If you ever sign out (from the account menu at the bottom of the sidebar), the **Welcome back** screen brings you back in. Your license key is prefilled as both username and password — just click **Sign In**.



## The first-run setup wizard

Right after activation, a short four-step wizard gets your studio ready. It takes about a minute, and you can click **Skip** at any time — everything here can be changed later in Settings.



1. **Welcome to Film Movie Studio** — a quick hello.
2. **How it works** — the three core ideas: the **Canvas** (double-click anywhere to spawn a node), the **Hub** (browse and install skills), and the **Chat** (describe what you want and the right agent generates it).
3. **Connect your providers** — unlock the managed **Our Providers (FMS)** service with a subscription code, and/or paste your own API keys (Z.AI is recommended for scripts; OpenAI, ElevenLabs, kie.ai and OpenRouter are optional).
4. **You're all set** — click **Get started**.

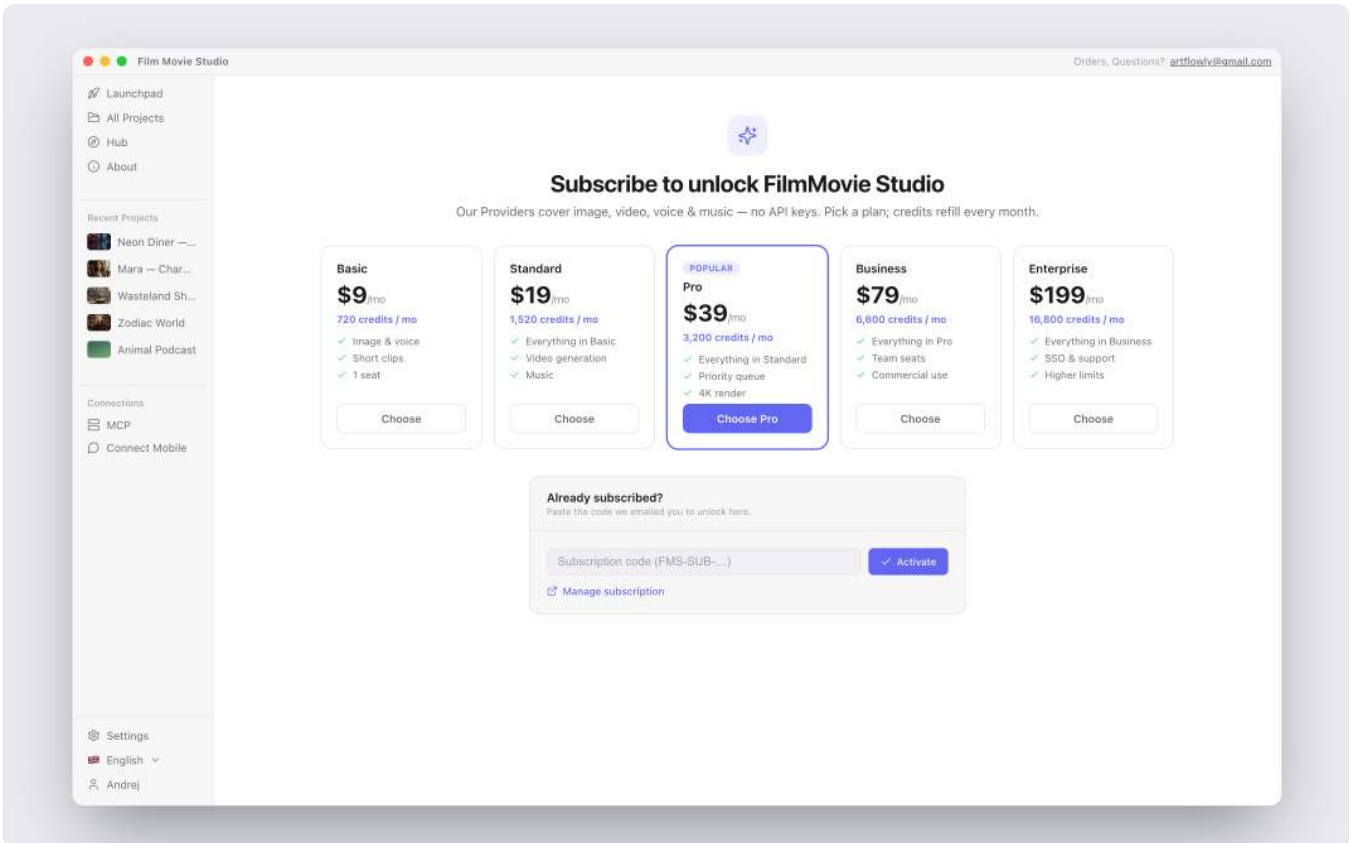
## Three ways to power generation

FilmMovieStudio doesn't generate media by itself — it orchestrates AI providers. You can power it in three ways, and mix them freely:

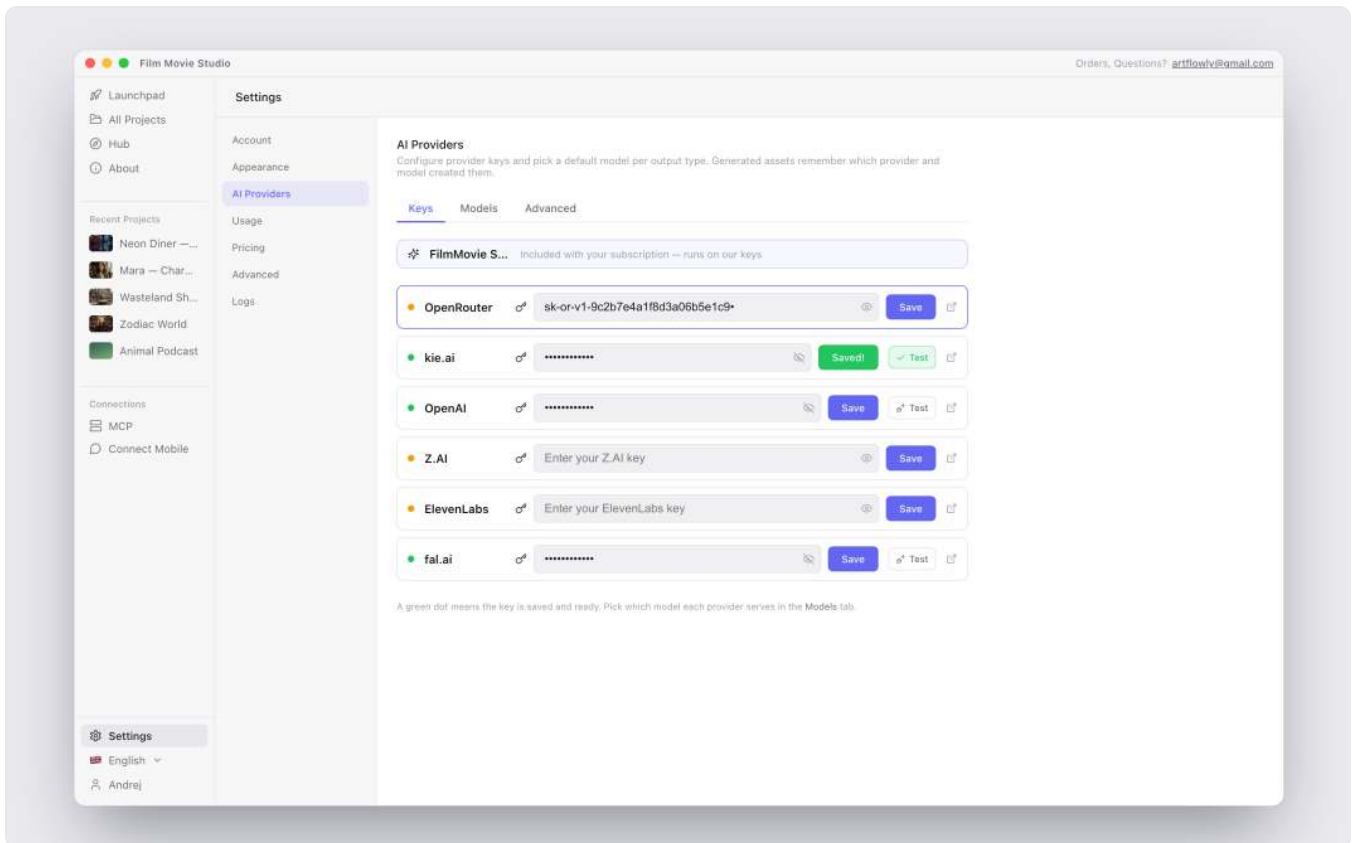
	FilmMovie subscription	Bring Your Own Keys	Free local models
<b>What it is</b>	Our managed provider — image, video, voice & music on our keys	Your own API keys for providers you choose	AI models that run on your computer
<b>What you need</b>	A subscription code (starts with sk-... )	API keys from kie.ai, fal.ai, OpenRouter, ElevenLabs, Z.AI, OpenAI...	Nothing — just disk space for the download
<b>Cost</b>	Monthly plan with included credits	You pay each provider directly	Free
<b>Works offline</b>	No	No	Yes, once downloaded

	FilmMovie subscription	Bring Your Own Keys	Free local models
Best for	The simplest setup — one code, everything works	Power users who want specific models and their own billing	Voiceovers and script drafts with zero cost and full privacy

**To activate a subscription:** buy a plan on [filmotvor.ai/studio](https://filmotvor.ai/studio), then open **Settings** → **Usage**, find the **Subscription** card, paste your code and click **Activate**. Your remaining **Credits** and renewal date show right there.



**To add your own keys:** open **Settings** → **AI Providers** and paste keys into the **Keys** tab. Each provider unlocks its own set of models in the **Models** tab.

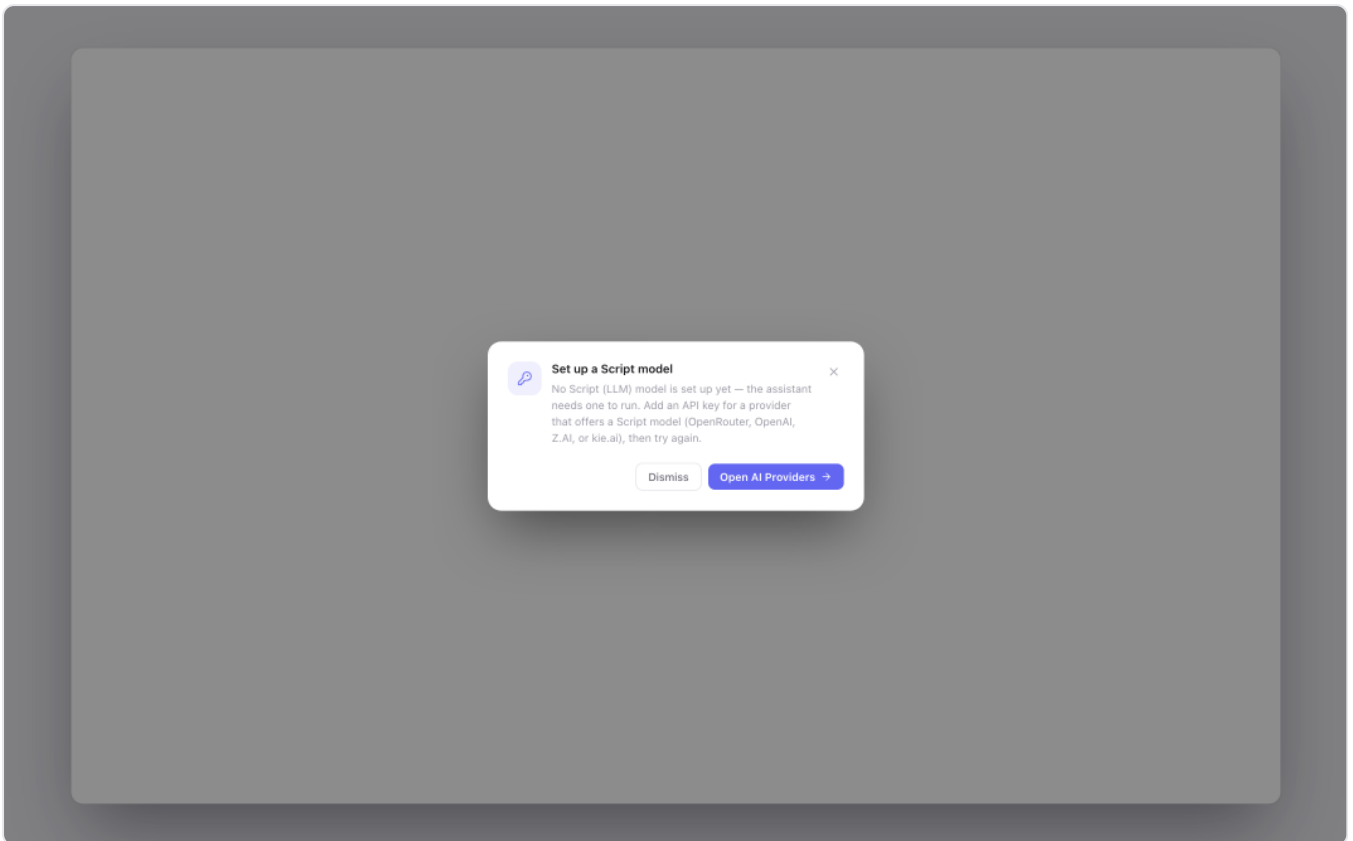


**To use free local models:** open **Settings** → **AI Providers** → **Local**. Two on-device models are available for download: **Supertonic** (text-to-speech narration in 31 languages, including Slovak) and **Qwen3 4B** (an offline script/chat model). No API key, no credits — and they keep working without internet once downloaded.

The full walkthrough of every provider, model selection and pricing lives in [Settings & AI Providers](#).

## "Set up a Script model" — the gate dialog

If you try to chat or launch a skill before any Script (LLM) provider is configured, this dialog appears:



It means the AI assistant has no language model to think with yet. Every conversation — even a simple image request — is driven by a Script model that reads your prompt and decides what to generate. To fix it, click **Open AI Providers** and do any one of these:

- Add an API key for a provider that offers a Script model (Z.AI, OpenRouter, OpenAI, or kie.ai), **or**
- Download the free local **Qwen3 4B** model in the **Local** tab.

The same dialog can also appear as "**API key needed**" when a skill requires an output type (image, video, voice) that has no provider yet — it lists exactly what's missing.

## Your first creation in 5 minutes

Ready? Here's the whole loop, end to end:

1. On the **Launchpad** (the home screen), type an idea into the big prompt box — for example: `/image a red sports car at sunset, cinematic lighting`. Or just describe it in plain words.
2. Press **Enter**. A new project is created automatically and its workspace opens with your prompt already running.
3. Watch the chat on the right: the agent thinks, picks the right tool, and the result appears as a card on the canvas within seconds.
4. Not quite right? Click **Remix** on the card for a fresh variation, or keep chatting — "make it night time", "add rain".
5. Drag the cards you like down onto the timeline, then click **Render** — your MP4 is saved and ready to share.

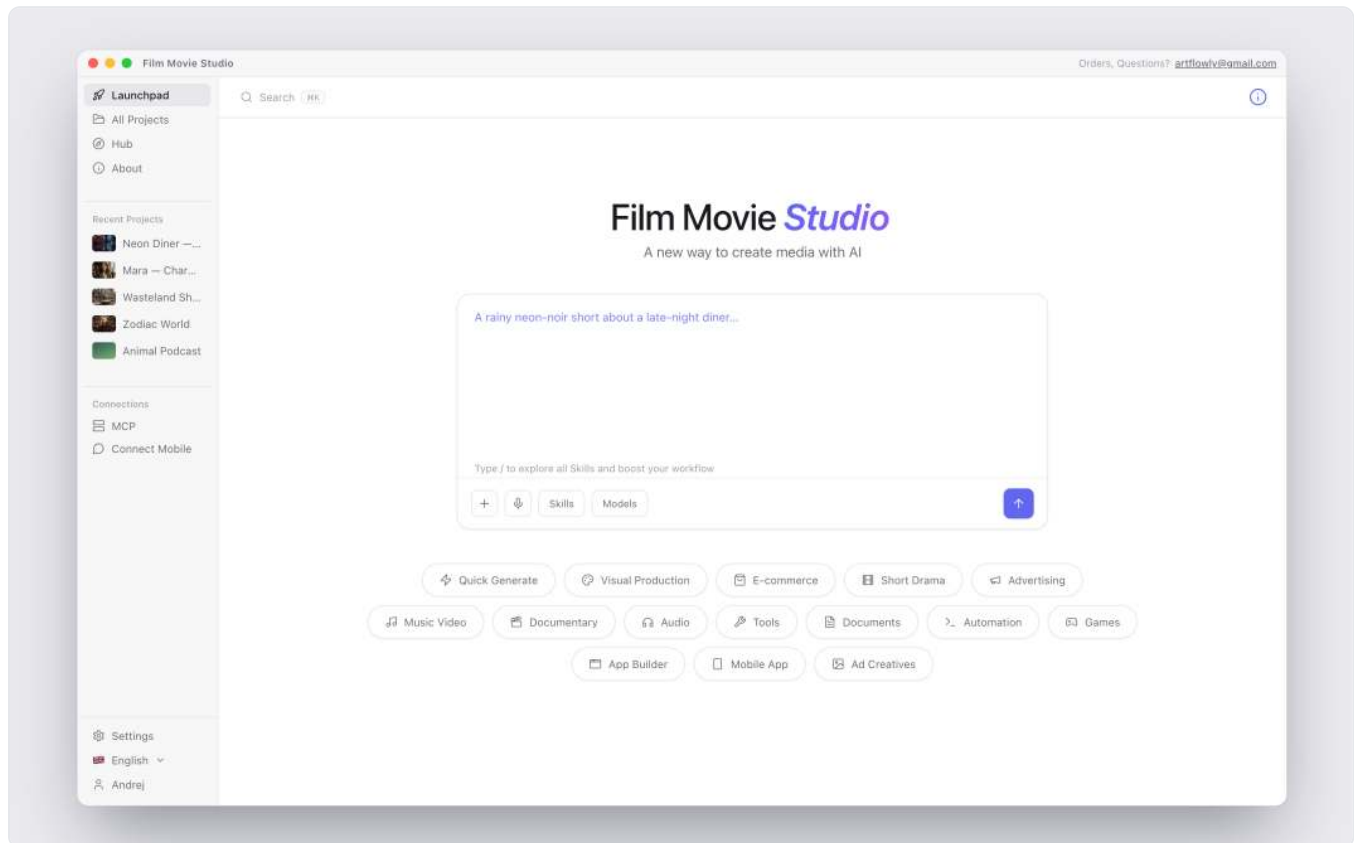
That's the entire creative loop: **prompt** → **canvas** → **timeline** → **render**. Everything else in this manual builds on it. Next up: the [Launchpad & Projects](#) chapter, then [Chat & the Agent](#) for going deeper with the assistant, and [Skills & the Hub](#) for multi-stage workflows like short dramas and music videos.

**Tip:** Stuck at any point? Ask the chat directly — "how do I add subtitles?" — the assistant knows the app and will walk you through it.

# Launchpad & Projects

The home screen, example prompts, preset categories and managing your projects.

The **Launchpad** is FilmMovieStudio's home screen — the first thing you see after signing in, and the fastest way to start anything. Type an idea, hit Enter, and a new project opens with the AI already working on it.



## The prompt box

The big command bar in the middle is the heart of the Launchpad. Under the headline — "A new way to create media with AI" — it cycles through typed-out example prompts so you can see the breadth of what's possible: `/image a red sports car at sunset...`, `/short-drama a 60-second sci-fi thriller in 6 scenes...`, `Design 4 logo concepts for a coffee roastery...` and more.

Around the box you'll find:

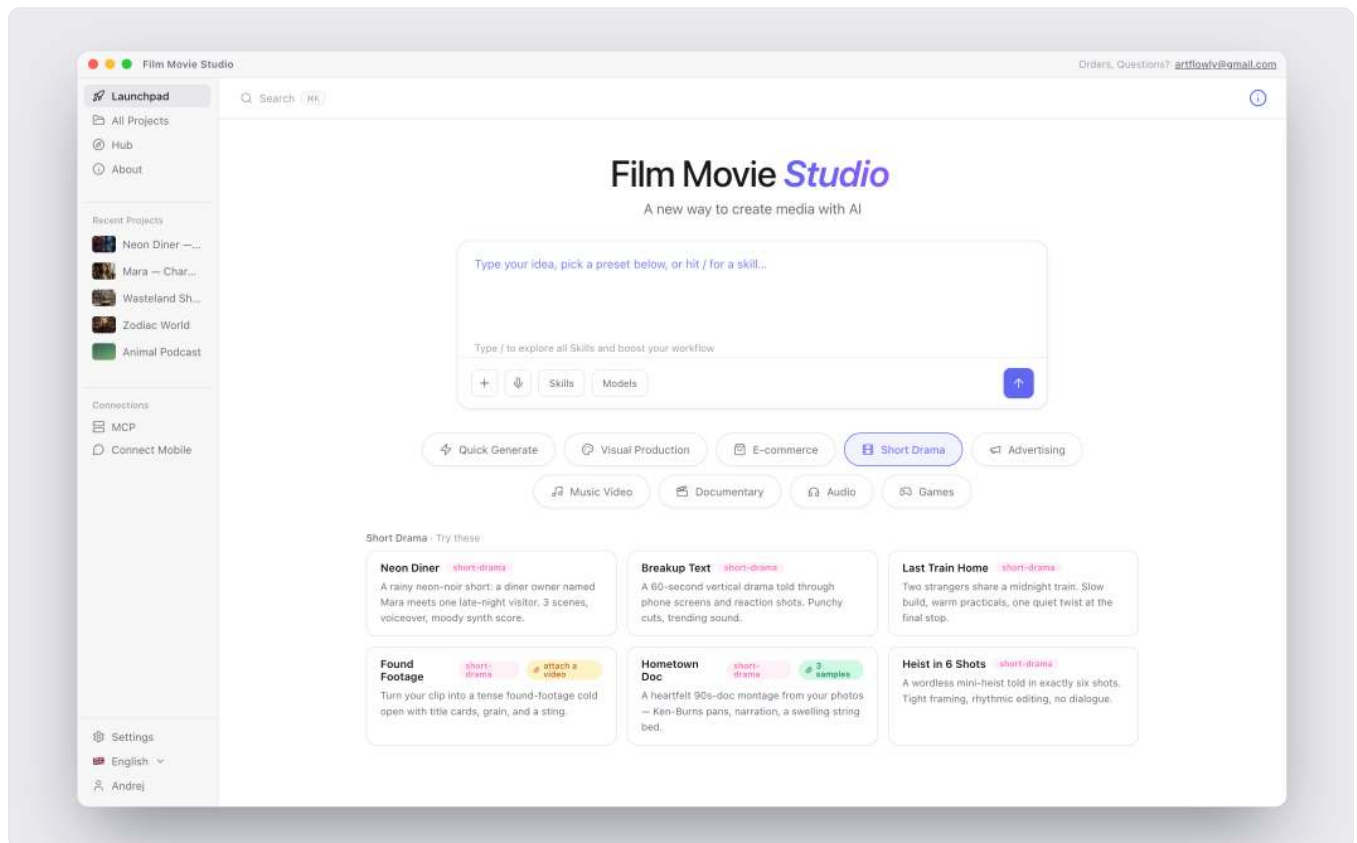
- **+** (**Attach file**) — attach up to 6 reference files (images, video, audio, PDFs, documents). They travel into the new project so the AI can actually see them.
- **Microphone** — dictate your prompt; the transcription lands in the box.
- **Skills** — opens the list of enabled skills; picking one inserts its `/command`. You can also just type `/` to browse and filter skills inline.
- **Models** — pick which AI models the agent is allowed to use, per category.
- **@ mentions** — type `@` to reference a saved asset or an attached image by name.

- The hint below the box reminds you: *"Type / to explore all Skills and boost your workflow."*

In the top toolbar, **Search** (or ⌘K) finds projects, skills, settings and quick links from one place. On the right, hovering the (i) icon shows a popover with your subscription **Credits**, the app version, and a **Guide** button that starts an interactive tour of the Launchpad.

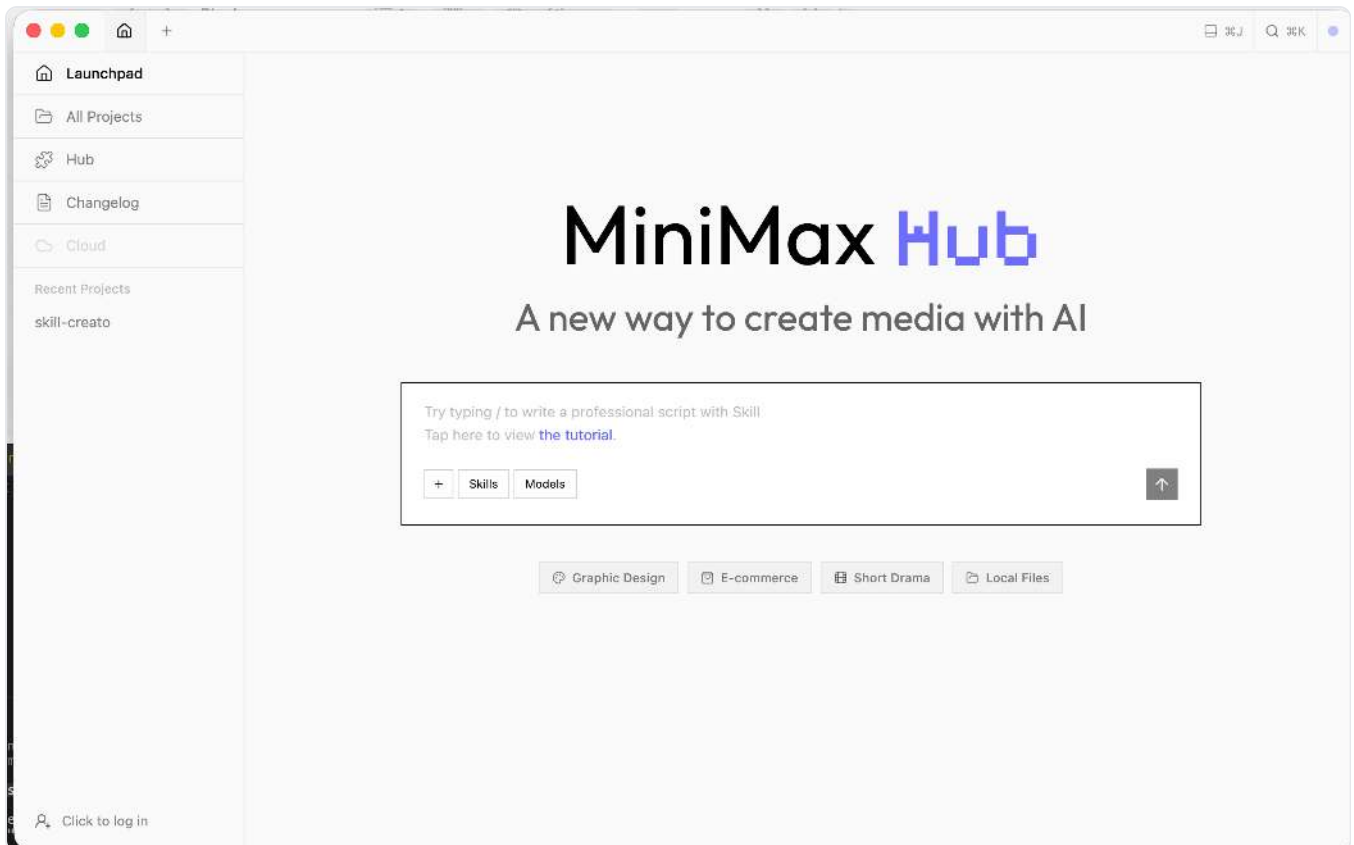
## Preset categories

Below the prompt box sits a row of category pills. Click one to reveal a **"Try these"** panel of ready-made preset prompts; click a preset to load it into the box, edit it if you like, and send.



The categories are: **Quick Generate, Visual Production, E-commerce, Short Drama, Advertising, Music Video, Documentary, Audio, Tools, Documents, Automation, Games, App Builder, Mobile App, Ad Creatives and Adobe Panel.**

Presets that need input media ship with sample files that attach automatically, so every preset works with a single click — no hunting for a matching photo or track.



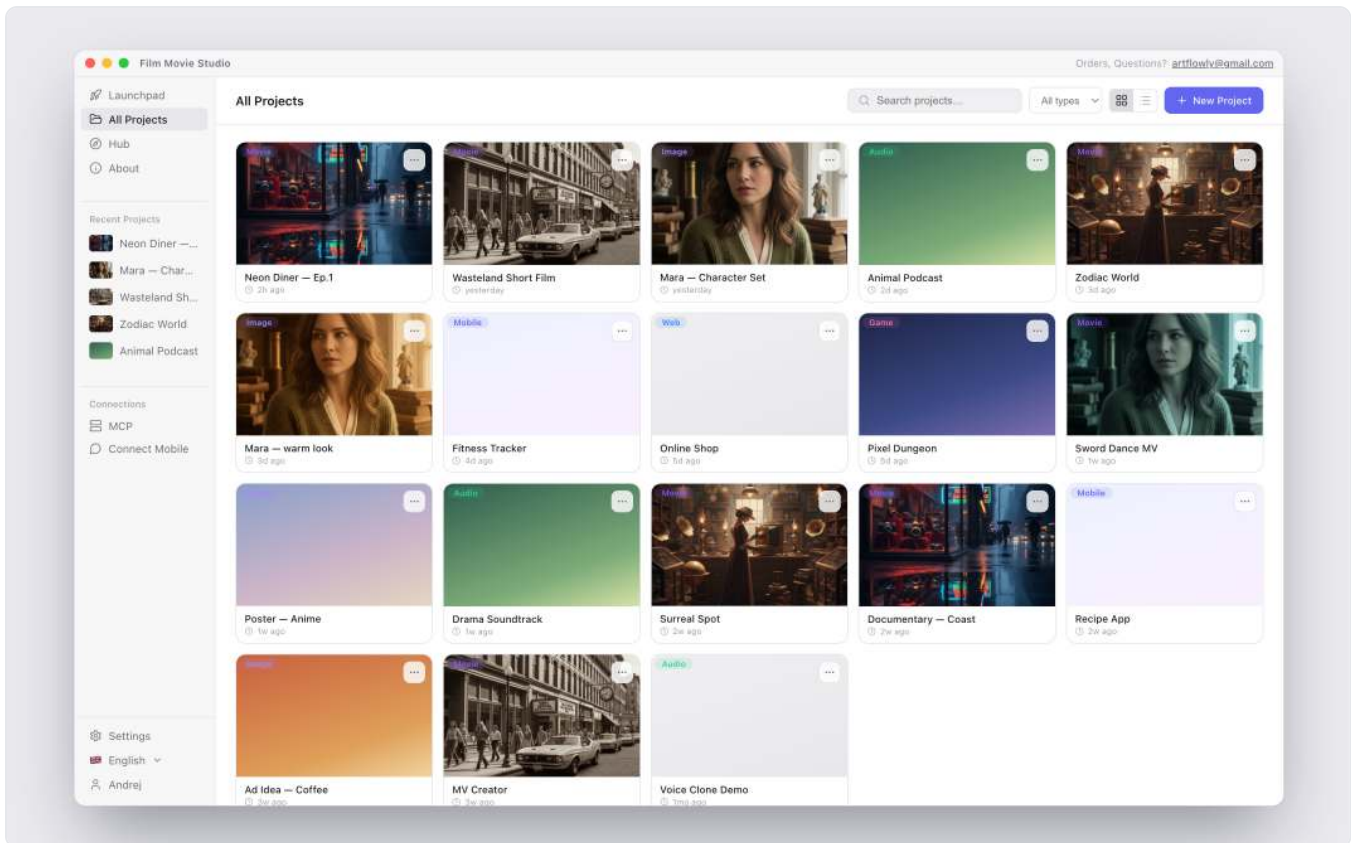
## What happens when you submit

1. A **new project is created automatically** — no dialog, no naming step. The app derives a title from your prompt instantly, then quietly asks the AI for a cleaner 2–5 word title a moment later.
2. The project **workspace opens** with your prompt (and any attachments) already handed to the agent.
3. The **agent starts working** immediately — results appear as cards on the canvas as they're generated. See [Chat & the Agent](#) for what happens next.

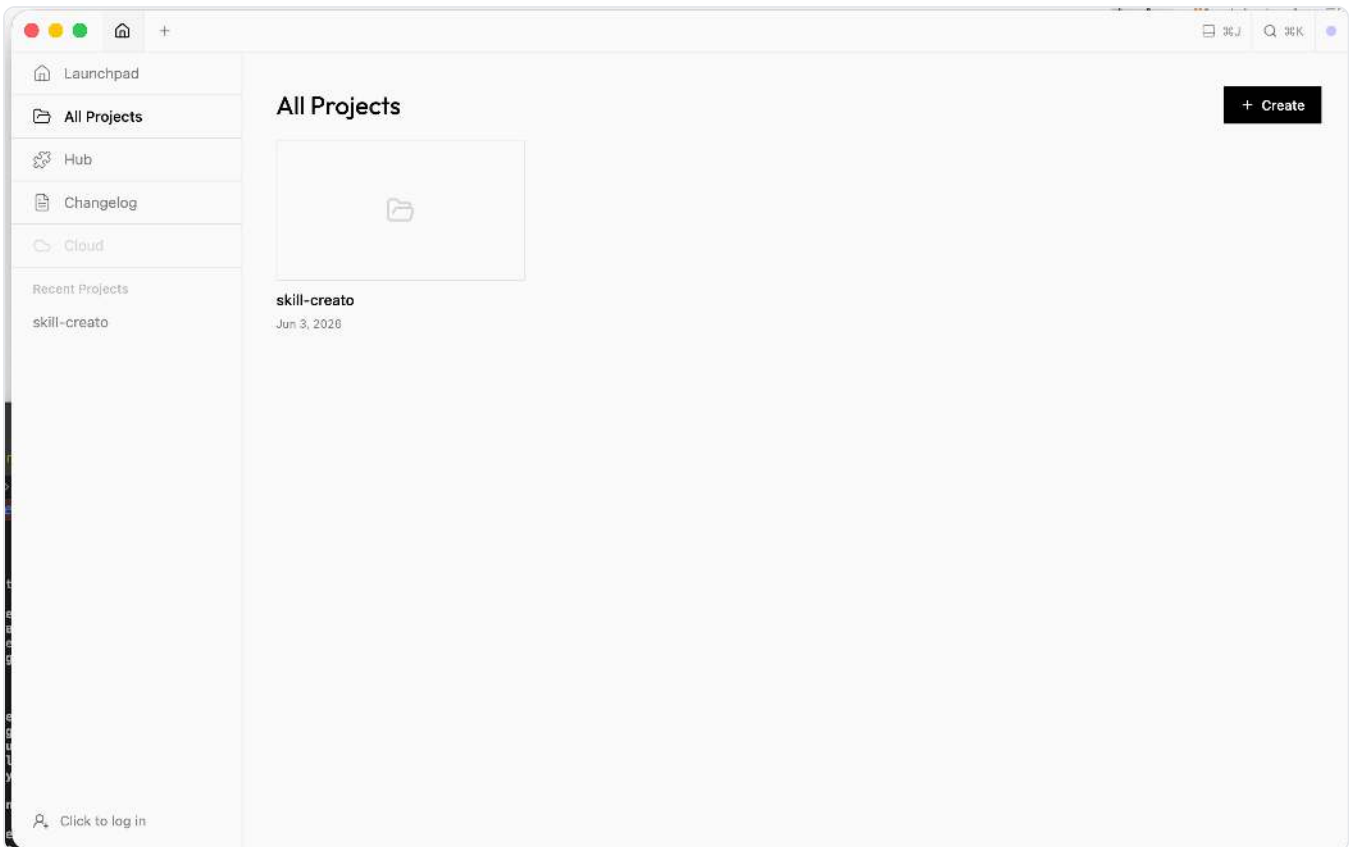
**Tip:** If a required provider isn't set up yet, a dialog appears instead of a broken project — see the "Set up a Script model" section in [Getting Started](#).

## All Projects

The **All Projects** page (in the sidebar) is your archive — every project you've ever created, as a grid of cards with cover thumbnails, searchable and paginated.

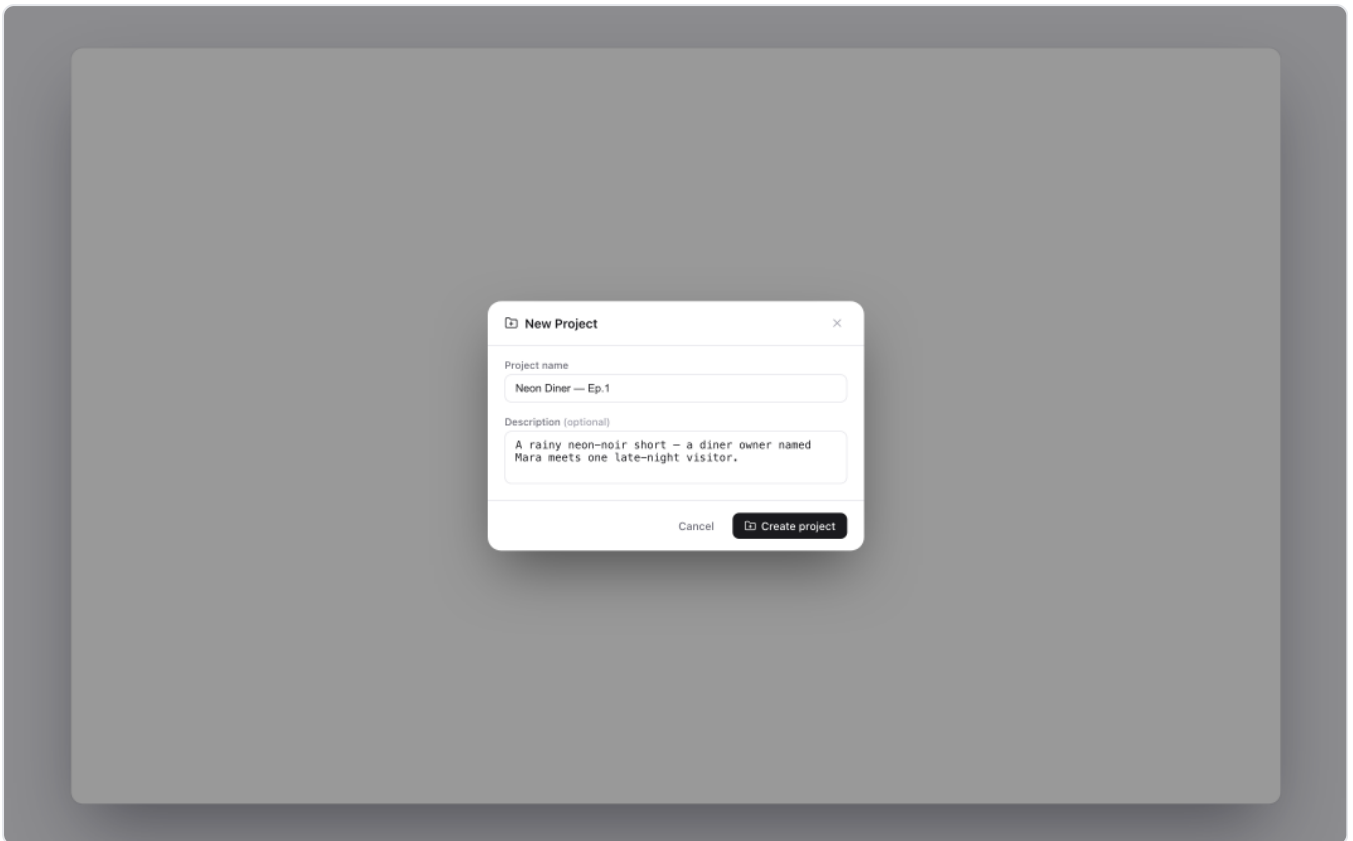


- **Open** a project by clicking its card. Hovering a card shows an enlarged preview of its cover.
- **Search** with the "Search projects..." field at the top.
- **Views:** switch between grid and list layout; choose how many projects show **Per page**, and page through with the arrows ("Showing 1–24 of 96", "Page 1 of 4").
- **Rename / edit:** open a card's menu and choose **Edit Project** — change the **Name** and **Description**, then **Save Changes**.
- **Delete:** choose **Delete Project**. You'll be asked to confirm — deleting removes all assets and chat history and **cannot be undone**.



## The New Project modal

Prefer to start with an empty canvas instead of a prompt? Click **New Project** on the All Projects page.



Give it a **Project name** (e.g. "Sunset Film"), optionally a **Description**, and click **Create project**. The empty workspace opens — you can add nodes by double-clicking the canvas or start chatting at any time.

## Sidebar tour

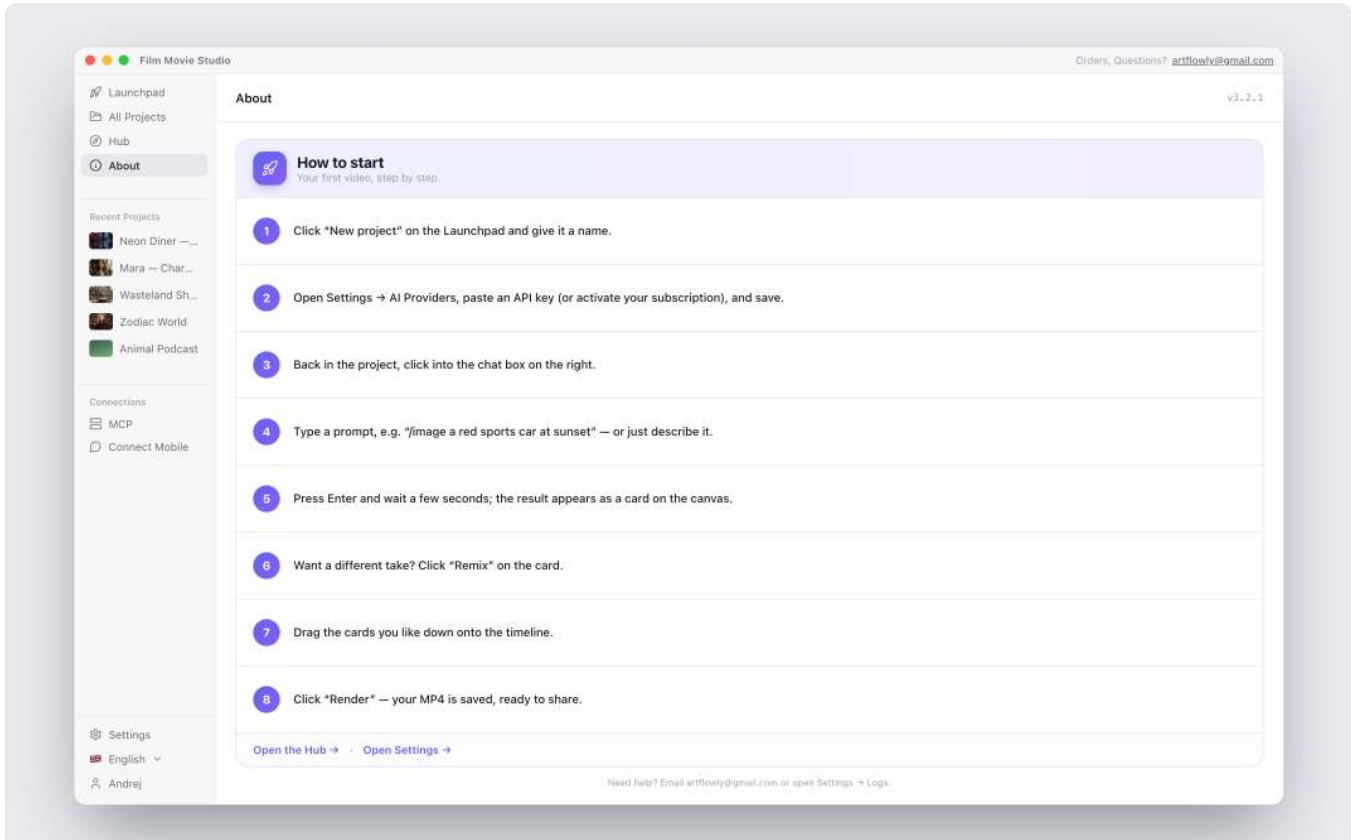
The sidebar on the left is your permanent navigation:

Item	What it does
<b>Launchpad</b>	The home screen described above
<b>All Projects</b>	The full project archive
<b>Hub</b>	Browse, install and manage Skills & Plugins — see <a href="#">Skills &amp; the Hub</a>
<b>Assets</b>	Your saved asset library, reusable across projects
<b>Cast</b>	The Cast library — reusable characters that keep the same face and voice in every project
<b>About</b>	Quickstart guide and app info
<b>Recent Projects</b>	Your 5 most recently updated projects, one click away
<b>Connections → MCP</b>	Let an external MCP client (Claude Desktop, Cursor) drive your open project
<b>Connections → Connect Mobile</b>	Drive the studio from your phone via a Telegram bot

At the bottom: **Settings**, the **language selector** (click the flag to switch languages — no restart needed), and the **account menu** showing your name, with **Account settings** and **Sign out**. Signing out keeps your license activated — the sidebar shows your masked key so you can sign back in with one click.

## The About page

The **About** page is a compact "How to start" card — eight concrete steps from creating a project to rendering an MP4 — with quick links to open the Hub and Settings. The header shows your exact app version (e.g. v3.3.5), handy when checking whether a newer release is available on [filmotvor.ai/studio](https://filmotvor.ai/studio) or when contacting support.

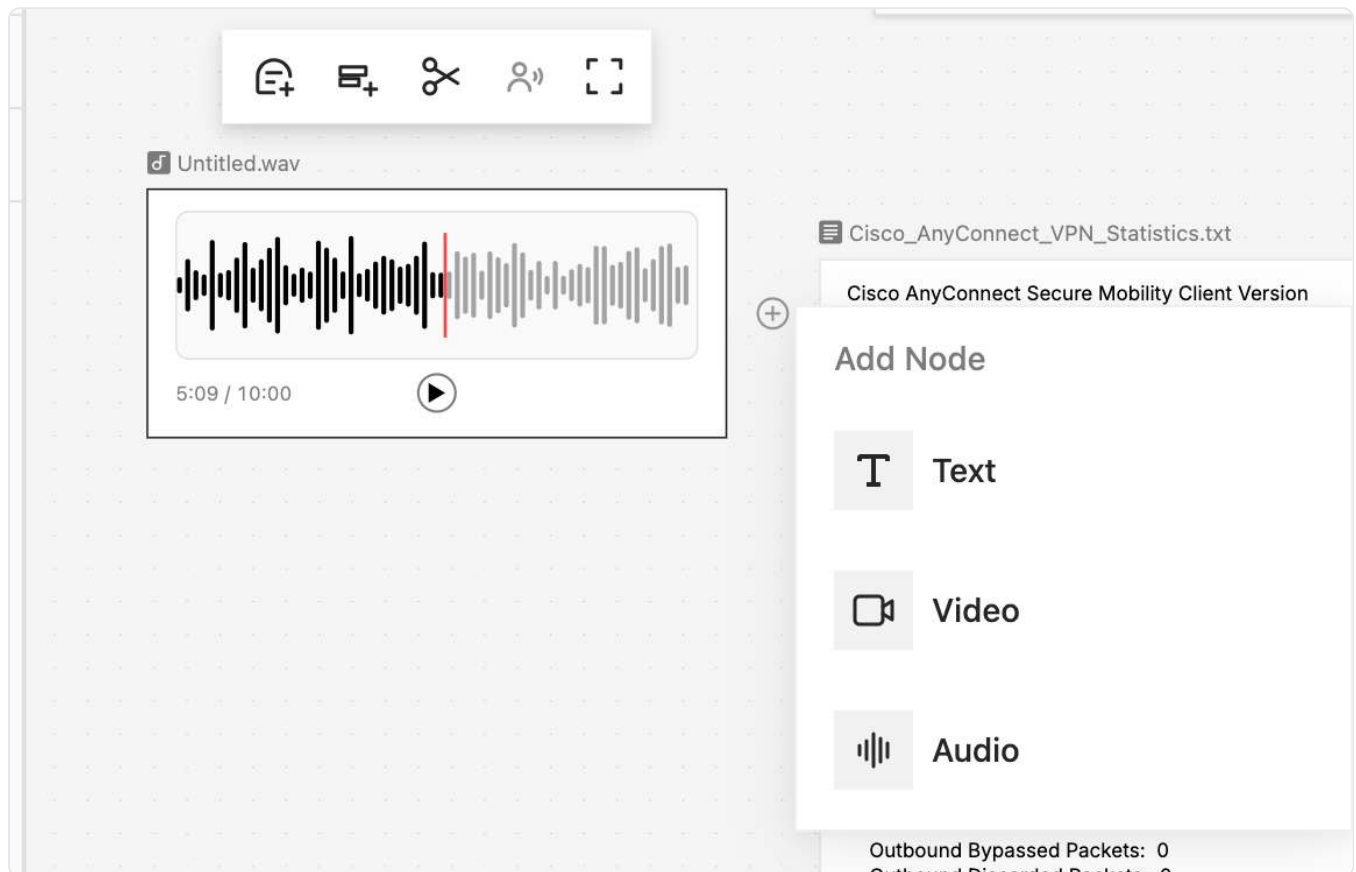


**Tip:** Forgot where something lives? The About page and the Launchpad's **Guide** tour are the two fastest refreshers — and you can always ask the chat.

# The Project Workspace

Anatomy of the workspace: icon rail, panels, tabs, and moving around the infinite canvas.

Every project in FilmMovieStudio opens into the same three-part workspace: a file and library panel on the left, an infinite canvas in the middle where your generated media lives as cards, and the chat panel on the right where you talk to the AI. This chapter walks you through each part and shows you how to move around comfortably.



## The project tabs bar

Across the very top of the window you'll find the tabs bar — it works just like browser tabs:

- **Home** — the house icon on the far left takes you back to the launchpad.
- **Project tabs** — every open project gets its own tab. Click a tab to switch projects; the canvas, timeline, and chat all swap together. Hover a tab and click the **x (Close tab)** to close it.
- **+ New** — the highlighted chip right after the last tab opens the **New Project** dialog. Give it a **Project name** (and an optional **Description**), click **Create project**, and the new project opens in a fresh tab.
- On the far right you'll see your **SUBSCRIBED** badge and live credit balance if you have a subscription. Without one, the balance sits blurred behind the subscribe prompt.

## The icon rail

The slim vertical strip on the left edge switches what the left panel shows:

- **Files** — the folder icon. Shows every file in the current project.
- **Hub — Skills & Plugins** — the compass icon. Browse and launch Skills and Plugins.
- **Assets Library** — the bookmark icon. Your cross-project library of saved cards (see [Cast & Asset Libraries](#)).

At the bottom of the rail:

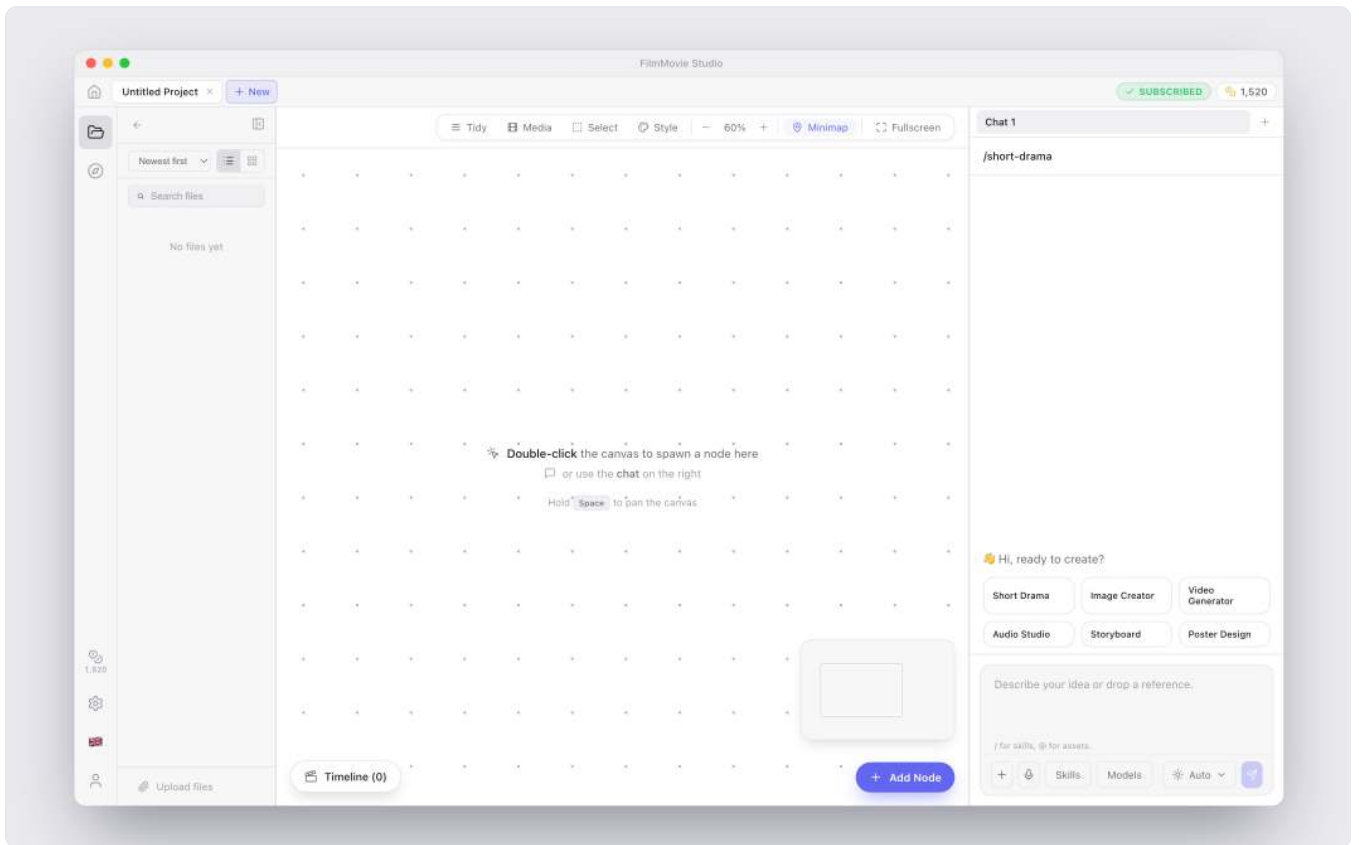
- **Credits** — the coin icon with your balance underneath; click it to open the Subscription section of Settings.
- **Settings** — the gear icon.
- **Language** — the flag icon opens the language switcher.
- **Account** — the user icon jumps to your account settings.

## The Files panel

With **Files** selected, the left panel lists everything the project has produced or imported. You can:

- **Search** files by name using the search box at the top.
- **Sort** by **Newest first**, **Oldest first**, or **Name**.
- Switch between **List view** and **Thumbnail view**.
- Click **Upload files** at the bottom to bring in your own images, videos, audio, or documents.

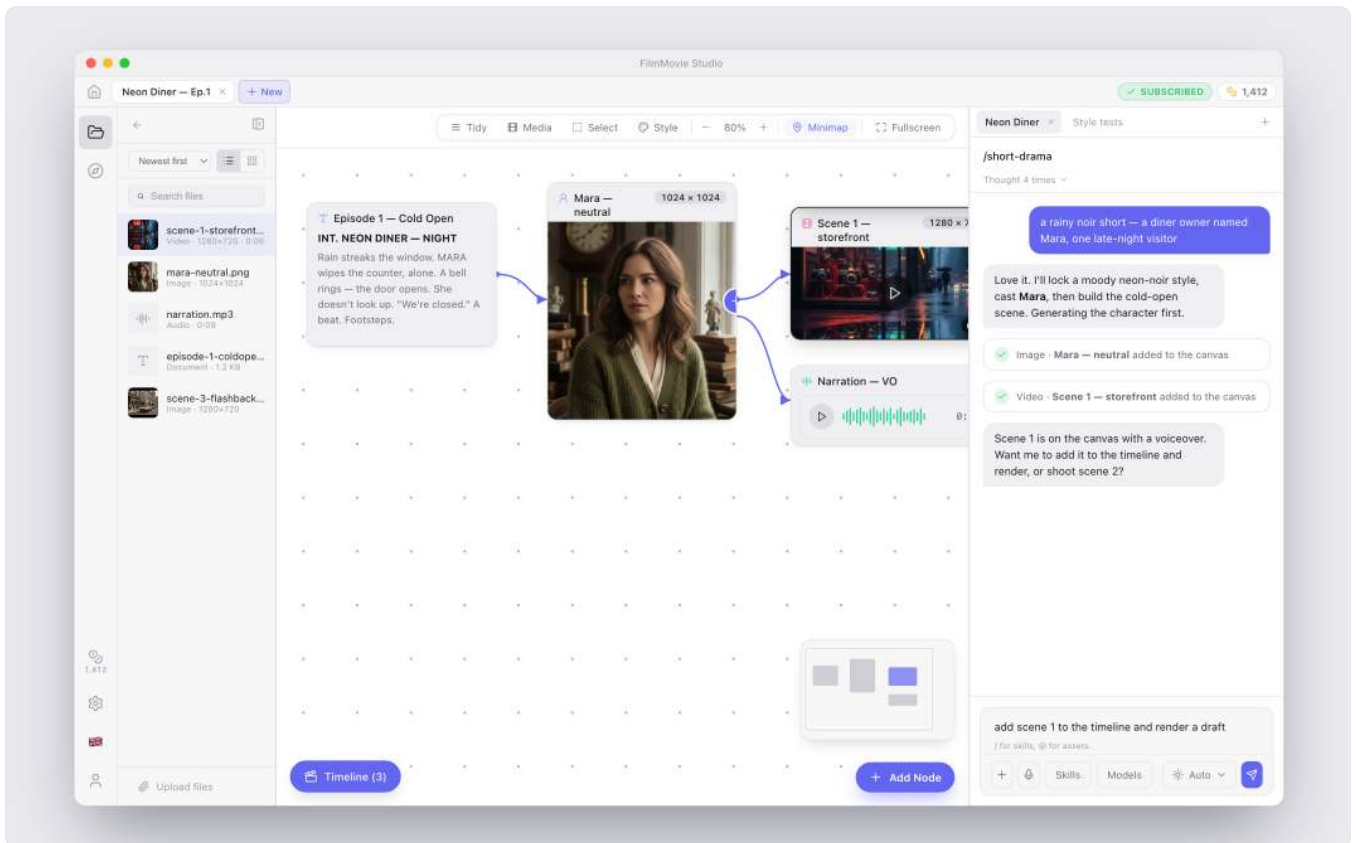
Anything you upload is copied into the project's local assets folder — your media always stays on your computer — and is immediately placed on the canvas as a card, so you can use it as a reference, remix it, or drop it on the timeline. You can also drag a file from the Files panel straight onto the canvas to place it exactly where you want it.



## The canvas and the chat panel

The big area in the middle is the **canvas** — a freeform, infinite surface where every generation lands as a card. The panel on the right is the **chat**, where you describe what you want and the AI agent builds it card by card. The chat is covered in depth in [Chat & the Agent](#); the cards themselves in [Canvas & Cards](#).

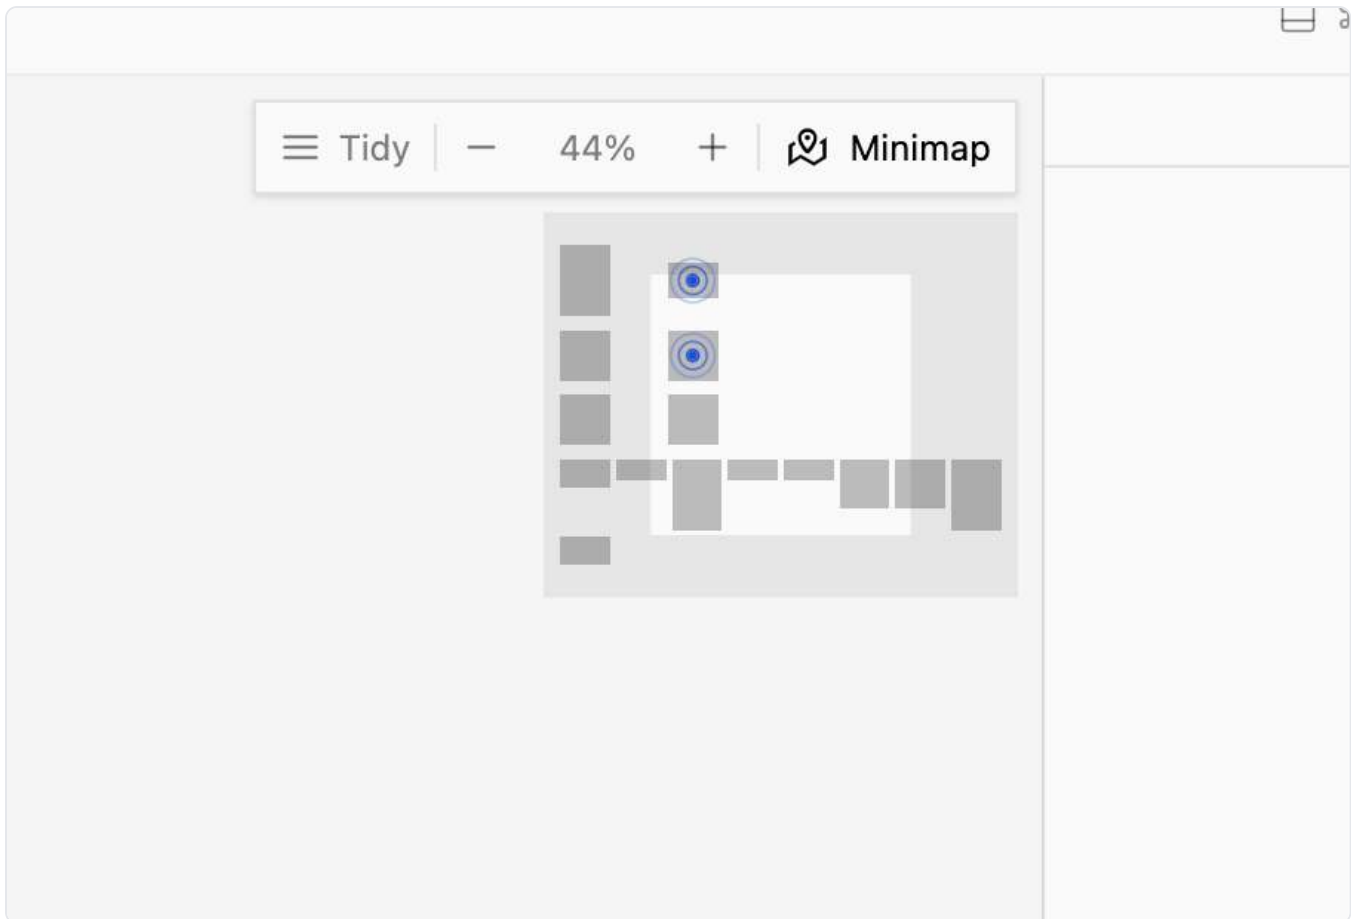
In the bottom-left corner of the canvas sits the **Timeline** pill (it shows how many clips are on it, e.g. "Timeline (5)") — click it to open the [Timeline Editor](#). In the bottom-right corner, the **+ Add Node** button lets you add a blank card at any time.



## Moving around the canvas

Once a project fills up, navigation matters. Here is everything at your fingertips:

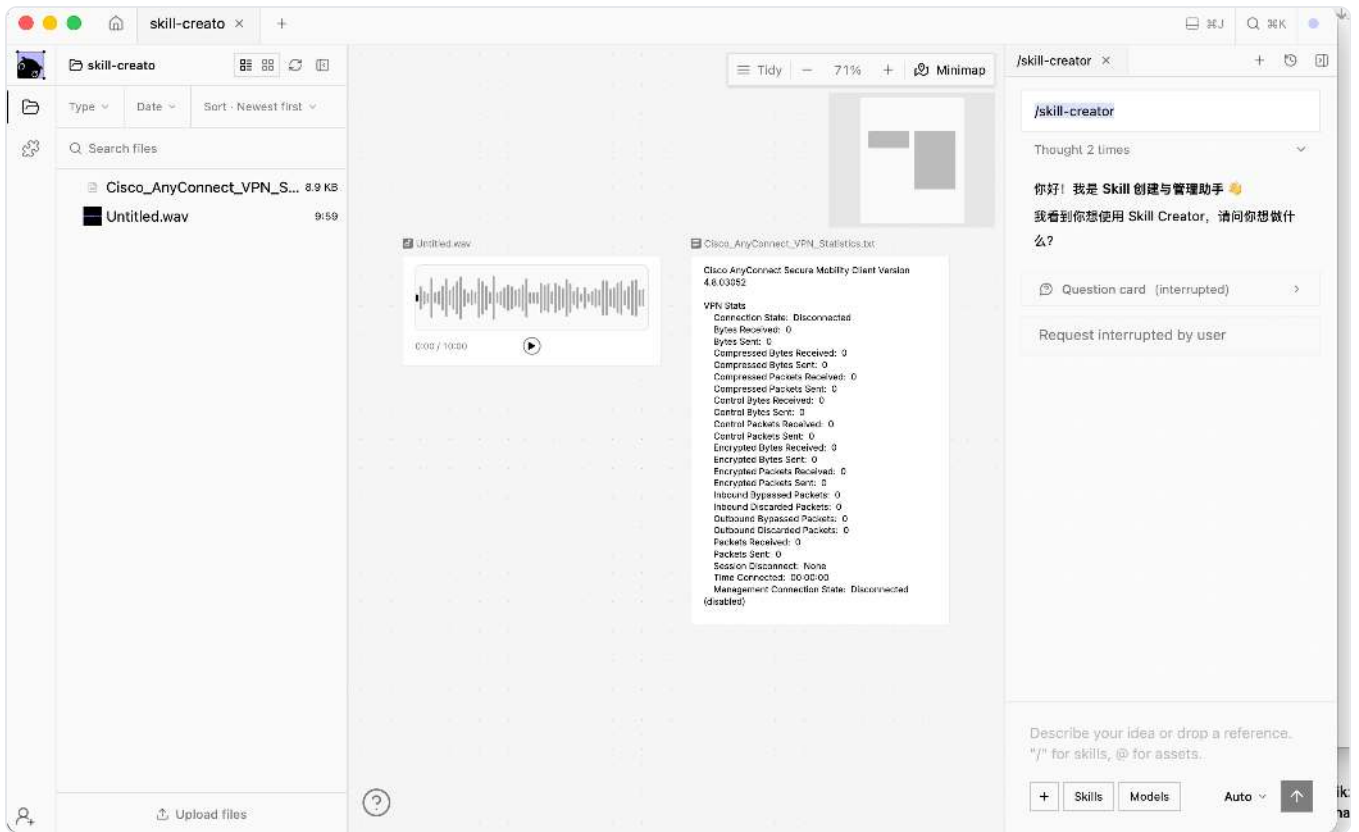
- **Zoom with the mouse wheel.** Scrolling zooms in and out, centered on your cursor — no modifier key needed. You can also use the **-** and **+** buttons in the top-right toolbar, which step the zoom 10% at a time between 5% and 100%. New projects start at 60%.
- **Pan by holding Space and dragging.** The cursor turns into a grab hand while you drag the canvas around.
- **Minimap.** The **Minimap** toggle in the top-right toolbar shows a small overview map in the corner of the canvas. The white rectangle is your current view — click or drag anywhere on the minimap to jump there instantly.



The top-left toolbar holds the canvas tools:

- **Tidy** — auto-arranges all cards into clean columns. If you don't like the result, press **⌘/Ctrl+Z** to undo it.
- **Media** — hides prompt and text cards so you see only images, videos, and audio. Click again to show everything.
- **Select** — turns a background drag into a selection box. Drag over several cards to select them all, then use the selection bar to **Group** them into a colored region, **Add all to chat**, **Add all media to timeline**, **Align into a row**, or **Delete all selected**. Press **Esc** to clear the selection and exit Select mode.
- **Style** — opens the **Project Style Lock**: a style sentence plus an optional reference image that gets applied to every image and video generation in the project, so all your visuals stay consistent. Pin a reference straight from any image card with its **Pin style** button.
- **Fullscreen** — the expand icon hides everything except the canvas. A floating toolbar keeps all the tools available; press **Esc** or click the fullscreen button again to exit.

**Tip:** Double-click any empty spot on the canvas to open the **Add Node** menu right there — the fastest way to drop a new Text, Table, Image, Video, or Audio card exactly where you want it.



## Where to go next

Now that you can find your way around, learn what the cards themselves can do in [Canvas & Cards](#), drive the whole pipeline from [Chat & the Agent](#), or assemble your finished clips in the [Timeline Editor](#).

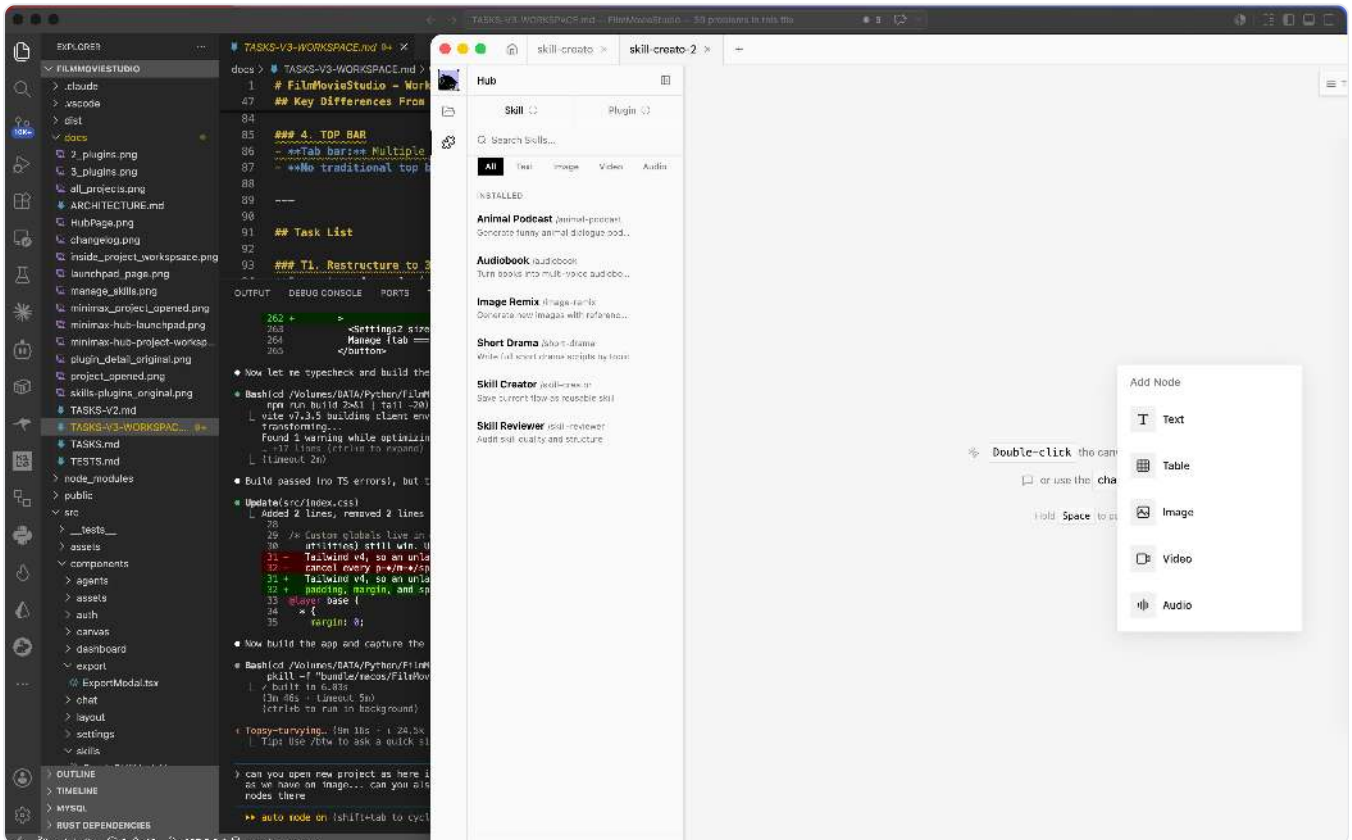
# Canvas & Cards

Every card type and every toolbar button — remix, variations, upscale, outpaint, style lock and more.

Everything you make in FilmMovieStudio lives on the canvas as a **card** (also called a node): every image, video, audio clip, script page, game, and document is a card you can move, connect, remix, and send onward to the chat or the timeline. This chapter is the full reference for every card type and every button on its toolbar.

## Cards and connections

When you ask the agent in [Chat & the Agent](#) to generate something, the result lands on the canvas as a new card — you never have to place it yourself. Cards created from other cards (a video made from a keyframe image, variations of an image, a scene generated from a script) are drawn with a **connection line** back to their parent, so you can always trace where something came from.

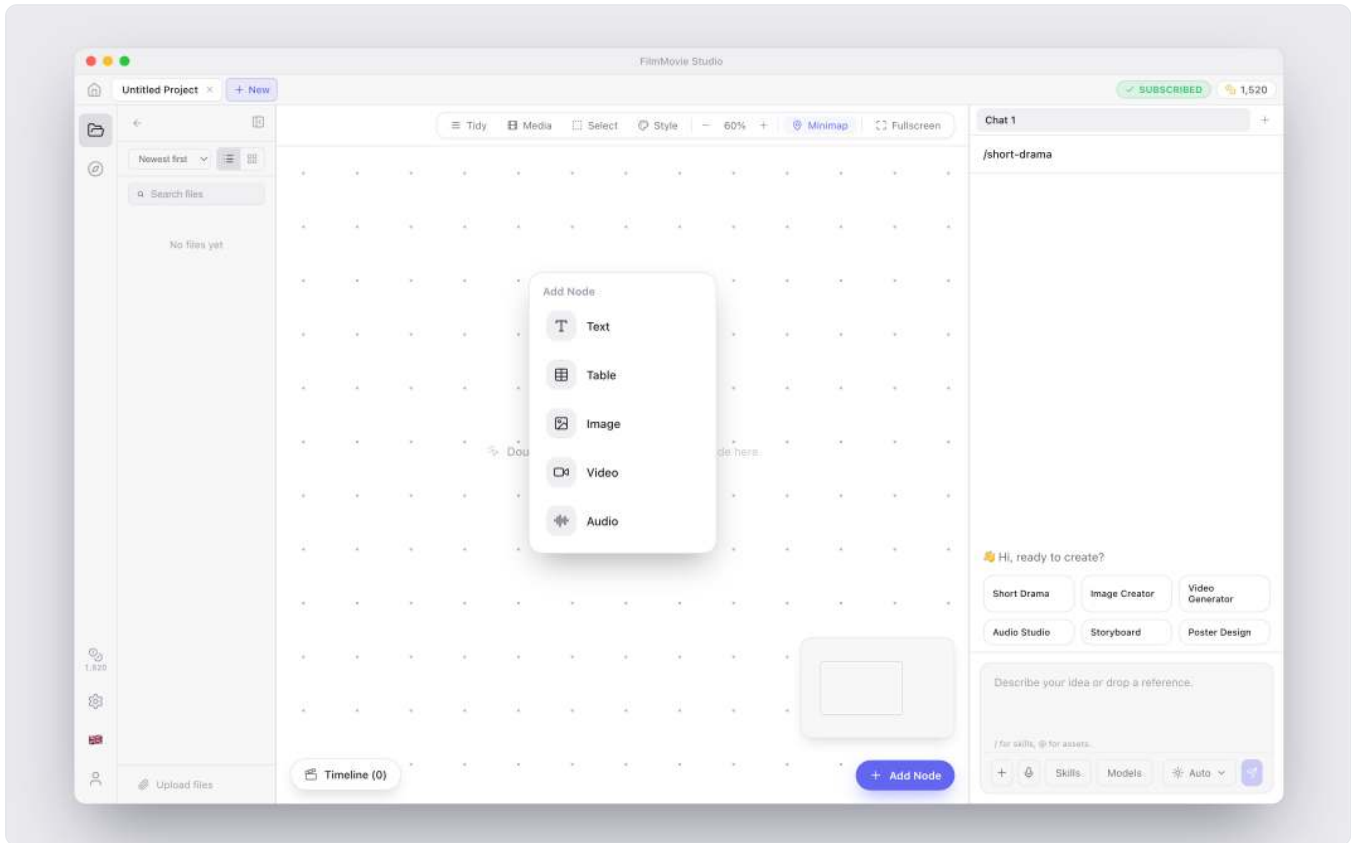


Hover any card and a  $\oplus$  handle appears at its edge (**Add connected node**). Click it — or drag it out onto empty canvas — to create a linked child card. When you branch off an image, the new card automatically carries the source image as its reference, so an "Image → Video" branch animates that exact picture.

## Adding cards manually

You don't need the chat to create a card. Double-click any empty spot on the canvas, or click the **+ Add Node** button in the bottom-right corner, to open the **Add Node** menu with five blank card types:

- **Text** — freeform notes, prompts, or scripts.
- **Table** — simple tabular text.
- **Image** — an empty image card with a generation panel.
- **Video** — an empty video card with a generation panel.
- **Audio** — an empty audio card with **Audio** (speech) and **Music** tabs.



## Generated card types

Skills and the chat agent create some specialized cards on top of the five basics. They behave like their base type but display extra content:

- **Character** — a portrait image plus a short profile; behaves like an image card and can be saved to the Cast library (see [Cast & Asset Libraries](#)).
- **Plan / Episode / Scene / Intent** — script and planning text produced by workflow skills; behave like text cards.
- **Final** — the rendered final video of a workflow; behaves like a video card (without Remix, since it's a composite).
- **Screen** — an app screen built by the App Builder skill, with its own export tools.
- **Game** — a playable game built by the game skills, playable right on the card.
- **Document** — a generated office file (PDF, Word, Excel) you can open with your system apps.

Select any card to reveal its **toolbar** above it. The sections below list every button per card type. Every toolbar ends with a trash icon — **Delete node** — which asks you to confirm (**Delete / Cancel**) before removing the card.

## Image cards

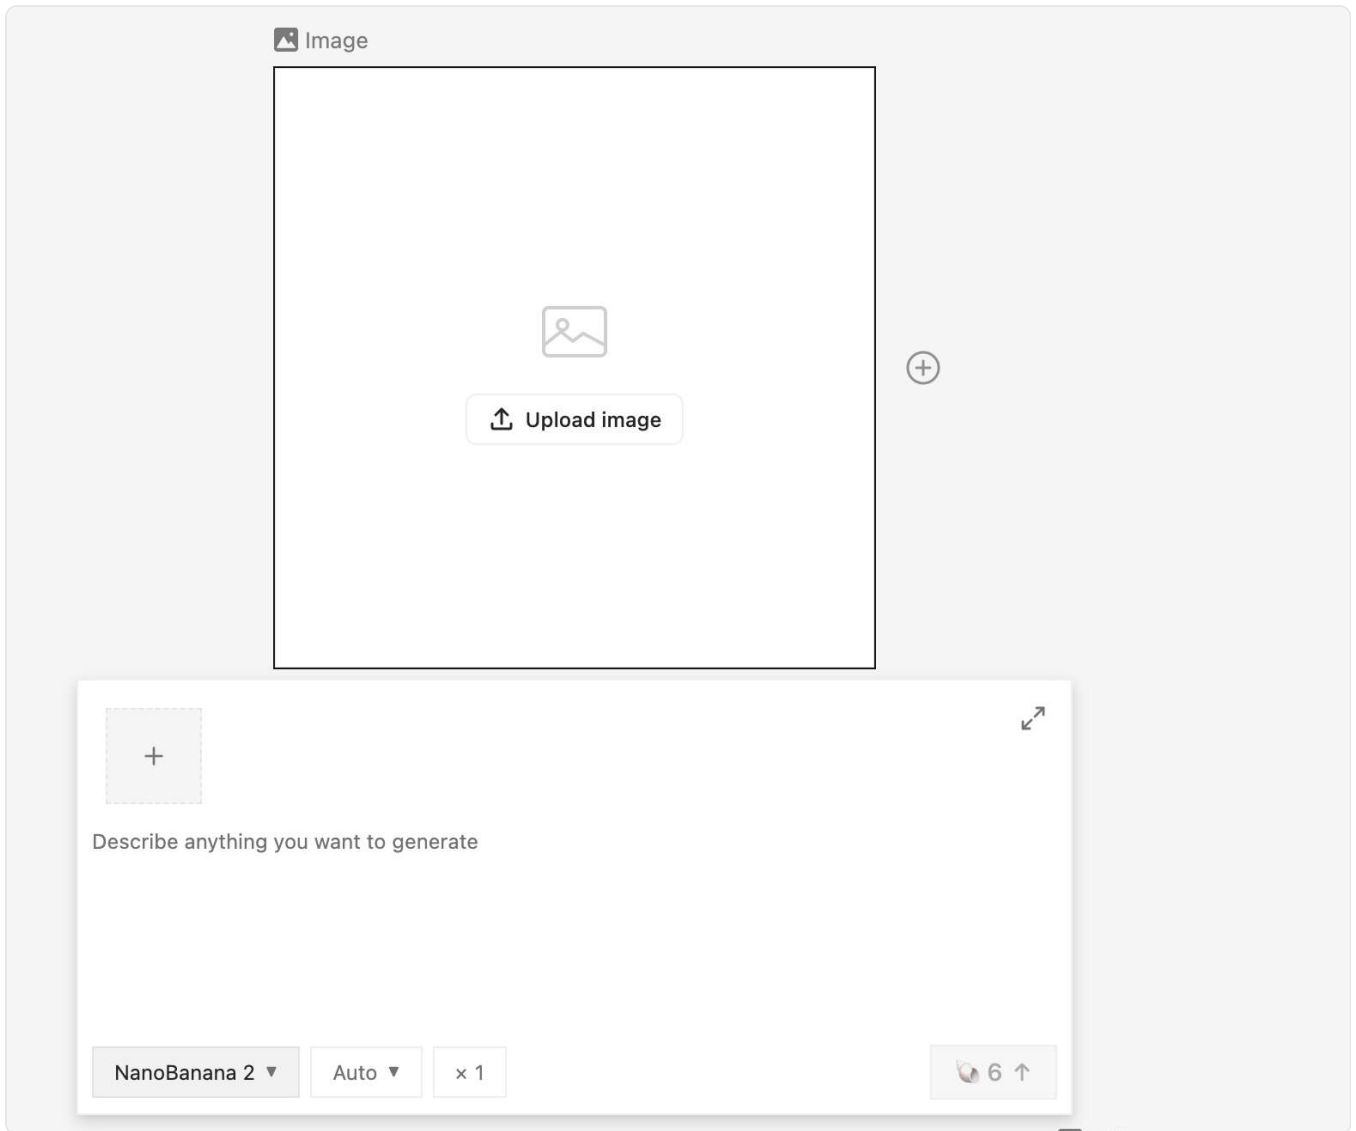



Image cards (and character cards) have the richest tool set — so many tools that only your **pinned** favorites sit in the main bar, while the rest live in the **More tools** menu behind the  button. By default the main bar shows **Remix**, **Color adjust**, and **Fullscreen**; use **Customize Toolbar** to pin the tools you use most.

Button	What it does	Credits
Remix	Regenerates a variation using the card's saved prompt	—
Variations	Generates <b>3</b> , <b>4</b> , or <b>6</b> fresh takes on the same prompt as child cards	—
Pin style	Pins this image as the project's Style Lock reference	—
Outpaint	Extends the picture beyond its borders to a new aspect ratio	60
HD Upscale	4x super-resolution upscale	100
Remove BG	Removes the background	20
Edit Image	Describe a change in words and apply it to the image	—
Erase & Redraw	Paint a mask, then regenerate just that region	—

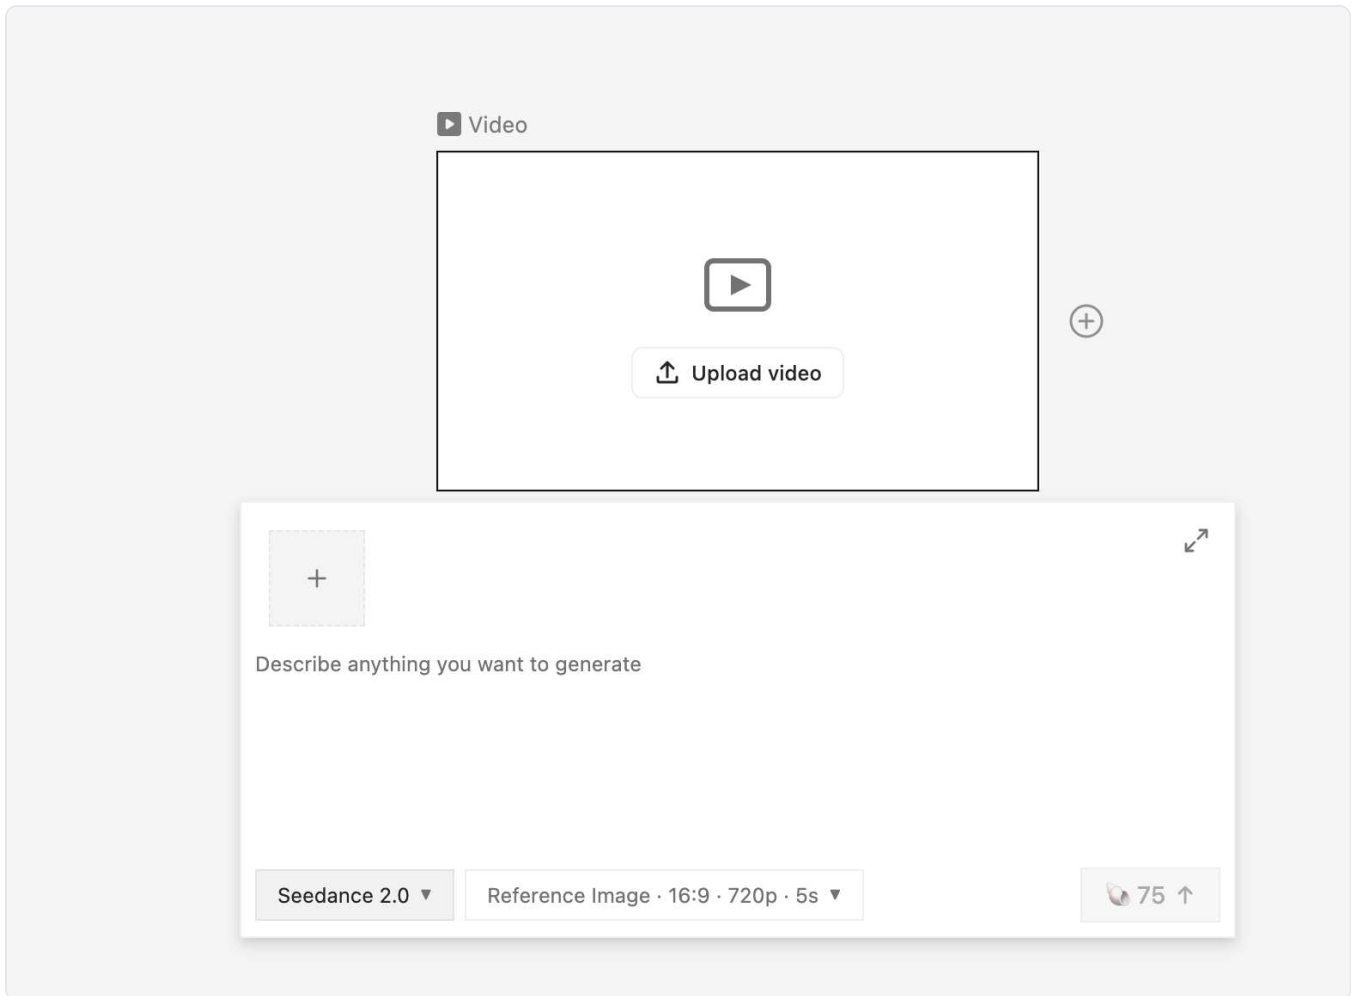
Button	What it does	Credits
Rotate	Rotate 90° clockwise / counter-clockwise, flip horizontally or vertically	—
Add to chat	References this card in your next chat message	—
Add to timeline	Sends the image to the <a href="#">Timeline Editor</a>	—
Color adjust	Opens the color adjustment sliders	—
Split Grid	Cuts a grid image into separate cards: 4-Grid (2x2), 9-Grid (3x3), 16-Grid (4x4), or 25-Grid (5x5)	—
Add to Assets	Bookmarks the card into your cross-project Assets Library	—
Save to Cast	Saves a character's face (and voice) to the Cast library	—
Fullscreen	Opens a full-window preview	—
Customize Toolbar	Choose which tools are pinned to the main bar	—

A few of these deserve a closer look:

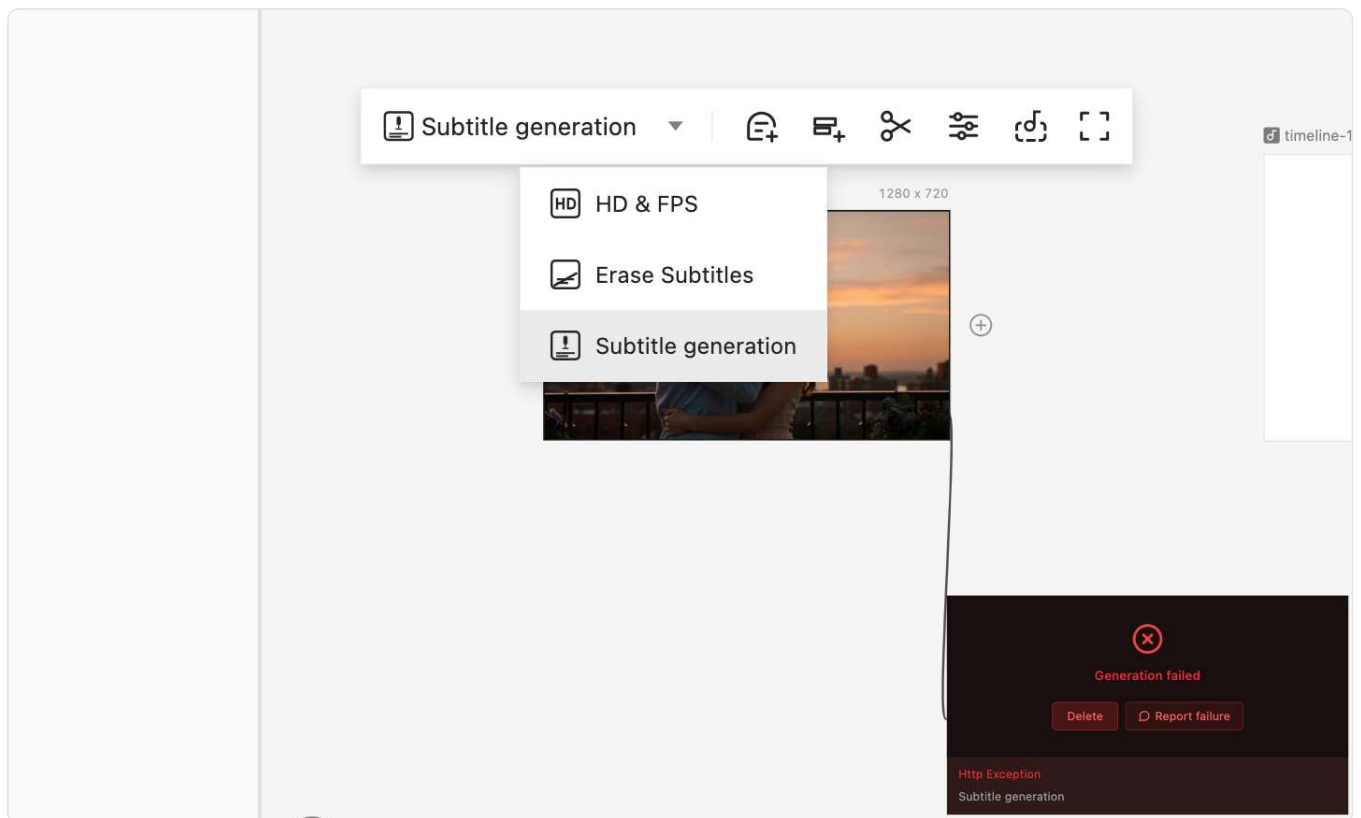
- **Remix** re-runs the exact prompt that made the card, so you get a sibling take without retyping anything. It works on any chat-generated image (and video).
- **Variations** fans out several children at once — great for picking the best of a batch.
- **Pin style** feeds the **Project Style Lock** (the **Style** button in the canvas toolbar): the pinned image plus your style sentence are applied to every image and video generation in the project, keeping the whole film visually consistent.
- **Outpaint** opens a panel where you choose **Scale**, **Resolution**, a **Preset**, and a target aspect ratio (1:1, 3:4, 2:3, 9:16, 4:3, 3:2, 16:9, 4:5, 5:4, or Original), then click **Generate** (60 credits) to grow the image outward.
- **Edit Image** is the quickest retouch: type into "**Describe the edit...**" and click **Apply**.
- **Erase & Redraw** gives you a **Brush** and **Eraser** with an adjustable size, plus **Undo** and **Clear mask**. Paint over what you want replaced, type "**Describe what to draw in the painted area...**", and click **Apply** — only the painted region is regenerated.
- **Split Grid** is perfect after generating a storyboard or character sheet as one big grid image — each tile becomes its own card.
- **Color adjust** opens sliders for **Color (hue)**, **Vibrance**, **Saturation**, **Temperature**, **Tint**, **Contrast**, and **Brightness**, with a **Reset** button.
- **Save to Cast** locks a character's face so they look (and sound) the same in every future project — see [Cast & Asset Libraries](#).

**Tip:** In **Customize Toolbar** you can also switch on **Show tool names** to display a label under every icon while you're learning the tools.

## Video cards

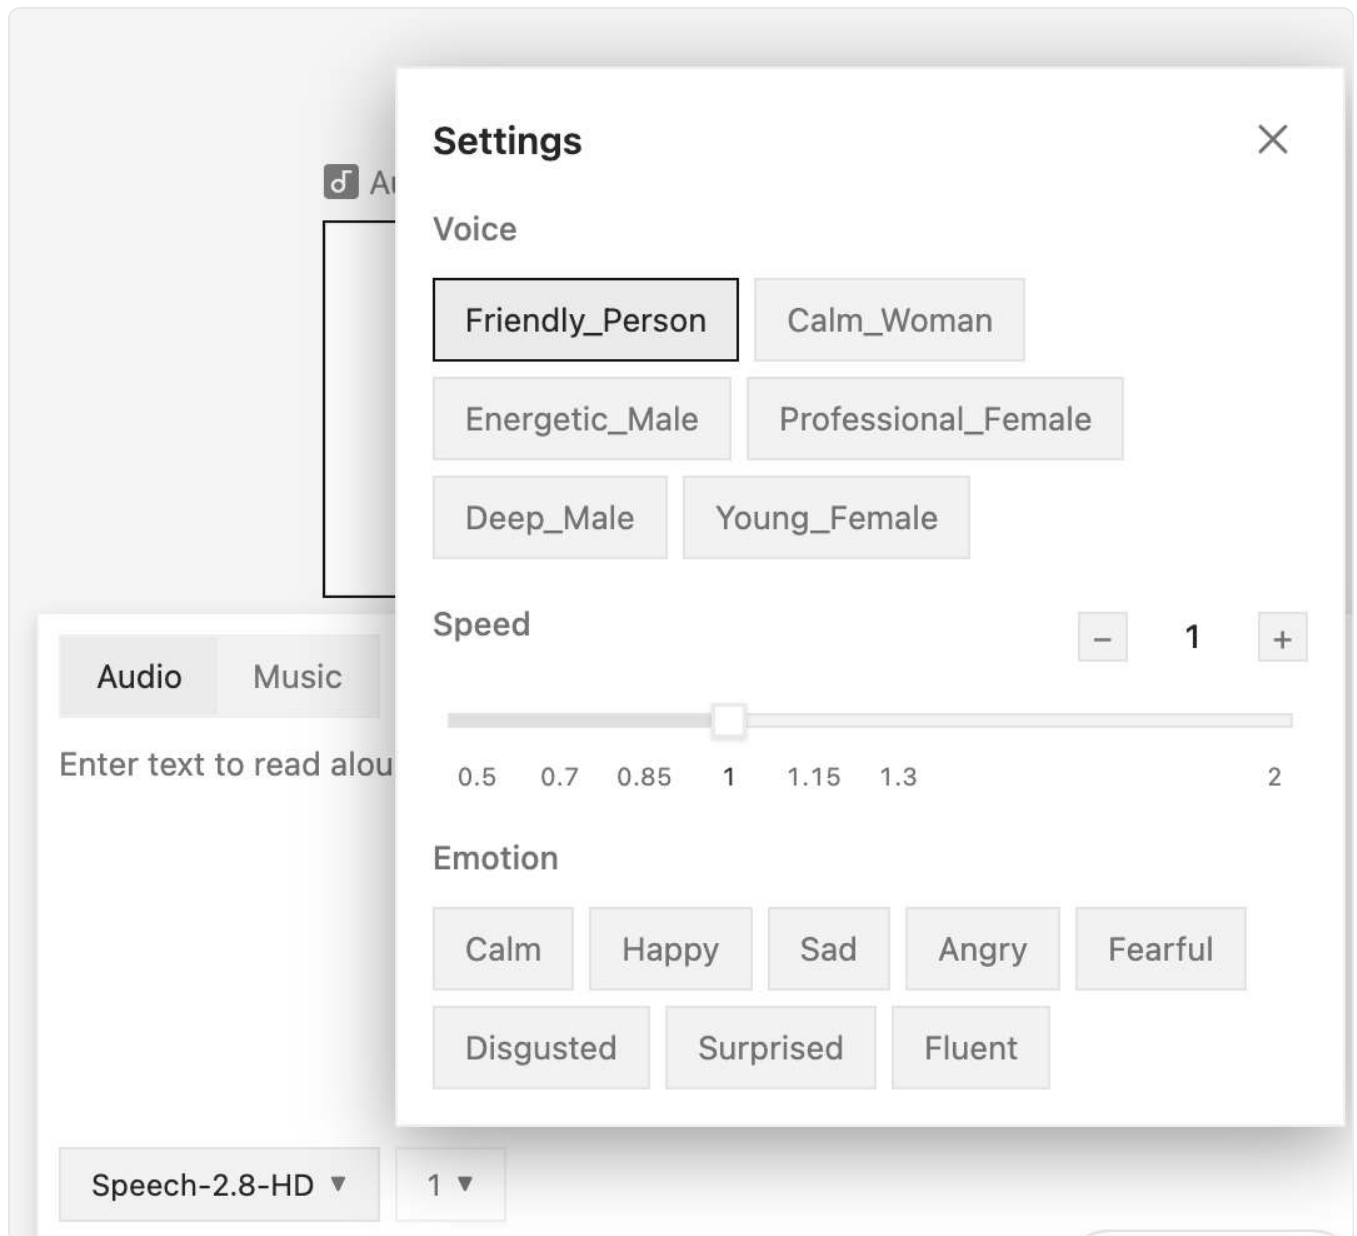


Button	What it does	Credits
Remix	Regenerates a variation from the card's saved prompt (chat-generated videos only)	—
Generate subtitles	Detects speech in the audio track and generates subtitles	—
Add to chat	References this card in your next chat message	—
Add to timeline	Sends the video to the <a href="#">Timeline Editor</a>	—
Clip	Trims the video to the section you choose	—
Extract audio	Pulls the soundtrack out into a new audio card	—
Color adjust	Opens the color adjustment sliders (same as images)	—
Add to Assets	Bookmarks the card into your Assets Library	—
Fullscreen	Opens a full-window preview with playback	—



**Generate subtitles** opens the **Subtitle generation** dialog where you pick the spoken **Language** (**English**, **Chinese**, or **Other**) before transcription starts. **Clip** lets you keep just the good part of a shot before it goes to the timeline, and **Extract audio** is handy when you want to reuse a voice line or soundtrack on its own. **Final cards** — the rendered result of a workflow — share this toolbar minus Remix.

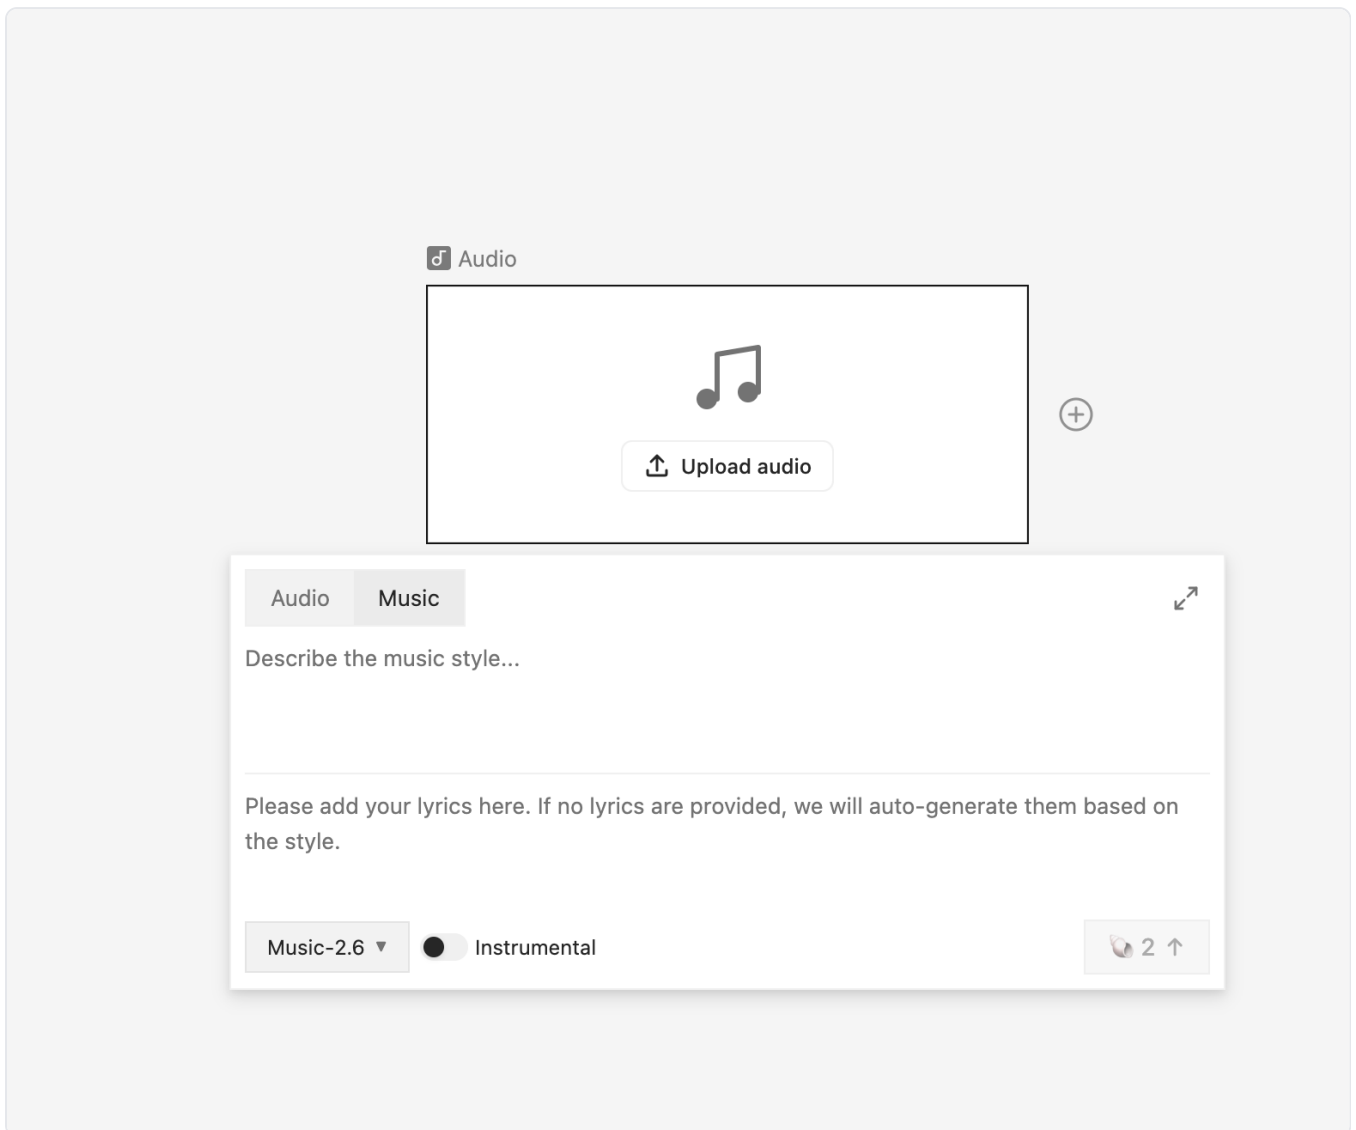
## Audio cards



Button	What it does	Credits
Add to chat	References this card in your next chat message	—
Add to timeline	Sends the clip to the <a href="#">Timeline Editor</a>	—
Clip	Trims the audio	—
Adjust volume	Raises or lowers the clip's loudness	—
Add to Assets	Bookmarks the card into your Assets Library	—
Fullscreen	Opens a full-window player	—

An empty audio card has two tabs: **Audio** for text-to-speech (type or paste the words to be read aloud and pick a voice) and **Music** for full tracks — describe the genre, mood, instruments, and tempo, tap the style

chips to build the description, and optionally click **Generate lyrics** (in your chosen language) or write your own. An **Instrumental** toggle skips vocals entirely, and you can choose **MP3** or **WAV** output.



## Text, Table, Document & script cards

Text cards — including **Table**, **Document**, and the script family (**Plan**, **Episode**, **Scene**, **Intent**) — share one compact toolbar:

Button	What it does	Credits
Copy text	Copies the card's full text to your clipboard	—
Add to chat	References this card in your next chat message	—
Add to Assets	Bookmarks the card into your Assets Library	—
Fullscreen	Opens the text in a full-window reader	—

## Screen cards (App Builder)

Button	What it does	Credits
Export React code	Exports the app as ready-to-run React code	—
Save as PNG	Saves the screen as an image	—
Record walkthrough video	Records a narrated walkthrough of the app	—
Add to chat / Add to Assets / Fullscreen	As above	—

## Game cards

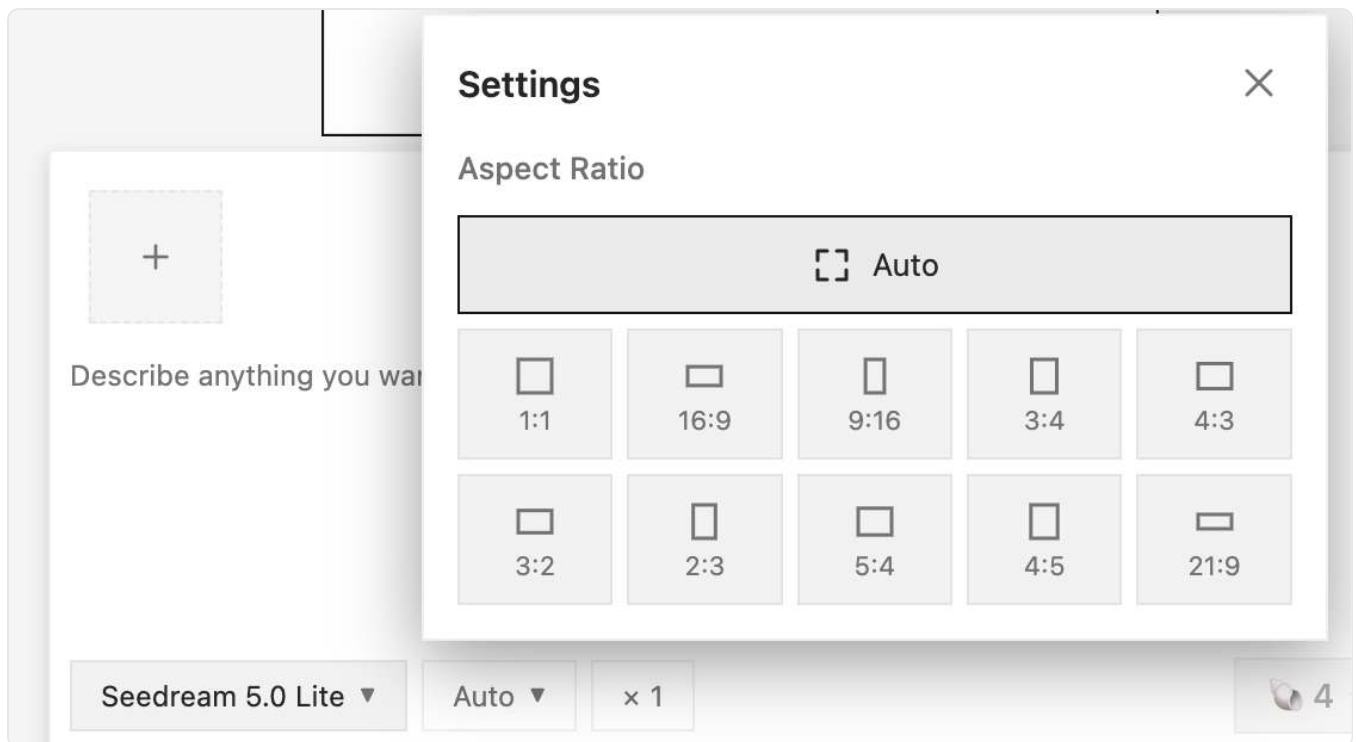
Button	What it does	Credits
Edit / remix game	Re-opens the game in chat so you can change it	—
Export as HTML	Saves the game as a single self-contained .html file	—
Add to chat / Add to Assets / Fullscreen	As above	—

## Document cards (office files)

Button	What it does	Credits
Open	Opens the file with your system's default app	—
Reveal in folder	Shows the file in Finder / Explorer	—
Add to chat / Add to Assets	As above	—

## The generation panel on a card

Select an empty (or previously generated) image, video, or audio card and its **generation panel** appears below it — the same power the chat has, right on the card.



- **Prompt box** — "Describe anything you want to generate". If the card was generated before, its original prompt is pre-filled so you can tweak and re-run.
- **Reference images** — connected image cards show up as thumbnails automatically, and the + button adds more from disk. A colored ring tells you whether the selected model will actually use them.
- **Enhance** — the sparkle button rewrites your prompt with richer detail (and spoken dialogue for video) using your script model; the undo arrow restores your original text.
- **Settings** — the sliders button opens the model's options: aspect ratio, resolution, duration, and whatever else the chosen model supports. Your picked aspect ratio even resizes the card to match.
- **Model** — the dropdown picks which AI model generates. Different models have different strengths, speeds, and prices.
- **Generate** — runs it. On subscription (managed) models the button shows the **estimated credit cost** up front; models using your own API keys just say **Generate**.

Reference Image

First & Last Frame

### Aspect Ratio



16:9



4:3



1:1



3:4



9:16



21:9

### Resolution

480p

720p

1080p

### Duration

4

5

6

7

8

9

10

11

12

13

14

15

### With Audio

With Audio

Silent

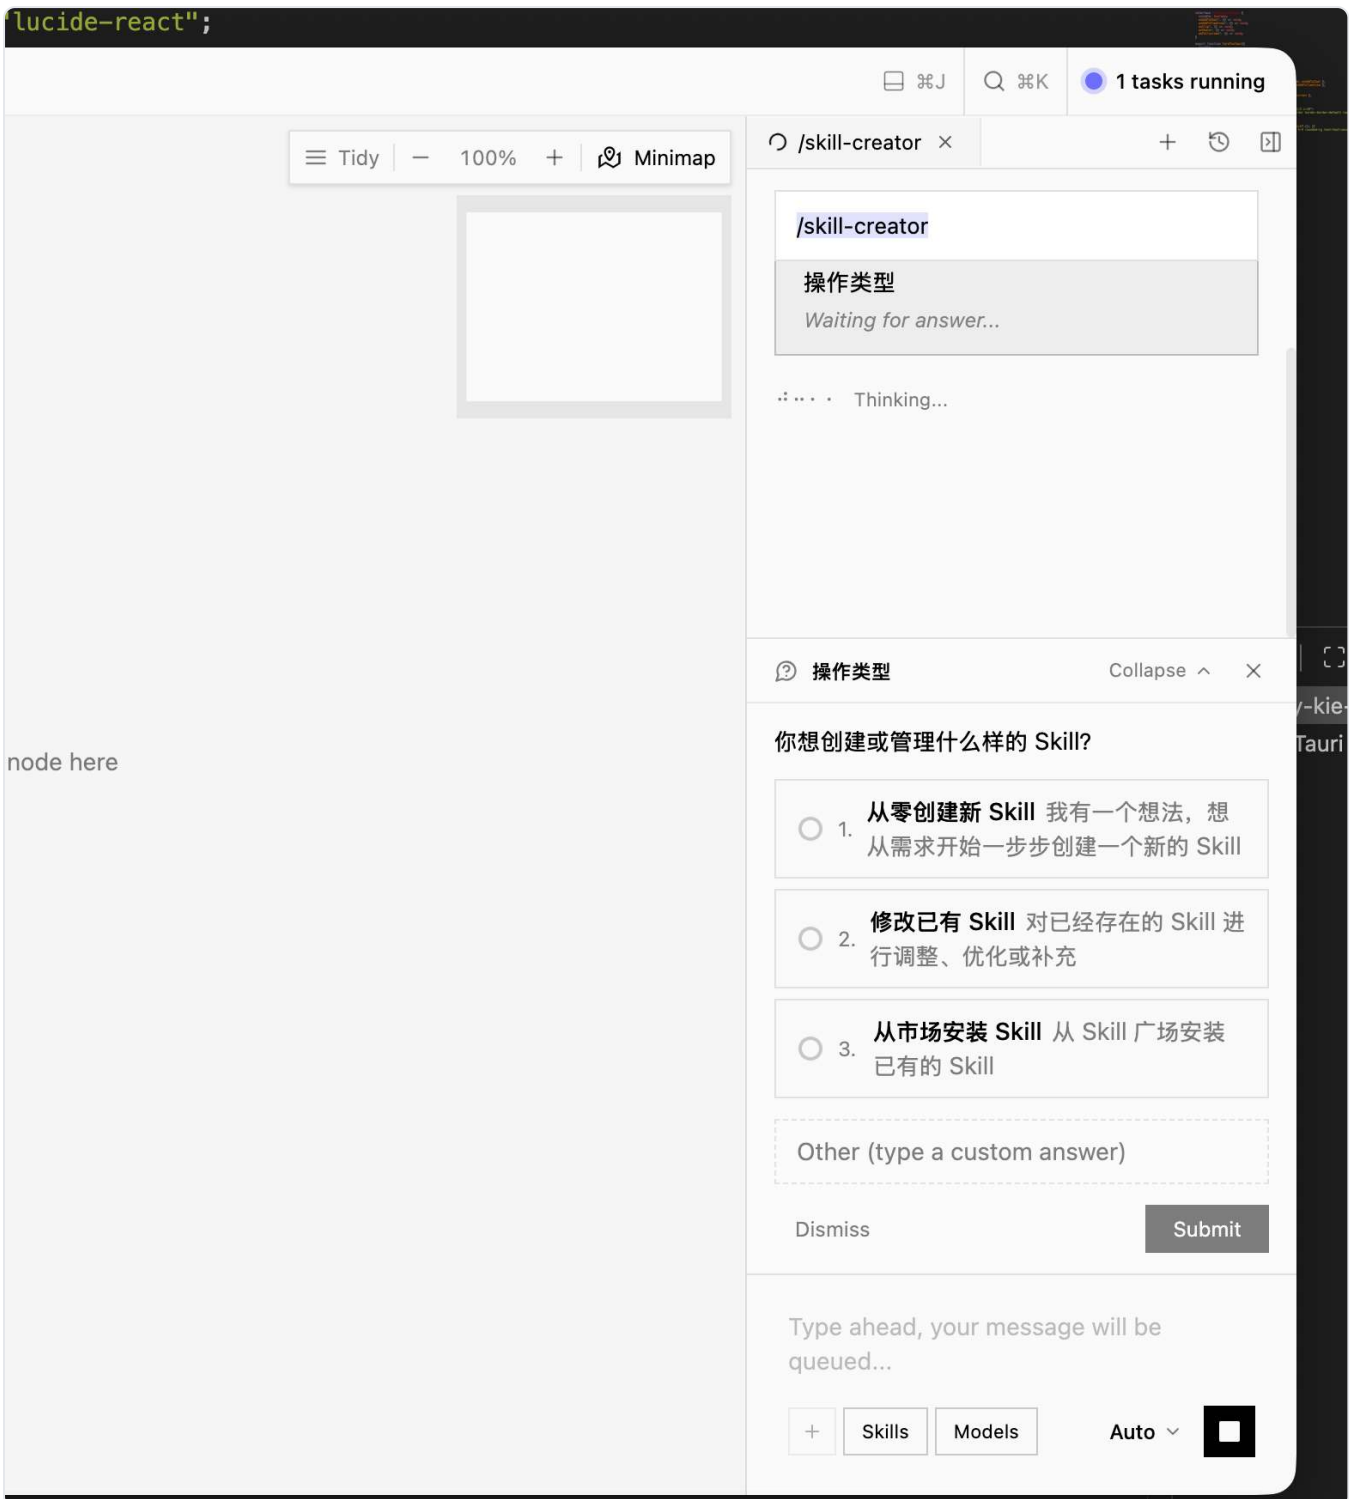
**Tip:** Text cards get a lighter version of this panel — a prompt box that asks your script model to rewrite or extend the card's text in place.

With cards mastered, put them in order in the [Timeline Editor](#), or make your characters permanent in [Cast & Asset Libraries](#).

# Chat & the Agent

Slash commands, session commands, @mentions, attachments, voice input and model selection.

The chat panel on the right side of every project workspace is where you direct the whole production. You describe what you want, and the AI agent does the rest: it plans the work, asks you questions when it needs a decision, generates images, video, audio and text, and places every result as a card on the canvas (see [Canvas & Cards](#)). One conversation can carry a project from a one-line idea to a rendered film.



## Just type — the agent understands intent

You don't need a command to get started. The input reads "**Describe your idea or drop a reference.**" — and that's genuinely all it takes. When you send plain text, the agent works out what you meant:

- **One-shot creative requests** — "make me an image of a red robot", "a 6-second clip of the moon over water" — get generated directly, and the result lands on the canvas.
- **Questions** — "what can you make?", "how does this app work?" — get a plain answer, no generation.

- **Big, multi-stage requests** — "let's make a short drama about cellists" — get routed to the right workflow skill, and the agent offers to start it for you.
- **Casual chat** just gets a friendly reply.

/short-drama A 3D animated film follows two young family friends, Shelly and Enz...

Expand



💡 Thought

🔧 skill { "name": "short-drama... done

💬 Response

I've loaded the short-drama skill. Now let me begin the workflow. You've provided a beautifully detailed story with 7 character reference images! Let me first analyze all your reference images to understand the characters.

**Stage 1 — Intent Recognition**

**Task type:** Short drama screenplay creation (scripting, not video

generation)

**Scale:** Multi-episode short drama series — 3D animated children's adventure film

**Genre:** Children's adventure / family drama — friendship, courage, kindness

**Domain keywords:** "3D animated", "seaside vacation" "resort villa"

Agent is running...



Below the input you'll always see the helper line **"/ for skills, @ for assets."** together with a short notice that you're chatting with an AI system and that generated content may be inaccurate.

## Quick slash commands

Type **/** to open the command menu. Alongside your skills, it lists five quick commands for single generations:

Command	Label	What it does
<code>/image</code>	Create Image	Generate an image
<code>/video</code>	Generate Video	Create a video from your prompt
<code>/audio</code>	Produce Audio	Create music or sound effects
<code>/script</code>	Write Script	Generate a script or text
<code>/all</code>	Run All Agents	Run all agents in parallel

For example: `/image a red sports car at sunset, cinematic light`.

## Skill commands

Every enabled skill is also a slash command — its name in kebab-case, like `/short-drama`, `/promo-video` or `/audiobook`. Type the command followed by your idea:

`/short-drama a lighthouse keeper who finds a message in a bottle`

Type a skill command **on its own** and a gray suggestion appears after it — press **Tab** to accept it and the skill replies with a short usage guide: good example prompts, what to include, and topic ideas. Some skills

(like `/app-builder`) also show clickable **"Try one:"** starter prompts under the input; click one to fill the box with a ready-to-send idea.

The **Skills** button under the input opens the same list as a browsable menu — see [Skills & the Hub](#) for how it's organized.

## Session commands

Workflow skills run as a *session*, and the session is saved automatically — even across app restarts. Four commands control it:

Command	Label	What it does
<code>/continue</code>	Continue Session	Resume the paused skill from where it stopped
<code>/redo &lt;what&gt;</code>	Redo Part	Regenerate only a subset, e.g. <code>/redo the images</code>
<code>/status</code>	Session Status	Show the current skill session + canvas state
<code>/start-over</code>	Start Over	Clear the saved session (canvas is preserved)

`/continue` and `/redo` only appear in the `/` menu when there is actually a paused or restored session to act on. `/status` works entirely offline — it reports without calling the AI. `/redo` is the surgical option:

`/redo scene 3's video` regenerates just that part and leaves everything else on the canvas untouched.

If you close the app mid-run, the next launch restores the session and the chat tells you so — type anything ("continue", "keep going", or a new instruction) and the agent picks up where it stopped.

## When the agent asks you questions

Workflow skills pause at creative decision points, and the question arrives as a card in the chat:

- **Text questions** — just type your answer in the normal input and send.
- **Choice cards** — click one of the option buttons, or use **"Other (type a custom answer)"** to write your own. You can also **Dismiss** a card.
- **Brief forms** — several questions grouped into tabs, with a progress counter like "2/4 answered" and a **Submit** button when you're done.

The run pauses while a question is open and resumes the moment you answer. If a question was pending when you closed the app, the card is re-posted on the next launch so you can answer it and continue.

## @mentions — reference your assets and Cast

Type @ in the input to open the mention menu. It lists your saved asset library and your Cast members (reusable characters with a locked face and voice — see [Cast & Asset Libraries](#)), filterable by category: **All**, **Characters**, **Scenes**, **Styles**, **Props**, **Backgrounds**, **Music**, **Scripts**.

Picking an entry inserts its @handle into your message and attaches its file, so the agent uses it as a real reference — mention a Cast member and generations keep that character's exact face and voice. Attached images can also be mentioned by position as `@image1`, `@image2`, ...; they render as small inline thumbnails right in your prompt, which is the clearest way to tell a model "put *this* subject into *that* scene".

## Attachments

The **+** button (**Attach image or file**) opens a file picker. You can attach:

- **Images** — PNG, JPG/JPEG, WebP, GIF
- **Video** — MP4, MOV
- **Audio** — MP3, WAV
- **Documents** — PDF, TXT, MD, DOCX, XLSX/XLS, CSV, ODS

Attached images are copied into the project's local assets folder the moment you send, so they stay valid forever — a reference photo you attach today still works when you resume the session next week.

Documents can be read by the agent (with the Documents skill or in free chat), and any attachment shows as a removable chip above the input before you send.

## Voice input

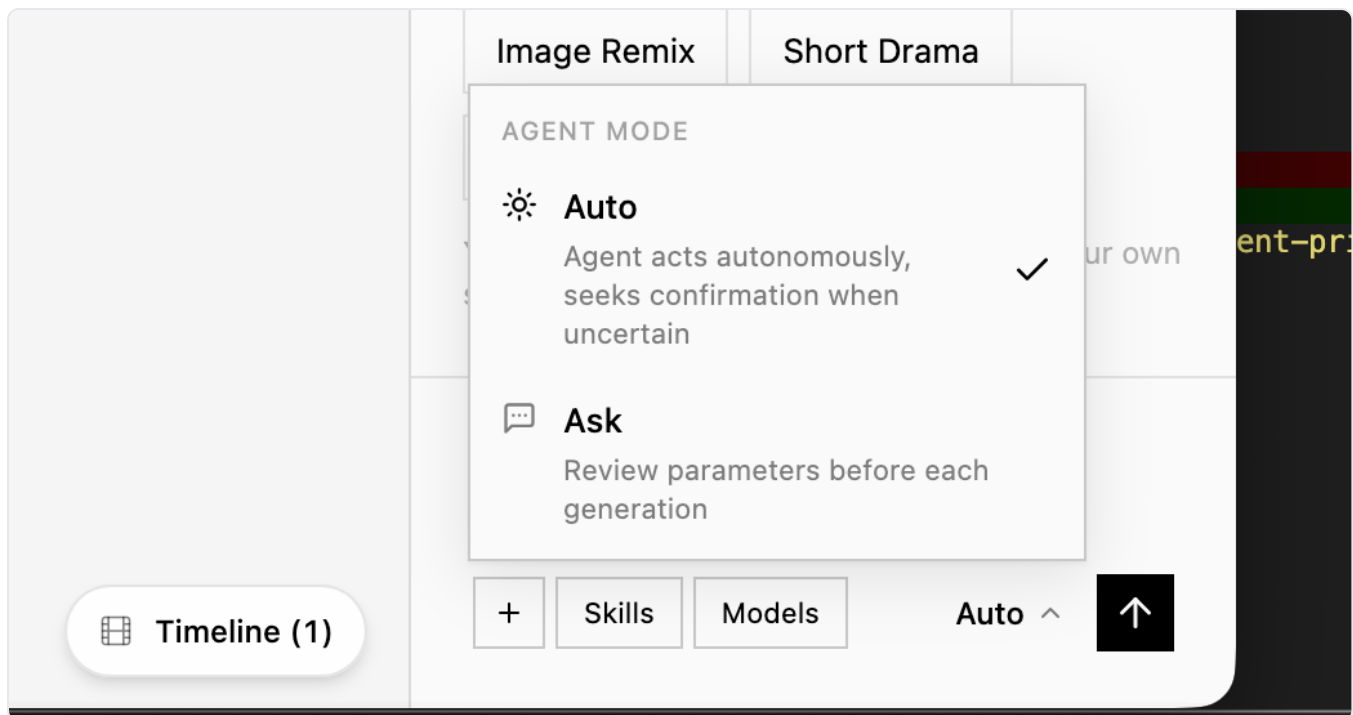
The microphone button records a voice command: click to start, click the square to stop. The recording is transcribed and **placed into the input for review — it is never sent automatically**, so you can edit before pressing Enter.

## Models and Agent mode

The **Models** button opens the model picker, with tabs for **Image, Video, Voice, Music and SFX**. Tick the models the agent is allowed to use — the footer reminds you that the *agent uses only selected models*. Each entry shows its provider and capability badges (max resolution, clip length).

Next to the send button sits the **Agent Mode** selector with two modes:

- **Auto** — Agent acts autonomously, seeks confirmation when uncertain.
- **Ask** — Review parameters before each generation.



Ask mode is great while you're learning what each model costs and produces; Auto is the hands-off mode for long workflow runs.

## Long runs, pauses and credits

A few practical notes for big productions:

- **Safety limit.** The agent has a built-in cap on how many tool steps it may take in one run. If it hits the limit mid-project, nothing is lost — your work is on the canvas, and typing anything (even just "continue") resumes from that exact point.
- **Provider pauses.** If a provider runs out of credits, rejects a key, or rate-limits, the workflow pauses with an explanation instead of failing. Fix the cause (top up, check Settings → AI Providers, or just wait a minute), then type `/continue`.
- **Credits.** Every image, video and audio generation consumes credits (or your own provider balance) — multi-scene workflows can add up, so keep an eye on the balance in the top bar.

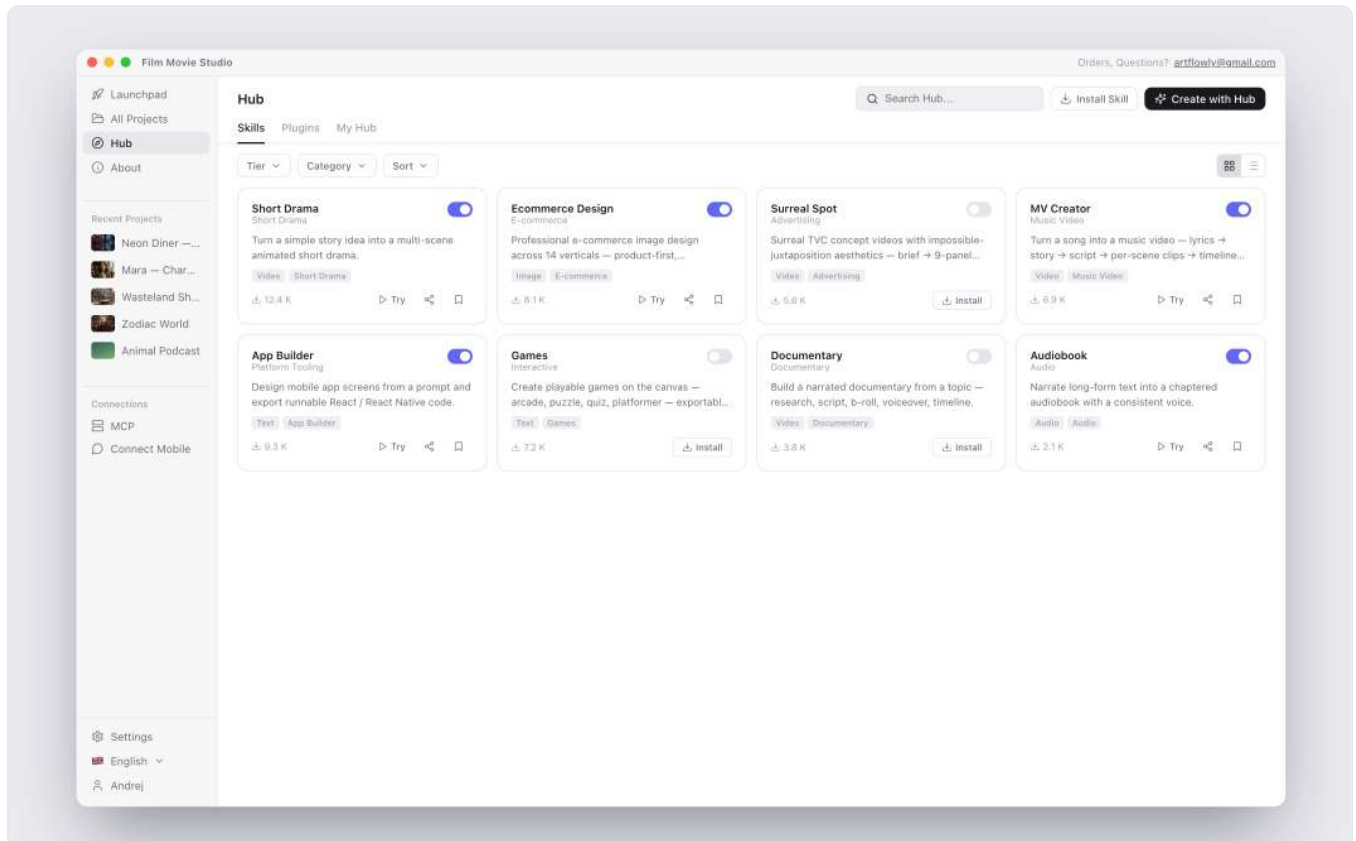
**Tip:** You don't have to babysit a workflow. Start `/short-drama`, answer the brief questions, and come back later — the session, the canvas, and any pending question will all still be there.

Ready to go deeper? Browse every skill in [Skills & the Hub](#), or jump straight to the full [Skills Reference](#).

# Skills & the Hub

Browse, enable and manage skills; workflow vs quick skills; creating your own.

Skills are guided AI workflows — playbooks written in plain language that the chat agent follows step by step. A skill knows the right order of work for its craft (characters before keyframes, keyframes before video), where to pause and ask for your approval, and which generation tools to use. You run one by typing its slash command in the chat, and everything it makes lands on the canvas.

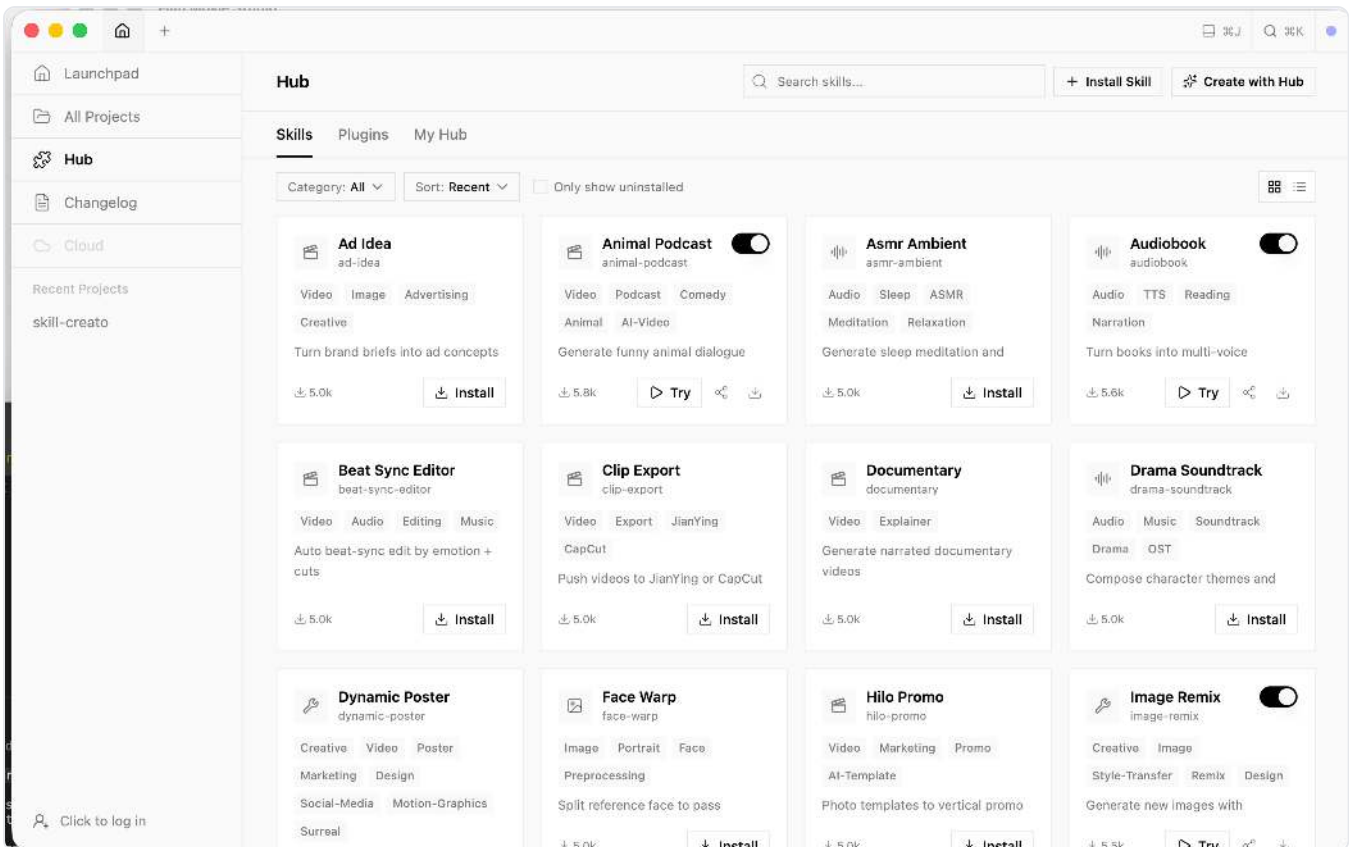


Skills come in two tiers:

- **Workflows** — multi-stage productions like **Short Drama**, **MV Creator**, **Documentary** or **Promo Video**. They plan, ask questions, generate in stages, and often end with a rendered video.
- **Quick** — single-purpose tools like **Quick Image**, **Quick Video** and **Quick Audio**: one prompt in, one result out.


## The Hub page

Open the **Hub** from the sidebar to browse everything. The top bar has a **Search Hub...** box plus two buttons — **Install Skill** and **Create with Hub** — and three tabs: **Skills**, **Plugins**, and **My Hub** (your installed skills).



The filter row narrows the grid:

- **Tier** — All, Workflows, or Quick.
- **Category** — All, Advertising, E-Commerce, Short Drama, Music Video, Documentary, Audio Content, Visual Production, Platform Tooling.
- **Sort** — Recent, Popular, or Name.
- **Only show uninstalled** — hide what you already have.
- A grid/list toggle switches the layout; long lists paginate 16 per page.

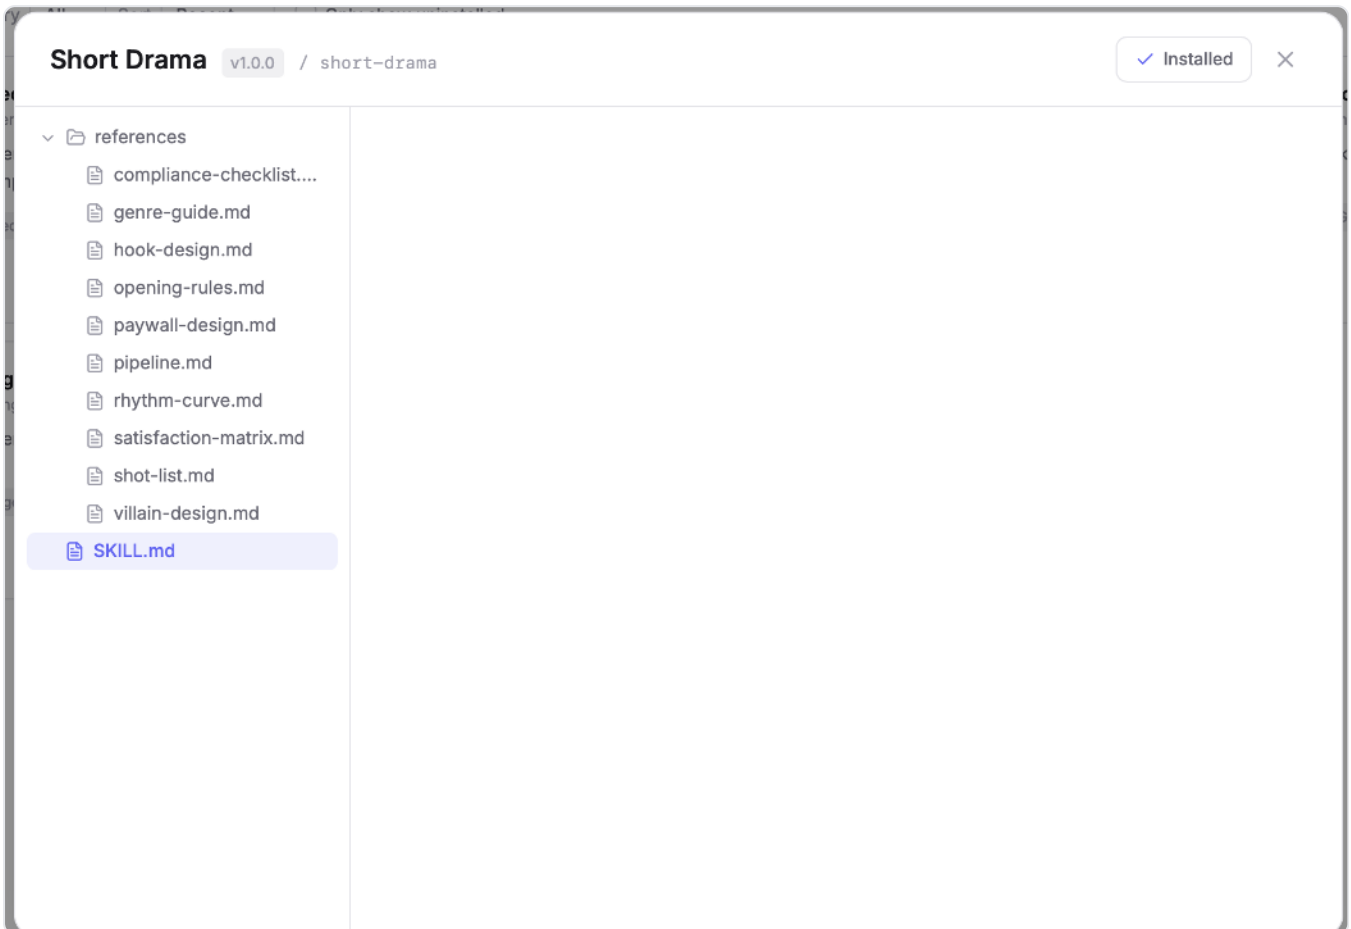


## Skill cards and their actions

Each card shows the skill's name, category, description and tags. What you can do depends on its state:

- **Not installed** — an **Install** button adds it to your library.
- **Installed but disabled** — the card says "**Enable to use** →"; flip the toggle in the card's corner to enable it.
- **Enabled** — you get **Try** (creates a new project with the skill command pre-filled and ready to run), **Share**, and **Save**.

- **Update available** — an accented **Update** button appears with a version chip like `v1.0.0 → v1.2.0`, so you can see what you're moving to before clicking.



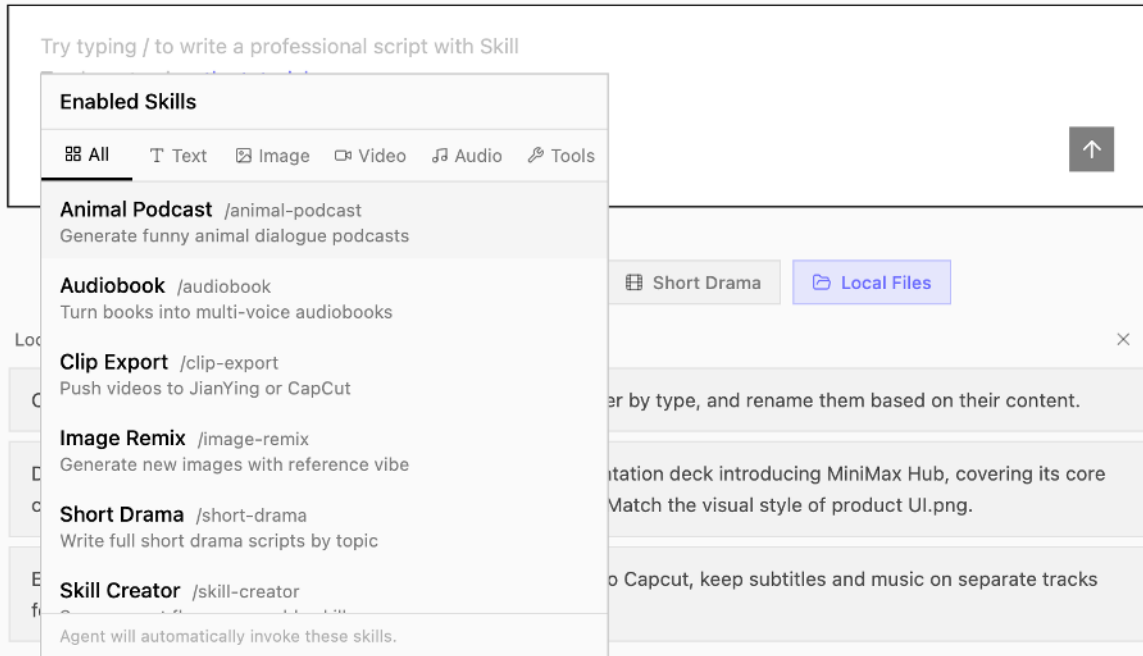
Only *enabled* skills appear in the chat's `/` command list and Skills menu — disabling a skill hides it there without uninstalling anything.

## The Skills menu in chat

In any project, the **Skills** button under the chat input opens the **Enabled Skills** menu. Category tabs — **All**, **Text**, **Image**, **Video**, **Audio**, **Tools** — filter the list, which is grouped into a **Workflows** section on top and **Quick Generation** below. Click a skill and its `/s lug` command drops into the input, ready for your prompt. The footer notes that the agent will also invoke these skills automatically when your request calls for one.

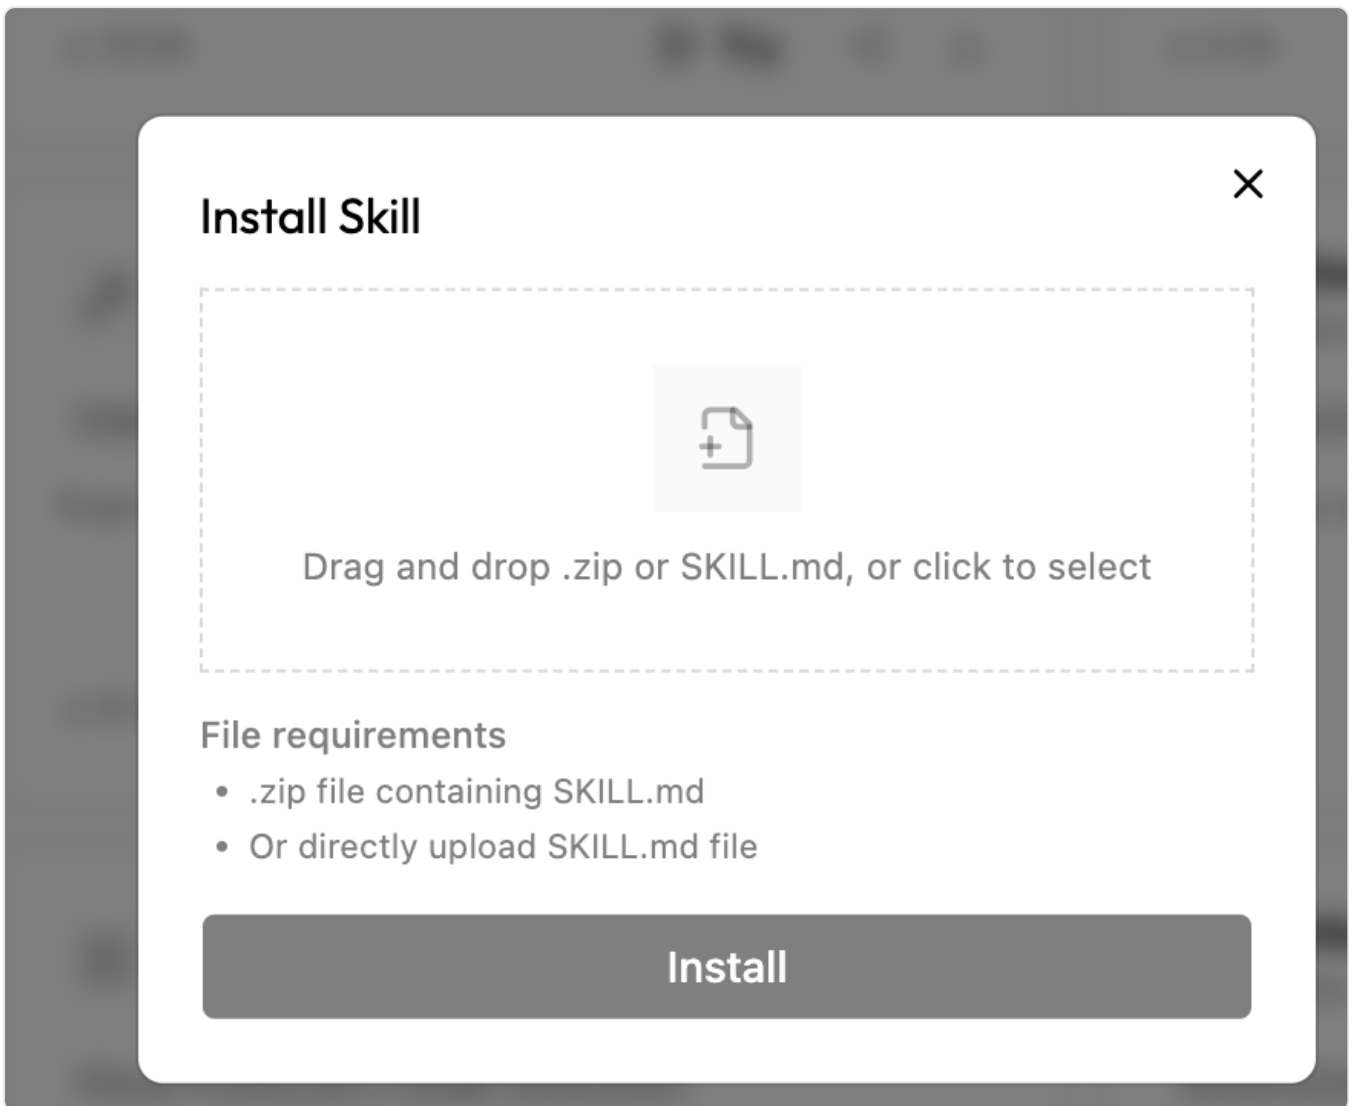
# MiniMax Hub

A new way to create media with AI



## Installing a skill package

**Install Skill** (in the Hub's top bar) opens a dialog where you can drop — or click to pick — a skill package: either a **.zip** bundle or a plain **SKILL.md** file. Once installed, it shows up in My Hub and as a slash command like any bundled skill.



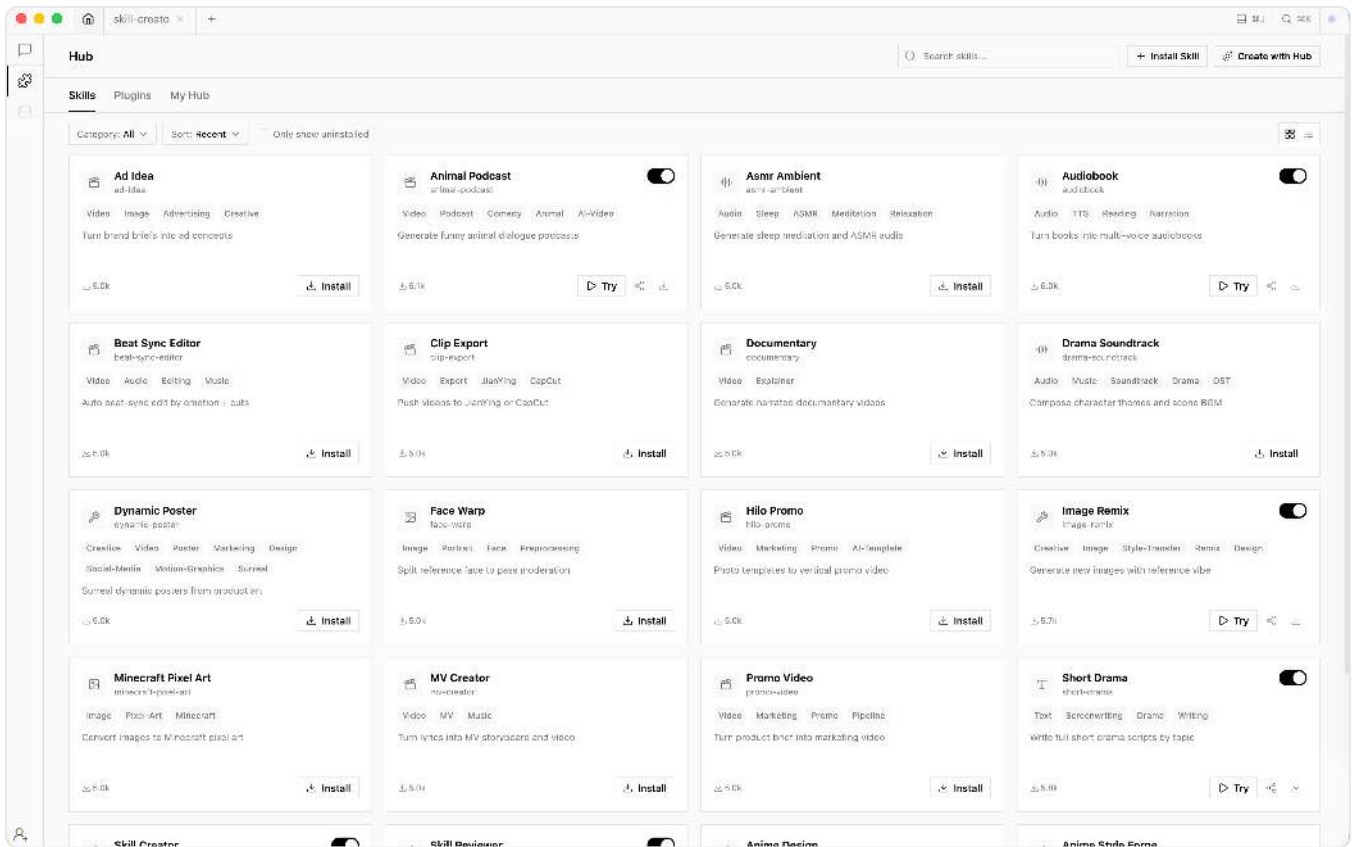
**Tip:** You can also drag a `.md` skill file straight into the chat panel to import it — handy when someone shares a skill with you.

## Creating your own skill

**Create with Hub** opens the **Create a Skill** dialog: give it a **Name** and **Description**, pick a **Category** and an agent type (**Text**, **Image**, **Video** or **Audio**), and click **Create**. New skills start in the Quick tier and immediately get their own slash command. Because skills are written in plain language, editing one is closer to writing a brief than programming.

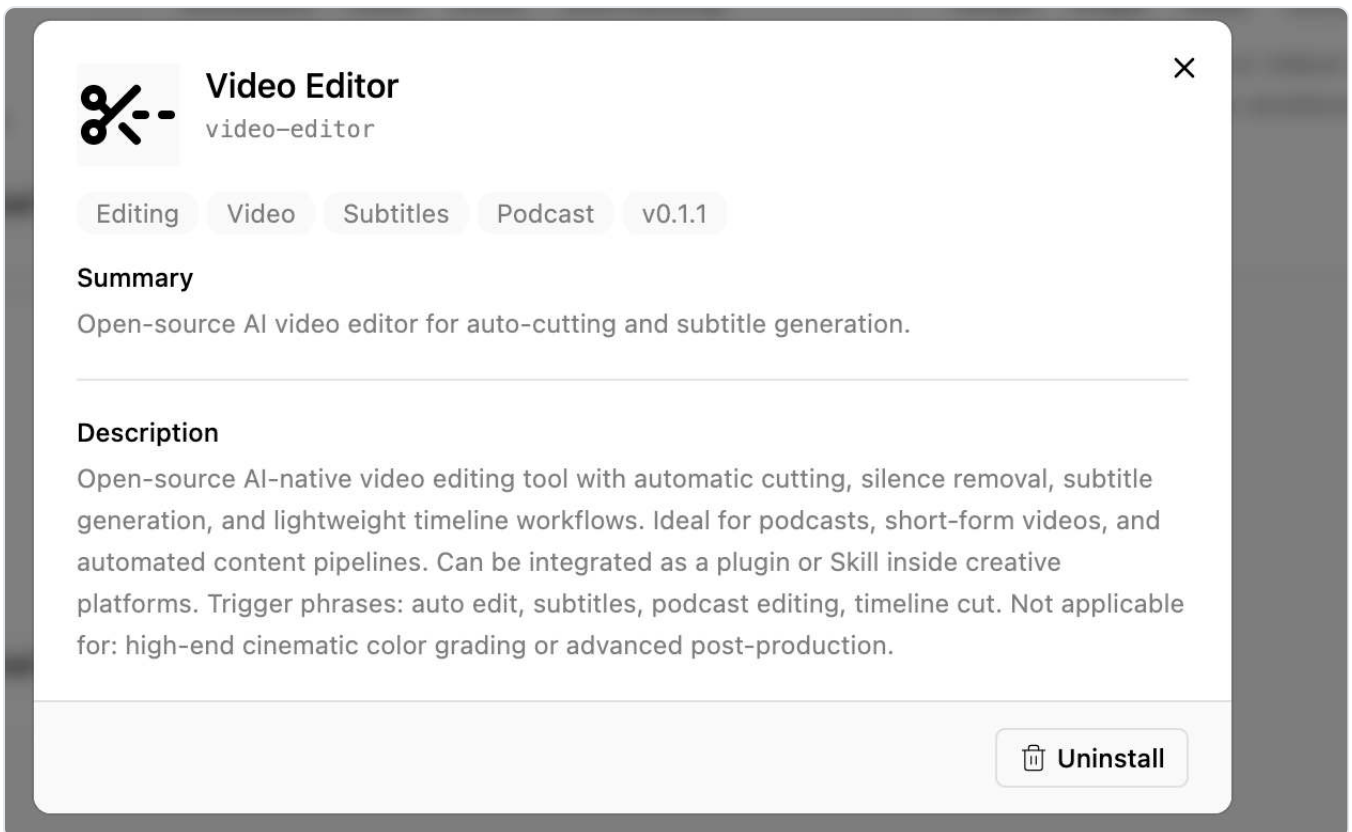
## Managing your skills

The **My Hub** tab is your installed collection. From here (or anywhere in the Hub) you can toggle skills on and off, apply updates as new versions ship with the app, and keep the chat's command list down to the skills you actually use.



## Plugins

The **Plugins** tab holds six bundled tools: **Multi-Panel Storyboard**, **Auto Color Grade**, **Voice Cloner**, **Scene Detector**, **Subtitle Burner**, and **Multi-Angle Generator**. Each has the same enable toggle as a skill; click a plugin card to open its detail view with the full description, tags, version — and a button to disable it (you can re-enable it any time from its card).

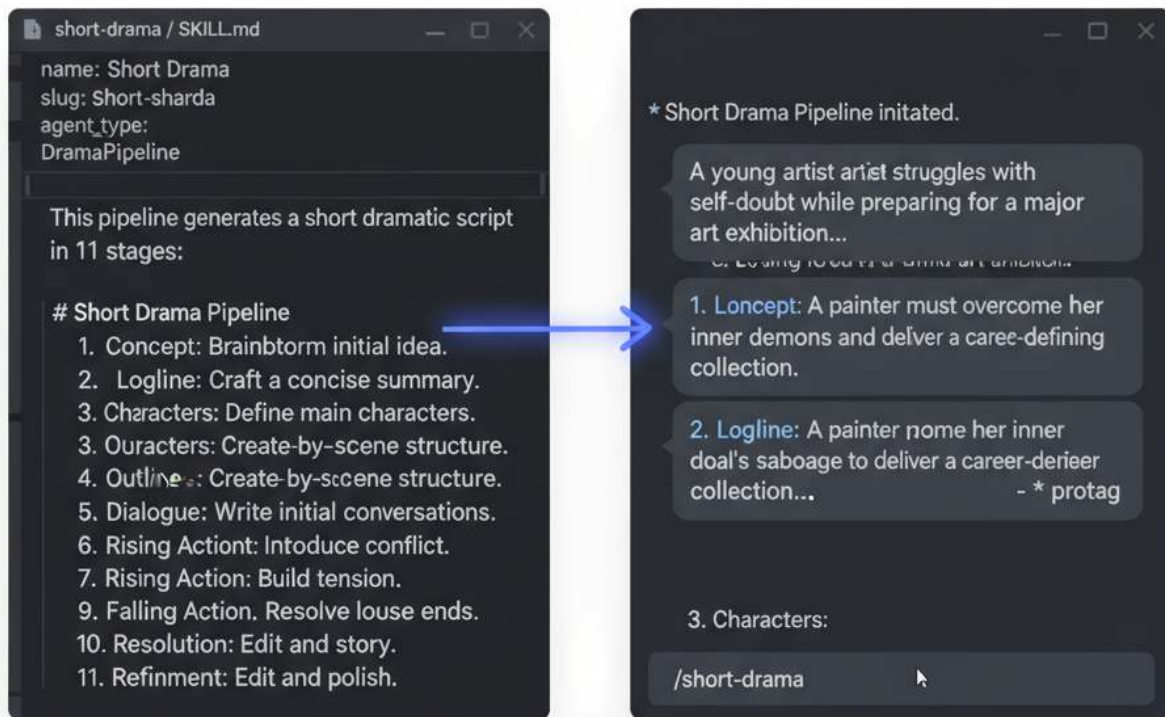


## How a skill actually runs

When you send `/short-drama a story about...`, the agent reads that skill's playbook and works through it one step at a time: it writes and generates, pauses to ask you questions at the decision points, places every intermediate result on the canvas, and — for video workflows — finishes by assembling the timeline and rendering. You stay in control the whole way: answer the questions, redo any part with `/redo`, or stop and resume days later.

# YOUR SKILLS. EDITABLE. FORKABLE

Same architecture artflowly@gmail.com uses for Claude Code  
— bring your own SKILL.md, share them, sell them



Every card a skill creates behaves like any other card — remix it, add it to the timeline, or save a character to your Cast so the next skill run reuses the same face and voice (see [Cast & Asset Libraries](#) and [Canvas & Cards](#)).

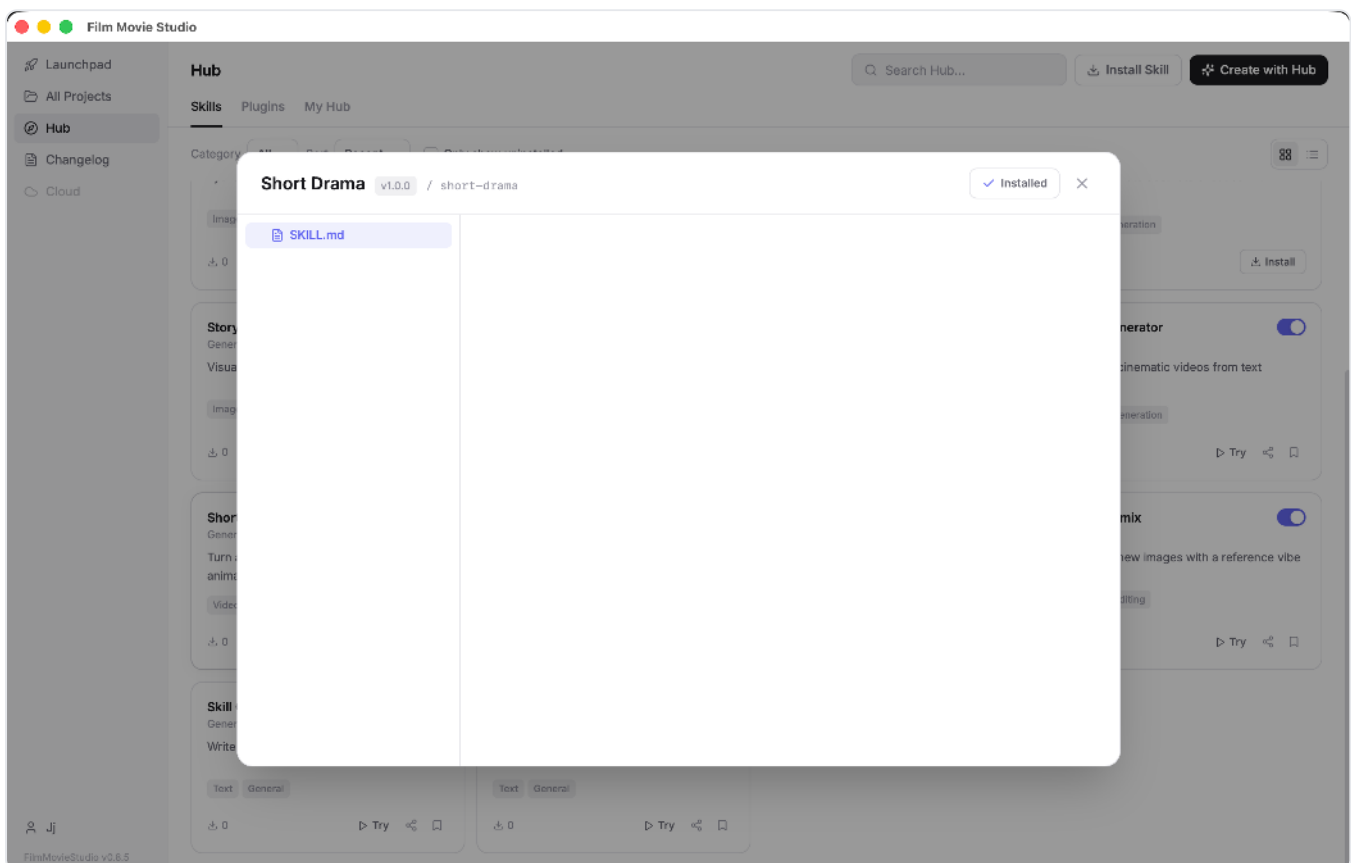
For a complete catalog of every bundled skill — what it makes, its stages, and example prompts — see the [Skills Reference](#).

# Skills Reference

Every bundled skill explained — what it does, an example prompt and how it runs, step by step.

Skills are FilmMovieStudio's ready-made creative pipelines: type a slash command in the chat, describe what you want, and the AI director walks the whole job — asking you the creative questions, generating the images, clips, audio, and documents onto your canvas, and (for video skills) rendering the final cut. This chapter is the complete reference for every bundled skill: what it makes, a realistic example prompt, and how a run actually unfolds.

There are two tiers. **Workflow** skills are multi-stage productions — they pause at creative gates so you approve the plan, the style, the keyframes. **Quick** skills are one-shot tools: one prompt in, one result out, no questions. You can run any skill three ways: type its `/command` in the chat, pick it from the Skills menu next to the chat input, or open it from the Hub page and click Run.



**Tip:** A workflow skill remembers where it was — even across restarts. Type `/continue` to resume an interrupted run, `/status` for a local progress report, `/redo <what>` to regenerate just one part, or `/start-over` to clear the session.

## Contents

- **Short Drama & Storytelling** — [Short Drama](#) · [Animal Podcast](#) · [Continue Story](#) · [Zodiac World](#) · [Wasteland Short Film](#)

- **E-commerce & Product** — [Ecommerce Design](#) · [E-commerce Image](#) · [Detail Page Designer](#)
- **Advertising** — [Ad Idea](#) · [Promo Video](#) · [Surreal Spot](#) · [Dynamic Poster](#) · [HiLo Promo](#) · [YouTube Creator](#)
- **Music Video** — [MV Creator](#) · [Sword Dance](#) · [Beat Sync Editor](#)
- **Documentary & Explainer** — [Documentary](#)
- **Visual Production** — [Avatar & Figure Studio](#) · [Storyboard](#) · [Poster Design](#) · [Anime Design](#) · [Anime Style Forge](#) · [Film Assets](#) · [Film Shot](#) · [Character & Scene Sheet](#) · [Face Warp](#) · [Film Style Picker](#)
- **Audio & Voice** — [ASMR Ambient](#) · [Narrate Scene](#) · [Drama Soundtrack](#) · [Audiobook](#) · [Voice Clone](#)
- **Games & Interactive** — [Games](#) · [3D Game Studio](#)
- **Platform Tools** — [App Builder](#) · [App Walkthrough](#) · [Documents](#) · [Automation](#) · [Skill Creator](#)
- **Quick Tools** — [Quick Image](#) · [Quick Video](#) · [Quick Audio](#) · [N-Grid Storyboard](#) · [Image Remix](#) · [Minecraft Pixel Art](#) · [Multi-Shot](#) · [Image Paint Edit](#) · [Relight](#) · [Video Prompting](#) · [Camera Path Prompter](#) · [Clip Export](#) · [Skill Reviewer](#) · [Voice Design](#)

## Short Drama & Storytelling

### Short Drama

The flagship skill. `/short-drama` is a full AI filmmaking director: it takes a one-line story idea and walks it through script analysis, character design, a locked visual style, episode plans, scene keyframes, animated clips, audio, and a final rendered cut — pausing at every meaningful creative gate so *you* make the calls.

**Example prompt:** `/short-drama A jaded Tokyo detective adopts a stray cat that can smell lies`

#### How it runs:

1. **Intent & characters.** The director reads your idea (and any reference photos you attached), identifies genre, tone, and audience, and writes a short profile card for each main character.
2. **The creative brief.** It pauses and asks you a tabbed brief shaped to *your* story: target market, episode count, and tone & style — each tab with concrete options to pick from.

/short-drama A 3D animated film follows two young family friends, Shelly and Enz...

Expand



🔍 Target Market 1/3

Collapse ^ x

**Target Market**

Episode Count

Tone & Style

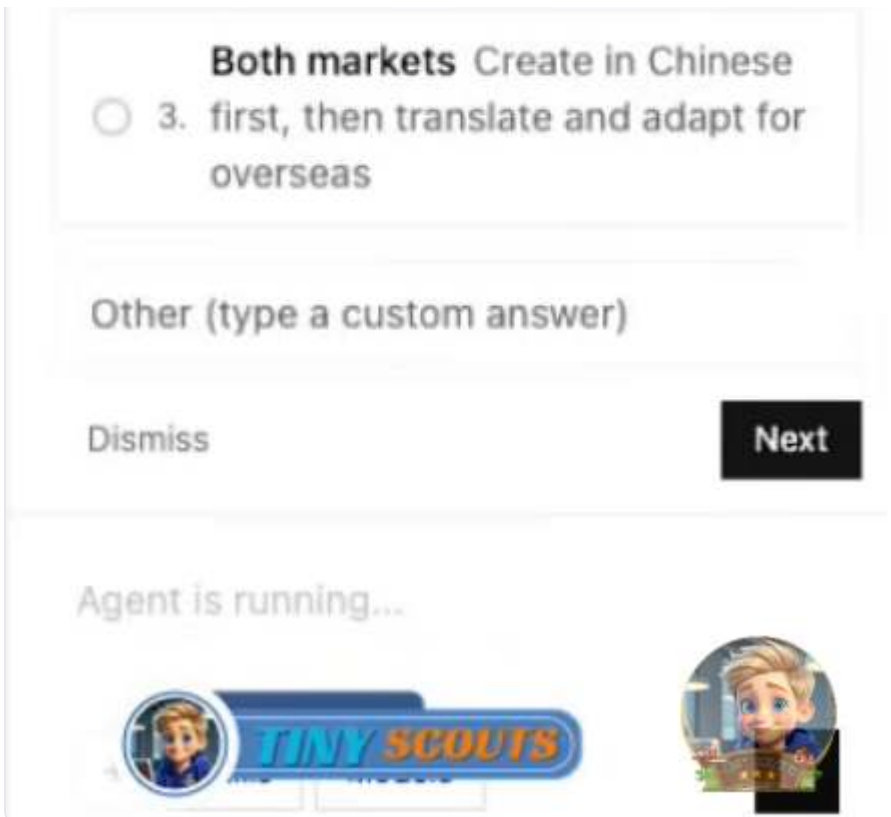
What is your target market for this short drama? This determines the script format and cultural adaption approach.

**Overseas (ReelShort/DramaBox) (Recommended)** English script,

1. Hollywood format (INT./EXT.), optimized for international platforms

**Domestic (China)** Chinese script,

2. vertical-screen format, optimized for Douyin/Kuaishou



- 3. Style lock & portraits.** From your style pick it writes a binding "Style Bible", generates a style anchor image, then a reference portrait for every character — so the look and the faces stay identical across every future shot. It offers to save your leads to the Cast library so they keep the same face and voice in future projects.
- 4. Creative plan & episodes.** It drafts the act structure, one card per episode, and three scene beats per episode — then pauses for you to approve the plan or change the brief.

/short-drama A 3D animated film follows two young family friends, Shelly and Enzo...

Expand



### Act I — Setup (Ep 1-3)

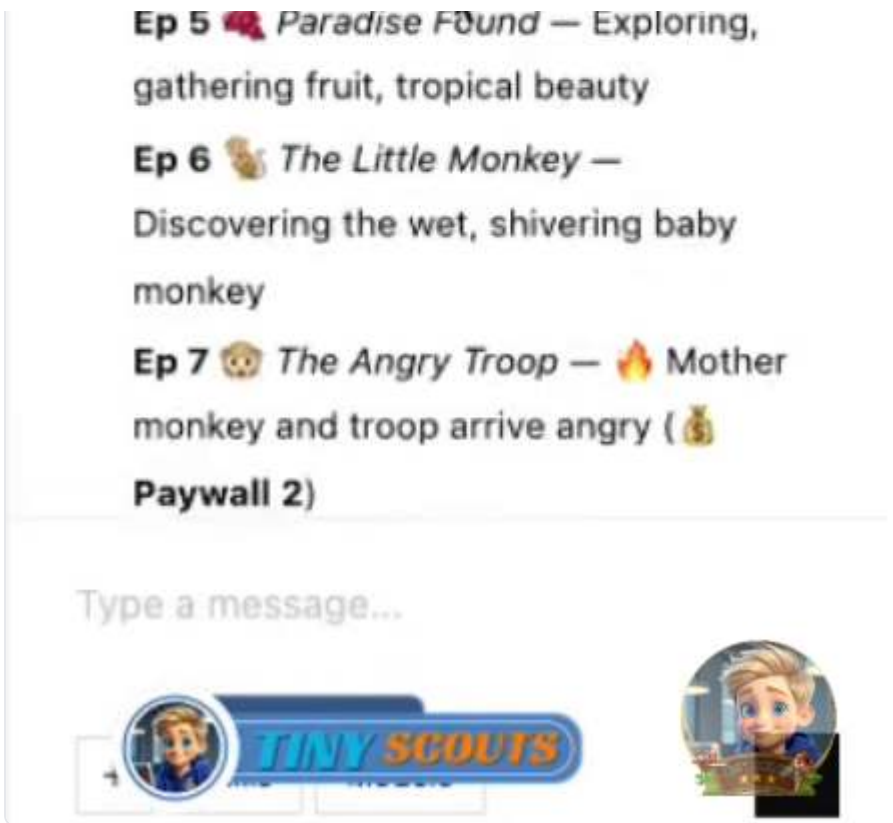
**Ep 1** 🌴 *Vacation Paradise* — Meet the families at the beautiful seaside resort

**Ep 2** 🚤 *The Little Boat* — Enzo discovers a boat, convinces Shelly to sail

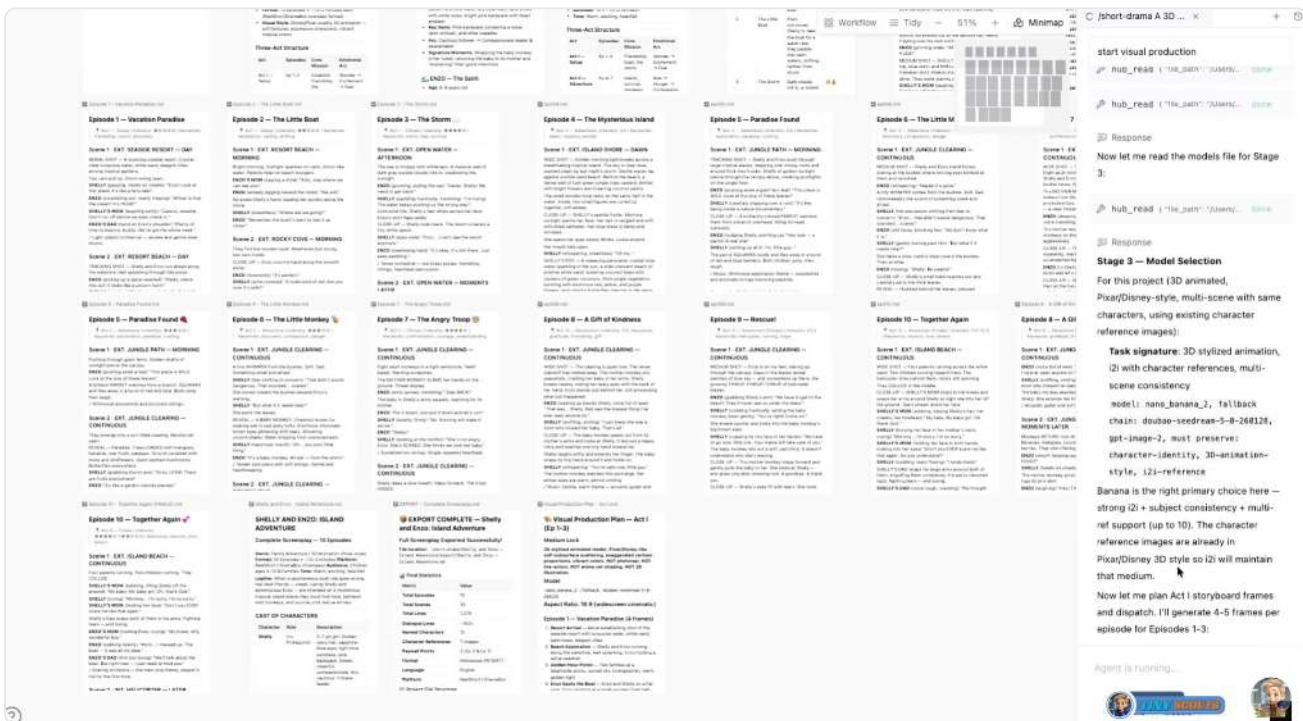
**Ep 3** 🌪️ *The Storm* — 🔥 Weather turns deadly, massive waves, the children hide in the boat (💰 **Paywall 1**)

### Act II — Adventure (Ep 4-7)

**Ep 4** 🏝️ *The Mysterious Island* — Morning after, Shelly discovers the island



5. **Keyframes & production brief.** After approval it generates a keyframe image for every scene (character portraits attached as references), asks you to confirm them, then asks a second tabbed brief: how much to produce now, what kind of audio bed, and how long each clip should be.



6. **Clips, audio, render.** Each keyframe is animated into a video clip (dialogue scenes get lip-synced speech), music is scored to your audio pick, and after a final "Open timeline / Render now" gate the whole thing is composited into an MP4 with crossfades and dropped on the canvas as the Final Render.

Thinking...

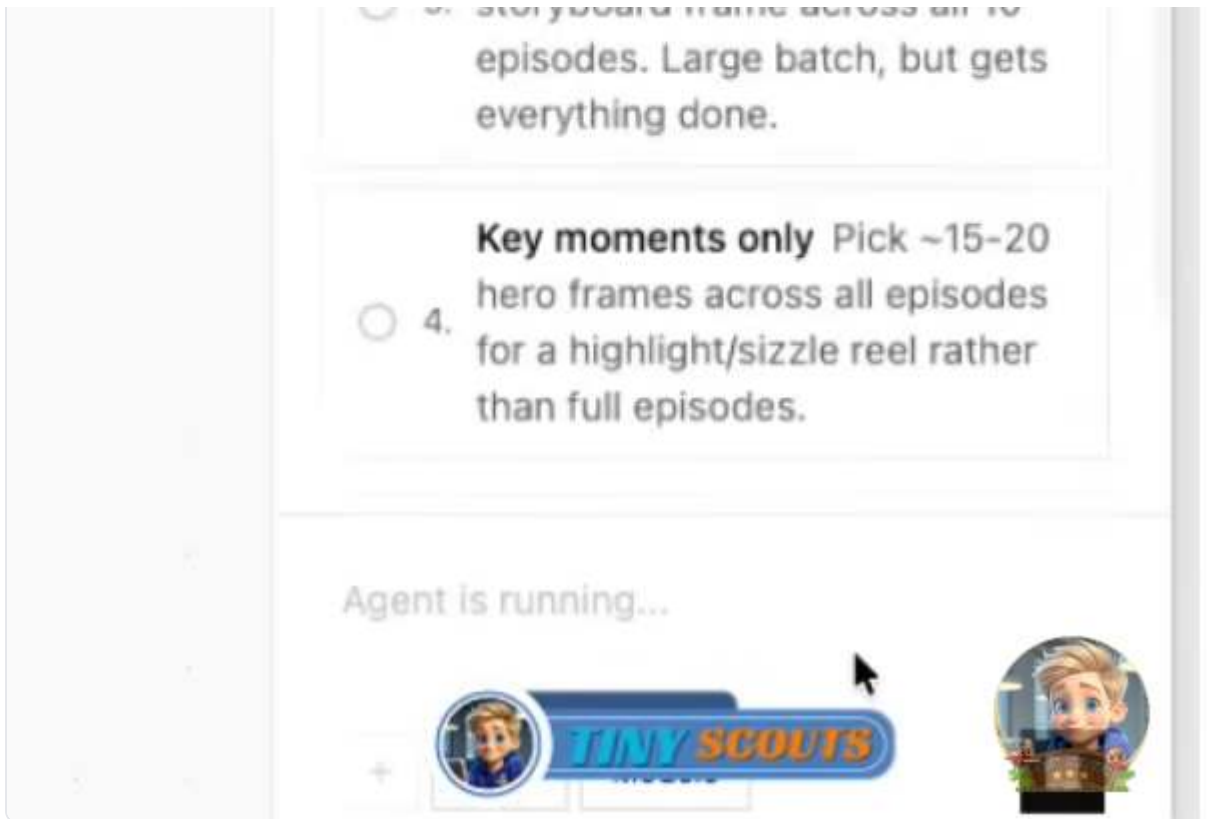
Scope 1/3

Collapse ^ x

**Scope** Audio Clip Duration

We have 45 storyboard frames across 10 episodes. How would you like to approach video generation?

- 1. **Pilot episode first (Recommended)** Start with Episode 1 (5 frames → 5 video clips) as a proof of concept. Review quality, then continue with remaining episodes.
- 2. **Full Act I (Ep 1-3)** Generate video clips for all 13 frames in Act I. Good for testing the full storm arc before committing to all 45.
- 3. **All 45 frames** Go all-in — generate video clips for every storyboard frame across all 10



**Tip:** The more flavor in your opening line — genre, setting, a twist — the better the brief options it invents for you. You can attach photos with the first message and the director will design the characters from them.

## Animal Podcast

`/animal-podcast` makes a comedic 30–39 second short where two contrasting animal hosts roast a human-world topic from an animal point of view — sized for TikTok, Shorts, and Reels, with burned-in subtitles and pop-text.

**Example prompt:** `/animal-podcast` Make a 9:16 podcast where a Shiba Inu and an orange cat discuss why humans pay for gym memberships. Tone: sarcastic but cute

### How it runs:

1. Designs two contrasting hosts (or offers you 3–4 fun pairings if you didn't pick animals) and asks for your aspect ratio (9:16 or 16:9).
2. Writes the comedy script — scene beats and dialogue lines — and checks the format before continuing.
3. Generates a style-locked portrait of each host to anchor the visuals.
4. Produces per-line voiceover, two ~15-second animated clips, background music, and reaction sound effects (laugh, surprise, applause...).
5. Assembles the timeline, burns in subtitles and pop-text, and renders the final vertical video.

**Tip:** The best topics are human habits animals would misunderstand — work meetings, dating apps, expensive coffee, productivity culture.

## Continue Story

`/continue-story` picks up an existing project and generates the *next* scene — visually continuous with the last clip, reusing your locked style and character portraits. One scene per run, so you stay in control of the plot.

**Example prompt:** `/continue-story The rivals are forced to share one umbrella when the storm finally breaks`

### How it runs:

1. Reads your canvas — finds the latest clip, the character portraits, and the story so far.
2. Extracts the last frame of the previous clip as a continuity anchor.
3. Writes the next scene beat (or proposes one and asks you to confirm if you gave no direction).
4. Generates a keyframe referencing the last frame and the character portraits, then animates it into the next clip.
5. Asks whether you want another scene, to open the timeline, or to render.

## Zodiac World

`/zodiac-world` produces a 30–60 second vertical children's English-learning episode (ages 3–8) set in the zodiac forest village of Kazefumori — a warm-3D story plus a sing-along closing song, with one core English phrase repeated so it sticks.

**Example prompt:** `/zodiac-world Rabbit learns to say "May I borrow this?" after taking Monkey's kite without asking`

### How it runs:

1. Asks a short brief: which zodiac friend to feature (Rabbit / Tiger / Rat / Dragon), the age band (3–5 or 6–8), episode length (30/45/60s), and language.
2. Writes the script and song lyrics and pauses for your approval.
3. Generates character turnaround sheets, a scene sheet, and a scale sheet, then a keyframe for every story beat.
4. Records the voiceover line by line plus the children's song, and times each clip to the measured audio.
5. Adds subtitles, assembles the timeline, and renders the finished 9:16 episode.

## Wasteland Short Film

`/wasteland-short-film` is a post-apocalyptic 3D-CG episode maker with a dark, game-cinematic look. Give it a premise and it delivers a style-locked wasteland short — batch keyframes, image-to-video clips, sparse voiceover, ambient score, and an optional unifying color grade.

**Example prompt:** `/wasteland-short-film A lone scavenger finds a still-working music box in a dead city, 60 seconds, quiet desolation`

### How it runs:

1. Takes your premise (attach character illustrations to continue a series) and asks a brief: length (30/60/90s), cast size (1–3), and mood (Tense / Quiet desolation / Hard action).
2. Writes the creative plan and a Style Bible, and pauses for your approval.

3. Designs the fixed cast, then batch-generates every first-frame image — with a second gate to confirm the frames before any video is made.
4. Animates the frames into clips, optionally applies a unifying color grade.
5. Adds sparse voiceover (0–4 lines max), subtitles, an ambient score, and renders the final 16:9 film.

## E-commerce & Product

### Ecommerce Design

`/ecommerce-design` is a professional e-commerce image designer covering 14 retail verticals — apparel to electronics to food. It crafts product-first, compliance-safe, platform-ready listing images one careful shot at a time, scoring each result and inspecting it for defects before delivering.

**Example prompt:** `/ecommerce-design Hero image for a matte-black ceramic pour-over coffee set, Amazon main image plus one lifestyle scene`

#### How it runs:

1. Checks the required facts and asks only for what's missing: category, the product's exact color/material/finish, target platform, how many images and their roles, with or without a model.
2. Builds an attribute card from your uploaded product reference and identifies the right category playbook.
3. Chooses the shot types and locks cross-image consistency (palette, lighting, product identity).
4. Generates each image from a precisely assembled prompt matched to platform specs.
5. Scores every image on six weighted dimensions and inspects for defects, regenerating a failed shot once.

### E-commerce Image

`/ecommerce-image` turns one real product photo into a complete, platform-compliant listing set — hero, white-background, lifestyle, detail macro, scale, and variant shots — with the original photo referenced on every generation so the product never drifts.

**Example prompt:** `/ecommerce-image Turn this into a full Amazon listing set, editorial premium styling (attach your product photo)`

#### How it runs:

1. Takes your product photo and analyzes it — materials, proportions, brand cues.
2. Asks a short brief: platform (Amazon / Shopify / TikTok Shop / Etsy / Instagram), scope (Hero only / Core 5 / Full 7), and styling direction (Minimal Clean / Editorial Premium / Warm Domestic / Street Fashion).
3. Proposes a shoot plan and pauses for your approval.
4. Generates the set, passing your real photo as a reference on every single image.
5. Runs a 19-point quality check, regenerates failures once, and delivers the set on the canvas.

### Detail Page Designer

`/detail-page-skill-course` designs a long-form e-commerce detail page for courses and knowledge products, structured as a 6-layer persuasion funnel — hook, pain, solution, proof, offer, CTA — rendered as a coherent series of portrait page images.

**Example prompt:** /detail-page-skill-course A detail page for my 6-week watercolor course for total beginners, \$49, includes 200+ student testimonials

#### How it runs:

1. Asks a brief: offer type, audience, price band, and what proof assets you actually have (numbers, testimonials, credentials, none).
2. Collects your course details and optional logo or instructor photo.
3. Plans the 6-layer funnel (merging thin layers when needed) and pauses for your approval.
4. Generates one 3:4 page image per layer with a shared palette and type system.
5. Checks legibility, consistency, and flow, regenerating weak pages once.

**Tip:** It will never invent proof numbers for you — if you have real stats, include them in the prompt.

## Advertising

### Ad Idea

/ad-idea is an advertising creative director for the concept stage: it digs out the audience insight, generates three genuinely different ad concepts, develops the one you pick into a full proposal, and ends with mood-board images. Concepts only — hand off to /promo-video or /surreal-spot to shoot.

**Example prompt:** /ad-idea A campaign for a sleep-tracking ring aimed at burned-out young professionals

#### How it runs:

1. Collects the brief: product, audience, goal (awareness / launch / reposition / conversion), and tone.
2. Writes the insight — the "audience wants X but Y because Z" tension the campaign will lever.
3. Presents three concepts built with three different ideation methods, internally scored — you pick one.
4. Develops the chosen concept into a full proposal.
5. Generates 1–3 mood-board images and stops.

### Promo Video

/promo-video turns a product photo and a short brief into a publish-ready promotional video — retention-structured script, product-locked keyframes, animated scenes, music, and a final render in 16:9, 9:16, or 1:1.

**Example prompt:** /promo-video A 20-second launch video for our smart water bottle, 9:16, aimed at gym-goers (*attach the product photo*)

#### How it runs:

1. Analyzes your product photo and description into a Product Brief with a locked color palette.
2. Asks the brief: goal, duration (~15s/3 scenes, ~20s/4, ~30s/5), and aspect ratio.
3. Writes a Hook → Feature → CTA retention script and pauses for approval.
4. Generates a keyframe per scene — your real product photo referenced on each — then animates every keyframe into a clip.
5. Adds music and renders the final cut.

## Surreal Spot

`/surreal-spot` makes TVC-grade surreal concept videos — "Apple x David Lynch" — built on impossible-juxtaposition aesthetics and lo-fi film mediums (VHS, 16mm, Super 8, Polaroid...). The product reveal is delayed, faces never shown.

**Example prompt:** `/surreal-spot A dreamlike spot for our new perfume, 9:16, VHS mood (attach the product photo)`

### How it runs:

1. Takes your product photo, analyzes the brand's visual world, and asks for the aspect ratio (9:16 / 4:5 / 16:9).
2. Presents three creative concepts — each with its own style-reference image — and you pick one.
3. Generates a 3x3 nine-panel storyboard grid in the chosen concept, quality-gated.
4. Animates the storyboard into the motion video.
5. Adds optional music and renders the final spot.

## Dynamic Poster

`/dynamic-poster` designs a motion poster: a single, beautifully composed poster frame where exactly one impossible thing moves — the crema pours upward, the neon bleeds, the mountain breathes — while the type stays crisp and static. Perfect for launches and announcements.

**Example prompt:** `/dynamic-poster A launch poster for our espresso machine where the crema pours upward into the cup`

### How it runs:

1. Reads your subject (plus optional product/logo image) and asks the aspect ratio (9:16 / 4:5 / 16:9).
2. Offers three motion concepts, each on a different juxtaposition mechanic — you pick one.
3. Generates the poster at rest as a style reference, then a 9-panel storyboard of the motion.
4. Animates it with a locked-off camera and the single impossible motion.
5. Adds music and renders the final motion poster.

## HiLo Promo

`/hilo-promo` builds a vertical 9:16 social promo fronted by a talking, lip-synced AI presenter — portrait, spoken script, voiceover, lip-sync video, supporting visuals, synced subtitles, and music, rendered at 1080x1920.

**Example prompt:** `/hilo-promo A 30-second vertical promo for my nail salon with an energetic female presenter, clean studio style`

### How it runs:

1. Asks the brief: what you're promoting, visual style (Clean studio / Lifestyle / Bold neon / Minimal luxury), presenter (generated, or a look-alike from your photo), and voice.
2. Generates the presenter portrait and pauses for your approval.
3. Writes a 25–40 second Hook → benefits → CTA script — you approve it before recording.

4. Generates the voiceover, then the lip-synced talking-presenter video (the slow step), plus 2–3 supporting visuals.
5. Assembles the timeline with transcribed, synced subtitles and music, and renders after your final go-ahead.

## YouTube Creator

`/youtube-creator` plans a YouTube video end to end: a retention-optimized script built on hook psychology and pattern interrupts, plus 2–3 high-CTR thumbnail images. It deliberately stops at the thumbnail — planning, not production.

**Example prompt:** `/youtube-creator A 10-minute video about why cities are quietly banning cars, for my urbanism channel`

### How it runs:

1. Takes your topic and channel/audience info, then asks: video type (educational / tutorial / listicle / story), duration, and format (16:9 or 9:16 Shorts).
2. Presents four title-and-angle candidates on different psychological levers — you pick one.
3. Writes the segmented script with retention tags: the three-part hook, pattern interrupts every 60–90 seconds, and CTAs.
4. Generates 2–3 thumbnail concepts using proven conversion patterns and squint-tests them.
5. You pick the thumbnail that ships.

## Music Video

### MV Creator

`/mv-creator` is a complete music-video production pipeline: from lyrics (pasted text, an LRC file, or an audio upload) through story, art style, per-scene script, character design, backgrounds, storyboard, and per-scene clips to a finished, subtitled MV on the timeline.

**Example prompt:** `/mv-creator Make a music video for these lyrics – melancholic indie folk, autumn palette (paste your timestamped lyrics or attach the song)`

### How it runs:

1. Breaks the lyrics into timestamped segments and validates the structure.
2. Proposes three story concepts — you pick one — then four art styles (plus "other"), and locks the style with a reference image.
3. Writes the per-scene script and designs the character portraits and background hero shots that anchor every scene.
4. Builds the storyboard, then animates each scene into an image-to-video clip with the characters referenced to prevent drift.
5. Assembles the timeline with your track and lyric subtitles, and renders the final MV.

**Tip:** Timestamped lyrics ( `0:00 – 0:08` per segment) give the cleanest scene timing — the skill will show you the exact format it wants.

## Sword Dance

`/sword-dance` creates a wuxia sword-dance music video in an ink-wash-meets-photoreal Chinese aesthetic: one identity-locked hanfu dancer, misty establishing shots, a named choreography, per-phrase clips, and a guzheng-and-drums score.

**Example prompt:** `/sword-dance A melancholic sword dance on a misty mountain ridge at dawn, 45 seconds`

### How it runs:

1. Takes your concept (and optional likeness photo, used as a reference — not a face swap) and asks: mood (Graceful / Fierce / Melancholic), setting, and length (30/45/60s).
2. Designs the dancer and pauses until you lock the character.
3. Generates two establishing shots, then writes the choreography as named movement phrases — approved by you.
4. Generates a keyframe per beat referencing the locked design, then one clip per phrase.
5. Composes the guzheng-and-drums score, assembles, and renders the 16:9 MV.

## Beat Sync Editor

`/beat-sync-editor` cuts a montage where every single cut lands exactly on a beat. It detects the track's beat grid, plans slots with shot lengths that follow the music's energy, fills gaps with generated stills or animations, and renders the cut.

**Example prompt:** `/beat-sync-editor Cut my canvas clips to this track – fast cuts on the chorus, breathe on the verses (attach the music)`

### How it runs:

1. Establishes the music: your attached track, one from the canvas, or a freshly generated one (long songs can be trimmed to a segment).
2. Analyzes the beat grid — BPM and exact beat times.
3. Proposes a cut plan with slots on the beats and shot lengths varied by energy — you approve, or ask for faster/slower.
4. Inventories your clips, trims them to slot length, and generates fills for any gaps.
5. Assembles the beat-locked timeline and renders after your confirmation.

## Documentary & Explainer

### Documentary

`/documentary` directs a narrated documentary or explainer from a single topic sentence: researched three-act outline, style-locked visuals, professional narration that drives the edit, subtitles, music, and a final render — 1 to 3 minutes, 16:9 by default.

**Example prompt:** `/documentary A 2-minute documentary about the last lighthouse keepers of the Atlantic, warm tone`

### How it runs:

1. Takes your topic and asks: tone (Authoritative / Warm / Investigative / Poetic), length (1/2/3 min), narration language, and look (Cinematic / Archival / Nature-doc / Illustrated).
2. Researches the subject and writes a three-act outline with the exact narration lines — you approve before production.
3. Generates style-locked keyframes for every beat, plus intro and outro title cards.
4. Records the narration line by line and times each clip to the measured audio, so the voice drives the cut.
5. Adds subtitles, an optional color pass, and music, then assembles and renders the film.

## Visual Production

### Avatar & Figure Studio

`/avatar-studio` turns a photo or a description into stylized identity keepsakes: avatar sets, the viral boxed action-figure shot, sticker/PFP packs, character turnaround sheets, or a drag-to-spin faux-3D turntable — with your face kept consistent across every output.

**Example prompt:** `/avatar-studio` Make me a boxed action figure called "The Deadline Slayer" with a laptop and triple espresso accessories (*attach a selfie*)

#### How it runs:

1. Works out the mode — avatar, action figure, sticker pack, character sheet, or 3D turntable — and whether identity comes from your photo or a description.
2. Asks for a style pick, and for the action figure, the name and title on the box.
3. Generates the set for the chosen mode: hero avatar plus variations, the templated figure box, ~6 stickers, turnaround views, or 6–12 rotation angles.
4. For the turntable, wraps the angles into an interactive drag-to-spin card on the canvas.
5. Confirms with you before any large batch spends serious credits.

### Storyboard

`/storyboard` breaks a story or script into a professional shot list — varied shot scales, camera notes, per-shot prompts — then generates one consistent storyboard frame per shot, all locked to the same style and characters.

**Example prompt:** `/storyboard` Storyboard this scene: a museum heist goes wrong when the power comes back on early

#### How it runs:

1. Takes your story (attach style or character reference images if you have them) and sets the style anchors.
2. Writes the shot breakdown table — 6 to 16 shots with scales from long shot to extreme close-up, never repeating a scale on adjacent shots — and saves it as a document.
3. Pauses: "Generate all frames" or "Revise the breakdown".
4. Generates every frame with the locked references.
5. Runs an automatic quality pass and retries any flagged frame once.

## Poster Design

`/poster-design` produces a static poster with real typographic hierarchy and clean, verbatim text rendering, across 12 poster subtypes — concert, film, product, event, minimalist and more — then scores its own result and retries once if it's below the bar.

**Example prompt:** `/poster-design A concert poster for a synthwave night called NEON DRIFT, June 21, 9:16, retro chrome type`

### How it runs:

1. Asks a two-tab brief — poster subtype (12 options) and aspect ratio — skipping anything you already specified.
2. Asks for the exact display text: title, subtitle, date, CTA (or "no text").
3. Writes a design direction: concept, composition mode, type hierarchy, and a hex palette.
4. Generates the poster with your text rendered verbatim.
5. Scores it 0–100 and regenerates once if it lands under 75 or any text came out garbled.

## Anime Design

`/anime-design` is a professional anime character designer with an identity lock: it generates an anchor portrait, then keeps the same recognizable person across poses, expressions, scenes, and an optional multi-view turnaround sheet. 14 substyles from 90s cel to modern webtoon.

**Example prompt:** `/anime-design A 90s-cel-style bounty hunter with silver hair, a long coat, and a scar over one eye`

### How it runs:

1. Asks you to pick one of 14 substyles (skipped if you already named one).
2. Collects the character description (plus optional reference image) and writes a Lock Card — the character's binding identity spec.
3. Generates the anchor portrait, verifies it, and retries once if it misses.
4. Generates the requested set — new poses, expressions, scenes — always referencing the anchor so nothing drifts.
5. Offers a multi-view turnaround sheet to finish.

## Anime Style Forge

`/anime-style-forge` handles anime stylization with provenance: it identifies exactly which school and era a target style comes from (Japanese, Chinese, Korean, Western families), writes it into a style spec you approve, then generates from text or converts your attached photo.

**Example prompt:** `/anime-style-forge Convert this portrait into a soft storybook anime with watercolor edges (attach the photo)`

### How it runs:

1. Works out the mode — generate from text, convert your image, or identify a style you attached.
2. Analyzes the target style across ten dimensions, collapses them into Form / Aesthetic / Mood axes, and names the "soul" of the style.

3. Shows you the Style Spec: approve, adjust, or pick a different style.
4. Generates with the soul axis weighted hardest.
5. Checks consistency against the spec and retries once if it drifted.

## Film Assets

`/film-assets` is a pre-production asset factory for film and TV: identity-locked character reference sheets, prop multi-view orbit sheets, and scene concept art, generated in internally consistent batches.

**Example prompt:** `/film-assets A character reference sheet for a weathered space-salvage captain – 5 views, neutral background`

### How it runs:

1. Asks which module you need: Character, Prop, or Scene (plus the character sub-type).
2. Takes an optional reference photo and a config brief — view counts, backgrounds, expression sets, pose counts.
3. Builds an Identity Card from the reference so every sheet shows the same subject.
4. Generates the module's sheets with silent quality gates and one retry on misses.
5. Presents the batch for review and offers extra rounds.

## Film Shot

`/film-shot` designs cinematic stills with real camera language: each shot is assembled from six dimensions of shot grammar across eight calibrated visual styles, with the subject locked so a 3–6 shot series reads as one film.

**Example prompt:** `/film-shot A lone astronaut at a roadside diner counter, melancholic, 3 shots`

### How it runs:

1. Takes an optional subject reference, then asks the brief: medium (Realistic / Anime / 3D), style (8 options), emotion (6 options), and scope (single shot, 3, or 5–6).
2. Assembles each shot from the six shot-language dimensions into a shot plan document.
3. Generates the shots — widescreen 2.39:1 or 16:9 — with the subject identity locked verbatim in every prompt.
4. Runs a drift check and retries failing shots once.
5. Offers an optional character card or three-view sheet of your subject.

## Character & Scene Sheet

`/character-scene-storyboard` is the one-pass pre-production package: from a story plus up to four character photos it delivers three composite images — a character design sheet, a scene concept board, and a storyboard panel grid — all in one shared style.

**Example prompt:** `/character-scene-storyboard Design sheets and a 12-panel storyboard for this story (attach the story text and 1–4 character photos)`

### How it runs:

1. Takes your story and character references, extracts appearance cards and a ~12-beat shot list, and pauses for your approval.
2. Asks the brief: style (6 options), panel aspect ratio (16:9 / 9:16 / 1:1), and panel count (6 / 12 / 16).
3. Generates the three composite sheets in one pass, characters anchored to your photos.
4. Quality-checks each sheet and retries once where needed.
5. Delivers the package and offers an iteration round.

## Face Warp

`/face-warp` deconstructs one portrait into two artistic variants: a "faceless" portrait where the features are erased but everything else — hair, light, wardrobe — stays intact, and a "feature puzzle" collage where the eyes, brows, nose, and lips are arranged as separate pieces.

**Example prompt:** `/face-warp Deconstruct this portrait` (*attach a clear, front-facing photo*)

### How it runs:

1. Takes your portrait (it will ask for a clearer one if the face is obscured).
2. Analyzes the features into a written character profile.
3. Generates Variant A — the faceless portrait — quality-checked with one retry.
4. Generates Variant B — the feature-puzzle collage — same check.
5. Offers a side-by-side composite of both.

## Film Style Picker

`/film-style-picker` helps you lock a target cinematic look from a built-in database of 36 styles across 12 categories — then hands you a reusable style guide: prompt template, palette, and usage notes you can apply in any other skill.

**Example prompt:** `/film-style-picker Something dreamy for a rainy rooftop love scene`

### How it runs:

1. Asks which of four style groups fits (skipped if you named a director, film, or style), optionally narrowing by mood words.
2. Presents the top three candidate styles as cards — with optional generated sample frames if you want to spend the credits.
3. You lock one in (or ask for different candidates).
4. Delivers the style-guide document: prompt template, palette, and notes.
5. Optionally generates a demo image of *your* scene in the locked style.

## Audio & Voice

### ASMR Ambient

`/asmr-ambient` produces layered sleep and relaxation audio — guided meditations, bedtime stories, or ambient ASMR — with a whisper-grade voiceover and a calm instrumental bed, laid out on the timeline for you to mix.

**Example prompt:** /asmr-ambient A 5-minute guided rain meditation for falling asleep, soft female whisper

#### How it runs:

1. Asks the session brief: type (Guided Meditation / Bedtime Story / Ambient ASMR), duration (~3/5/10 min), voice, and ambience (rain, fireplace, ocean, forest, drone, none).
2. Takes your theme and writes the script by sleep-writing rules — energy only ever goes down — then pauses for approval.
3. Records the voiceover in gentle 60–90 word segments.
4. Generates the matching instrumental background bed.
5. Opens the timeline with Dialogue and Music tracks so you can balance and render the mix yourself.

### Narrate Scene

/narrate-scene adds a voiceover and burned-in synced subtitles to one existing video clip and renders the finished narrated MP4. The fastest way to turn a silent clip into a shareable one.

**Example prompt:** /narrate-scene Narrate this clip like a nature documentary, calm male voice (*attach the clip*)

#### How it runs:

1. Finds the target video — your attachment or a clip on the canvas.
2. Writes narration sized to the clip length (~2–3 words per second) and shows it to you to edit or approve — or uses the text you supplied.
3. Picks a voice with you and generates the voiceover.
4. Fits the picture to the narration (looping if the clip is shorter — no black tail).
5. Adds word-synced subtitles via auto-transcription (or evenly placed captions as a fallback) and renders the final clip.

### Drama Soundtrack

/drama-soundtrack composes a full original-soundtrack suite for your script or story: an opening theme, an ending theme, one theme per main character, and scene BGM cues — every track justified by a musical analysis you approve first.

**Example prompt:** /drama-soundtrack Score my 3-episode noir — a detective, a femme fatale, rain-soaked Tokyo (*paste the script or story*)

#### How it runs:

1. Reads your script or show description — genre, setting, main characters, emotional beats.
2. Writes the musical analysis: the emotional arc, each character's musical personality, and a track-by-track Soundtrack Plan.
3. Pauses for approval before generating anything (music is the expensive part).
4. Generates every track — themes and instrumental-only BGM cues.
5. Lays the suite on the timeline for audition; you decide what to keep.

## Audiobook

`/audiobook` turns a prose chapter into multi-voice narrated audio: it segments the text, casts a distinct voice for the narrator and every character — one character, one voice, forever — and records the chapter in reading order onto the timeline.

**Example prompt:** `/audiobook Record this chapter – a gruff dwarf blacksmith, a young elf apprentice, and a wry narrator` *(paste the chapter text)*

### How it runs:

1. Reads your chapter and splits it into ordered narrator and character speech segments.
2. Proposes the casting sheet — one voice per role, matched on gender and age — and pauses for your approval (you can recast any voice).
3. Records the segments verbatim in reading order, with a listen-check after each character's first line.
4. Pauses to report progress on long chapters (60+ segments).
5. Opens the timeline with the chapter assembled on the Dialogue track.

## Voice Clone

`/voice-clone` clones a voice from a reference recording and gives you a reusable voice you can cast in any other skill — with a mandatory consent gate first, and a spoken demo so you can verify the result.

**Example prompt:** `/voice-clone Clone this voice and name it "Grandpa Jozef"` *(attach a 1–5 minute clean recording)*

### How it runs:

1. Starts with the consent gate: you must confirm you have the legal right to clone this voice and will use it lawfully. Celebrity and unowned voices are refused.
2. Locates and quality-checks the reference audio (mp3/wav/m4a/aac/ogg/flac, up to 10 MB; overlong clips are trimmed).
3. Clones the voice.
4. Verifies it with a spoken demo line in your language.
5. Saves a voice registry card on the canvas so the voice is easy to reuse later.

**Tip:** Voice cloning is a paid-tier voice-provider feature — the skill will tell you plainly if your current provider plan doesn't include it.

## Games & Interactive

### Games

`/games` builds real, playable games — arcade, puzzle, quiz, platformer — that run instantly in a window on your canvas and export as a single self-contained HTML file you can send to anyone. Keyboard and touch controls, HUD, win/lose, and restart included.

**Example prompt:** `/games A retro arcade snake game with power-ups and increasing speed`

### How it runs:

1. Builds straight away if you named a game; otherwise asks one genre question (Arcade / Puzzle / Quiz / Platformer) and whether you want a retro code-drawn look or AI-generated art.
2. Picks the right engine for the job — from a simple 2D canvas up to physics and 3D engines.
3. Optionally generates sprites, backgrounds, and audio (it asks before spending credits on art).
4. Builds the game — complex ones as clean modular systems — and runs an automated test pass to make sure it actually boots and plays.
5. Iterates with you ("make the enemies faster", "add a boss") and exports the standalone HTML when you're happy.

## 3D Game Studio

`/threejs-game` is the 3D-specialized sibling of Games: premium Three.js browser games — endless runners, racers, shooters, physics toys — built in disciplined phases with a mandatory generated-asset pass (skyboxes, textures, sprites, sound effects, music) and an AAA graphics pass (lighting, materials, bloom).

**Example prompt:** `/threejs-game An endless runner through a neon cyberpunk canyon with a speed boost mechanic`

### How it runs:

1. Locks the game loop with you — genre, art direction, target device — asking at most one clarifying question.
2. Builds the playable core slice (world, player, enemies) and boot-tests it.
3. Runs the full asset pass: ~8–14 generations covering skybox, tileable textures, billboards and sprites, SFX, and a music bed.
4. Applies the AAA graphics pass — lighting, materials, tone mapping, shadows, bloom — then UI/HUD and game-feel tuning.
5. Scores itself against a quality rubric, fixes what fails, and exports the standalone HTML.

**Tip:** 3D models are built from procedural geometry — the skill is honest that it doesn't generate mesh files. Its generated images become textures, skies, and sprites.

## Platform Tools

### App Builder

`/app-builder` is a no-code app designer: describe an app and it plans the screens, declares the data models with realistic seed data, builds every screen with live phone-frame preview on the canvas, wires the navigation, and exports runnable React (web) and React Native (Expo) code. It also does LinkedIn/Meta ad creatives and desktop layouts.

**Example prompt:** `/app-builder A habit tracker app with streaks, a today checklist, and a stats screen – calm emerald theme`

### How it runs:

1. Clarifies the format if needed (phone app / desktop web / ad creative) and the essentials: purpose, screens, style.

2. Plans the screens, sets the theme, and declares the data models with realistic seed rows.
3. Builds each screen — lists bound to real data collections, forms that capture input, working navigation between screens.
4. Optionally wires app state, auth flows, and conditional logic.
5. Reviews its own work with a screenshot critique pass, then exports the runnable code.

## App Walkthrough

`/app-walkthrough` records a narrated walkthrough video of an app you built with App Builder: it screenshots every screen, tours them in navigation order with an animated tapping cursor and transitions, and narrates each screen with a generated voiceover.

**Example prompt:** `/app-walkthrough` Record a walkthrough of my app with a warm, friendly voice

### How it runs:

1. Checks your project actually has an App Builder app with screens.
2. Optionally asks one question about the narration style (warm / crisp professional / captions only).
3. Records the whole tour in one pass — screenshots, cursor animation, transitions, one voiceover line per screen.
4. Drops the finished MP4 on your canvas and reports the duration and screen count.

## Documents

`/documents` reads and writes office files: extract the text from a PDF, Word, Excel, or CSV file to summarize or analyze it, and generate real .docx, .pdf, and .xlsx deliverables that land as document cards on the canvas.

**Example prompt:** `/documents` Summarize this contract and give me a one-page Word summary of the key obligations (*attach the PDF*)

### How it runs:

1. Extracts the content of your attached file (PDF / DOCX / XLSX / CSV) onto the canvas.
2. Does the thinking you asked for — summary, analysis, data extraction.
3. Generates the deliverable: a Word or PDF document for prose, a real spreadsheet for tabular data.
4. Drops the finished file as a card you can open or share.

**Tip:** Scanned image-only PDFs can't be read yet (no OCR), and PowerPoint generation lives in `/automation` — Documents will point you there and offer a PDF deck instead.

## Automation

`/automation` runs Python on your own computer to automate files, spreadsheets, documents, and data — and can even drive your screen. Every code block appears on an approval card with a one-sentence explanation before it runs; nothing executes without your click.

**Example prompt:** /automation Rename all the invoice PDFs in my Downloads folder to "YYYY-MM-DD\_total.pdf" and build me a summary spreadsheet

**How it runs:**

1. Tells you what it's about to do, then proposes a small code block — you approve or deny each one.
2. Reads the result and decides the next step; work builds up incrementally.
3. Uses the right library for the job — Excel, Word, PowerPoint, PDF — so outputs are real office files.
4. For desktop control, works in a careful look → act → verify loop with screenshots.
5. Saves everything it produces into your project's assets and attaches it to the canvas.

**Tip:** Enable Automation in Settings first; on macOS, desktop control additionally needs Screen Recording and Accessibility permissions.

## Skill Creator

/skill-creator turns a workflow you liked — or an idea you describe — into a brand-new installable skill. It drafts the complete skill definition as a document on your canvas and walks you through installing it from the Hub.

**Example prompt:** /skill-creator Turn what we just did into a reusable skill called "podcast-teaser"

**How it runs:**

1. Asks whether the new skill should be reconstructed from this session's workflow or built from an idea (with a short brief for the latter).
2. Confirms the reconstructed workflow with you.
3. Drafts the full skill definition — staged flow, questions, generations — as a document card.
4. Runs a review round: install as-is, revise, or add a stage.
5. Hands you the install steps: Hub → Install Skill → paste the markdown.

## Quick Tools

One-shot skills: no briefs, no gates — one prompt, one result. Perfect when you know exactly what you want.

### Quick Image

/image-creator generates one polished image from a single prompt. If your request is a rough sketch, it silently expands it into a well-composed prompt — subject, action, environment, lighting, camera — and generates once.

**Example prompt:** /image-creator a red fox curled in fresh snow at golden hour, low camera angle

**How it runs:**

1. Reads your request; uses it verbatim if concrete, otherwise fleshes it out.
2. Generates the image and names the card.

3. Done — use the card's Remix button for variations.

## Quick Video

`/video-generator` generates one short cinematic clip from a single prompt — text-to-video, or image-to-video if you attach a still to use as the first frame.

**Example prompt:** `/video-generator a paper boat drifting down a rain-soaked gutter stream, slow tracking shot`

### How it runs:

1. Shapes your request into a motion brief: what's in frame, what it does, how the camera moves.
2. Animates from your attached still if you provided one (better continuity), else generates from text.
3. Delivers one clip, roughly six seconds or less.

## Quick Audio

`/audio-studio` generates one audio asset — music, background score, a sound effect, or a spoken voiceover line — from a single prompt.

**Example prompt:** `/audio-studio lo-fi study beat with vinyl crackle and soft piano, 60 seconds`

### How it runs:

1. Classifies your request: music/BGM or spoken voice.
2. Music gets a one-line genre/instrumentation/mood brief; voiceover speaks your written line verbatim.
3. Generates the asset and stops.

## N-Grid Storyboard

`/n-storyboard` tells a whole visual story as ONE composite image: an N×N grid of sequential panels with continuous characters, lighting, and style — an instant story-at-a-glance.

**Example prompt:** `/n-storyboard A 3x3 grid of a small robot learning to garden, wide establishing first panel`

### How it runs:

1. Breaks your story into the grid's key beats — first panel establishes, last one pays off.
2. Locks one style profile across all panels.
3. Generates the entire grid in a single image (2×2 up to 4×4).

## Image Remix

`/image-remix` extracts what makes an inspiration image *work* — composition, palette, lighting, mood, never the subject — and generates one new image that carries the same vibe onto a different subject.

**Example prompt:** `/image-remix Same vibe, but a lighthouse at dusk instead (attach the inspiration image)`

### How it runs:

1. Analyzes the inspiration and names its "soul" — the one or two dimensions that carry the feel.
2. Rebuilds those qualities purely as a text description (your inspiration is never copied directly).
3. Generates the new subject wearing the old vibe.

## Minecraft Pixel Art

`/minecraft-pixel-art` converts a photo or a described subject into Minecraft-style blocky voxel art — cubic geometry, 16×16 texture feel, and the classic block palette, with source colors mapped to named block materials.

**Example prompt:** `/minecraft-pixel-art Blockify my dog` (*attach the photo*)

### How it runs:

1. Takes your photo (locks the composition) or your described subject.
2. Rebuilds it from cubic blocks with flat shading and screenshot-style framing.
3. Delivers the image and tells you which block materials it mapped your colors to.

## Multi-Shot

`/multi-shot` re-renders a subject image from a different camera angle — identity, lighting, and style locked to the original. Seven angle presets, one angle per pass, loop as long as you like.

**Example prompt:** `/multi-shot Show her from a low-angle hero shot` (*attach the subject image*)

### How it runs:

1. Takes your subject image.
2. You pick an angle from 7 presets (front, three-quarter, profile, back, low-angle hero, top-down, dutch tilt) — skipped if you named one.
3. Re-renders the same subject from that viewpoint.
4. Asks if you want another angle and links the series on the canvas.

## Image Paint Edit

`/image-paint-edit` changes exactly one region or object of an existing image — replace an object, fix a detail, drop a product into a scene — while everything else stays pixel-identical.

**Example prompt:** `/image-paint-edit Replace the coffee mug on the desk with a small bonsai tree`

### How it runs:

1. Finds the source image (your attachment or a card on the canvas).
2. Pins down the edit: which region, what replaces it, what must not change.
3. Runs the edit — masked inpainting or a described-region edit, including multi-reference product placement.
4. Shows the result and offers one refinement round.

## Relight

`/relight` changes only the light in an existing image — direction, intensity, color temperature, shadows, mood — while the subject and composition stay identical. Fourteen named presets from Golden Hour to Neon Noir, all blendable.

**Example prompt:** `/relight Golden hour from the left, long warm shadows` (*attach or pick the image*)

### How it runs:

1. Identifies the source image and reads its current lighting.
2. Collects the lighting intent — a named preset, your description, or a short 4-tab brief (direction, temperature, mood, background).
3. Generates the relit image with a strict "change nothing but the light" clause.
4. Offers 1–3 variants (always relit from the original) and an optional color trim.

## Video Prompting

`/video-prompting` is a video prompt engineer: it turns your plain-language clip description into one production-grade generation prompt — subject, action, camera, lighting, style, duration — saved as a document you can reuse, with an optional "generate now".

**Example prompt:** `/video-prompting A drone shot rising over a foggy pine forest at dawn`

### How it runs:

1. Picks the mode — text-to-video, or image-to-video if you attached a reference frame.
2. Writes the prompt as one disciplined paragraph: visual verbs, one camera move, no AI-tell buzzwords.
3. Saves it as a document with mode, duration, and notes.
4. Asks: generate the clip now, or just keep the prompt.

## Camera Path Prompter

`/path-guided-camera-move` turns an image with a drawn arrow or path (or a described route) into an FPV follow-camera video prompt — one continuous shot along your route, with a mandatory speed profile and scene-native sound only.

**Example prompt:** `/path-guided-camera-move Fly the camera along the arrow I drew, ending at the tower door` (*attach the marked-up image*)

### How it runs:

1. Reads your image and traces the drawn route — start, waypoints, end (your markings never appear in the video).
2. Decides subject-follow vs pure camera flythrough.
3. Writes the structured prompt document: route segments, camera grammar, speed profile, diegetic audio.
4. Offers to generate — with defaults, an adjusted duration (4/6/8/15s), or prompt-only.

## Clip Export

`/clip-export` is the finishing editor: it inventories every clip and audio asset on your canvas, proposes a cut order, assembles the timeline, syncs subtitles if you have a voiceover, and renders the shareable MP4.

**Example prompt:** `/clip-export Assemble everything on the canvas into a final cut, tightest clips first`

### How it runs:

1. Inventories your canvas — videos, audio, subtitles — and stops honestly if there's nothing to cut.
2. Proposes the clip order (or lets you pick and exclude), with optional trimming of long takes.
3. Assembles the timeline and offers voiceover-synced subtitles.
4. Waits for your "render", then renders the MP4 with crossfades and points you at the app's export options (SRT file, project bundle).

## Skill Reviewer

`/skill-reviewer` is a read-only auditor for skill definitions: paste a SKILL.md and it scores it on five dimensions, quotes the exact problem lines, and hands you prioritized fixes — without changing or installing anything.

**Example prompt:** `/skill-reviewer Review this skill before I install it (paste the SKILL.md)`

### How it runs:

1. Reads your pasted skill definition.
2. Scores five dimensions 0–10: frontmatter, tool correctness, self-containment, flow design, safety & honesty.
3. Writes the review — Critical / Important / Nice-to-have findings, each quoting the offending line with a concrete fix.
4. Drops the review card on the canvas and offers to draft corrected sections.

## Voice Design

`/voice-design` creates a custom AI voice from a plain-language description — age, gender, tone, accent, pacing — and gives you a reusable voice plus a preview audio card, ready to cast in audiobooks, promos, or dramas.

**Example prompt:** `/voice-design A gravelly 60-year-old sea captain, slow and warm, slight Irish accent`

### How it runs:

1. Reads your description and asks at most one clarifying question if it's vague.
2. Designs the voice and generates a preview line.
3. Saves a voice registry card on the canvas so the voice is easy to reuse.
4. Offers a longer demo, a tweak-and-redo, or done.

**Tip:** Voice design is a paid-tier voice-provider feature and works best in English, Chinese, and Japanese — the skill warns you up front about other languages.

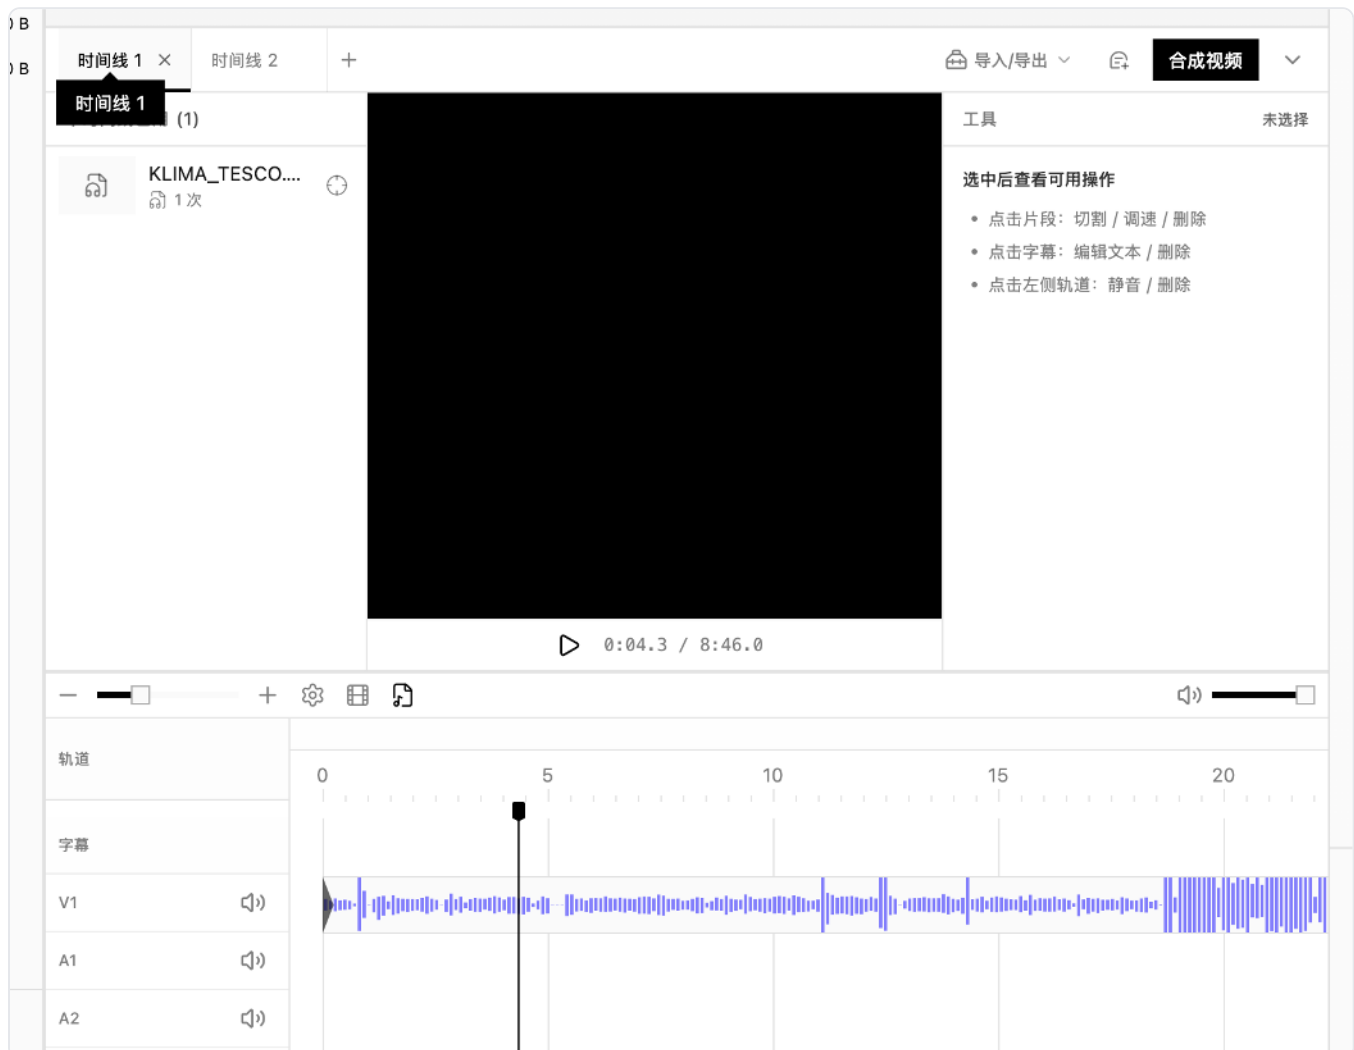
## Build your own

Every skill you just read is a markdown playbook the AI director follows — and you can write your own. Run [/skill-creator](#) after a session you liked ("turn what we just did into a skill") or describe a new workflow from scratch, review the draft it produces, and install it from the Hub in one paste. For how skills, the Hub, and installation fit together, see the [Skills & the Hub](#) chapter.

# Timeline Editor

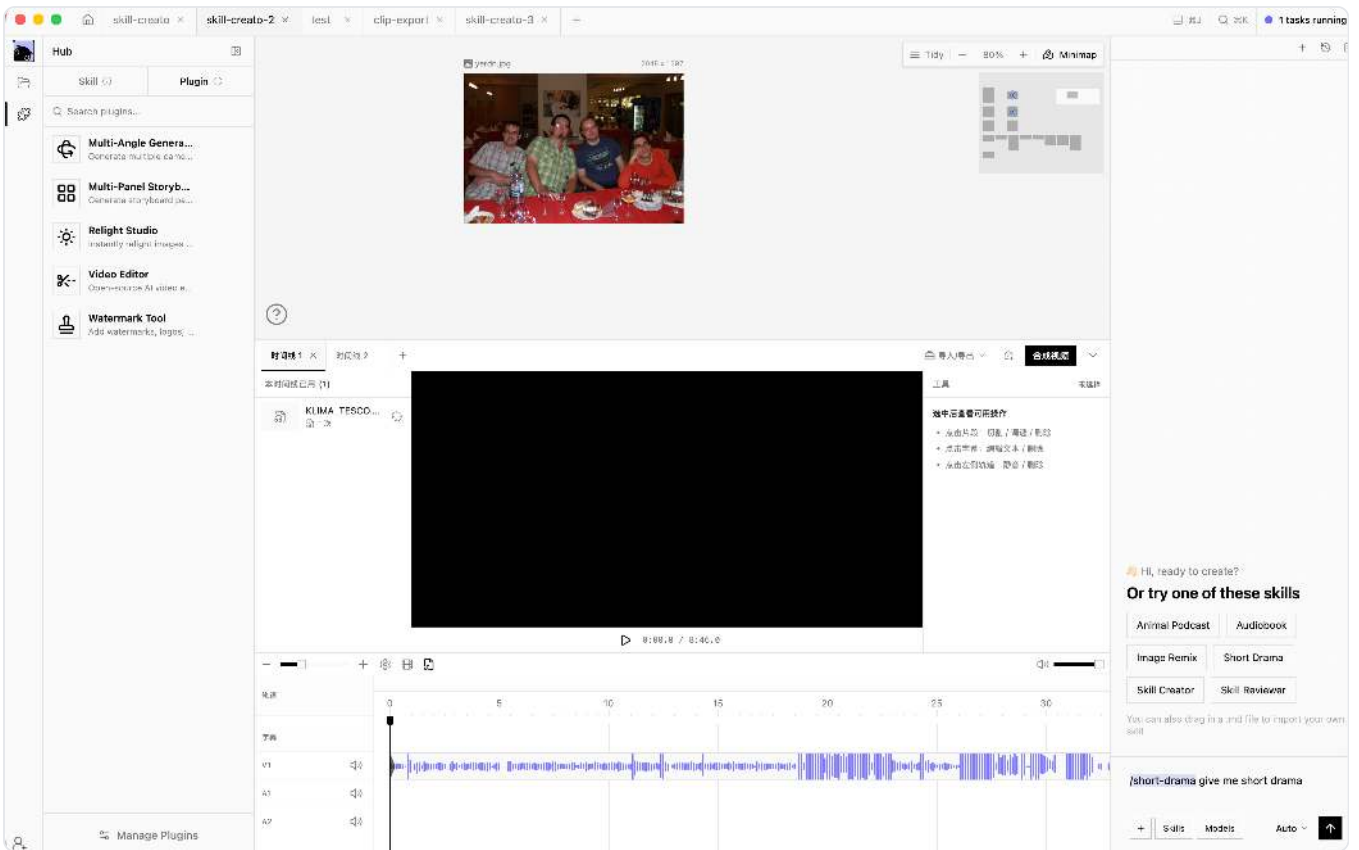
Cut, arrange and render your final video: tracks, transitions, subtitles, audio tools and export.

The Timeline Editor is where your generated clips become a finished film. Arrange video, dialogue, sound effects, music and subtitles across tracks, fine-tune the cuts, then compose everything into a real MP4 — all inside the app, no external editor required.

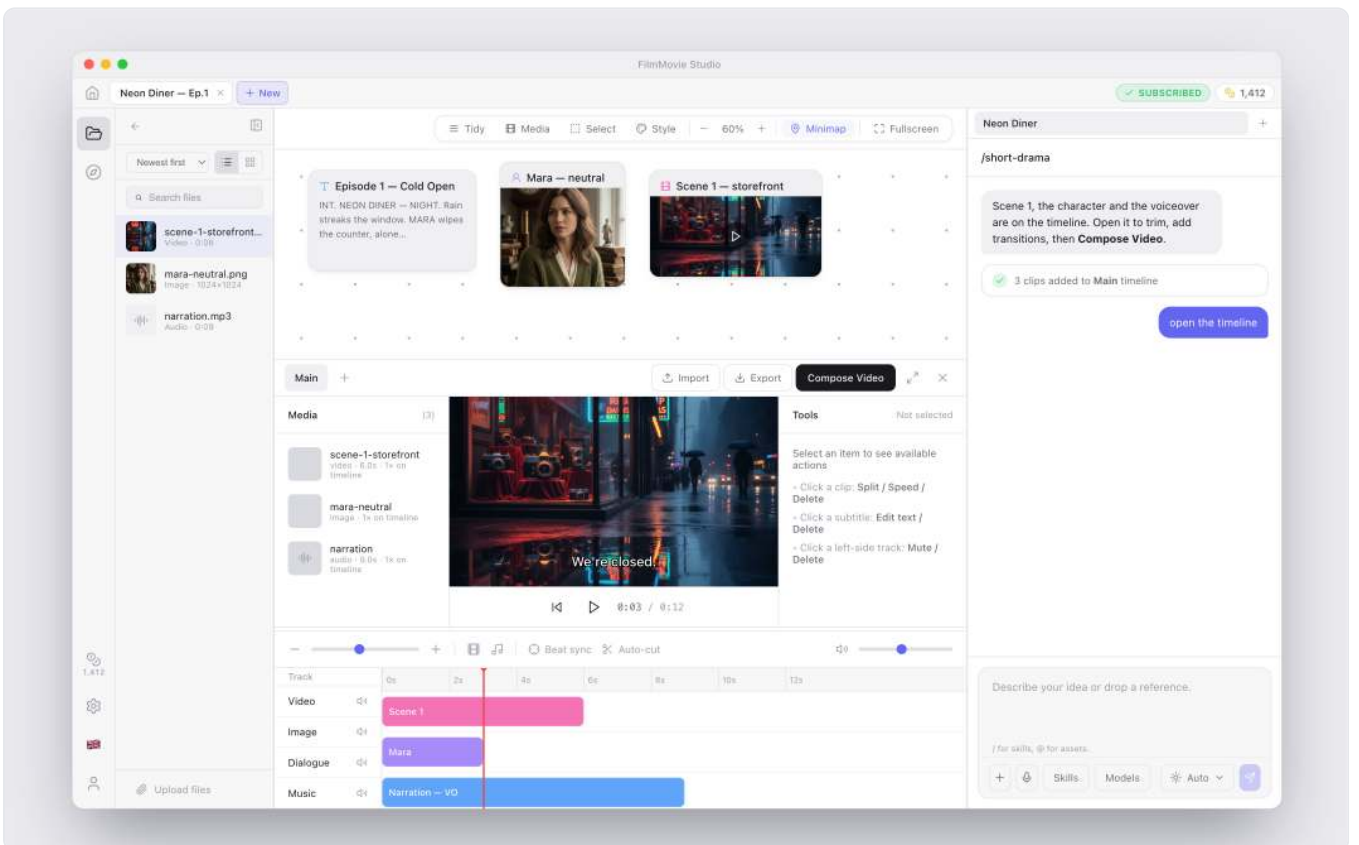


## Opening the timeline

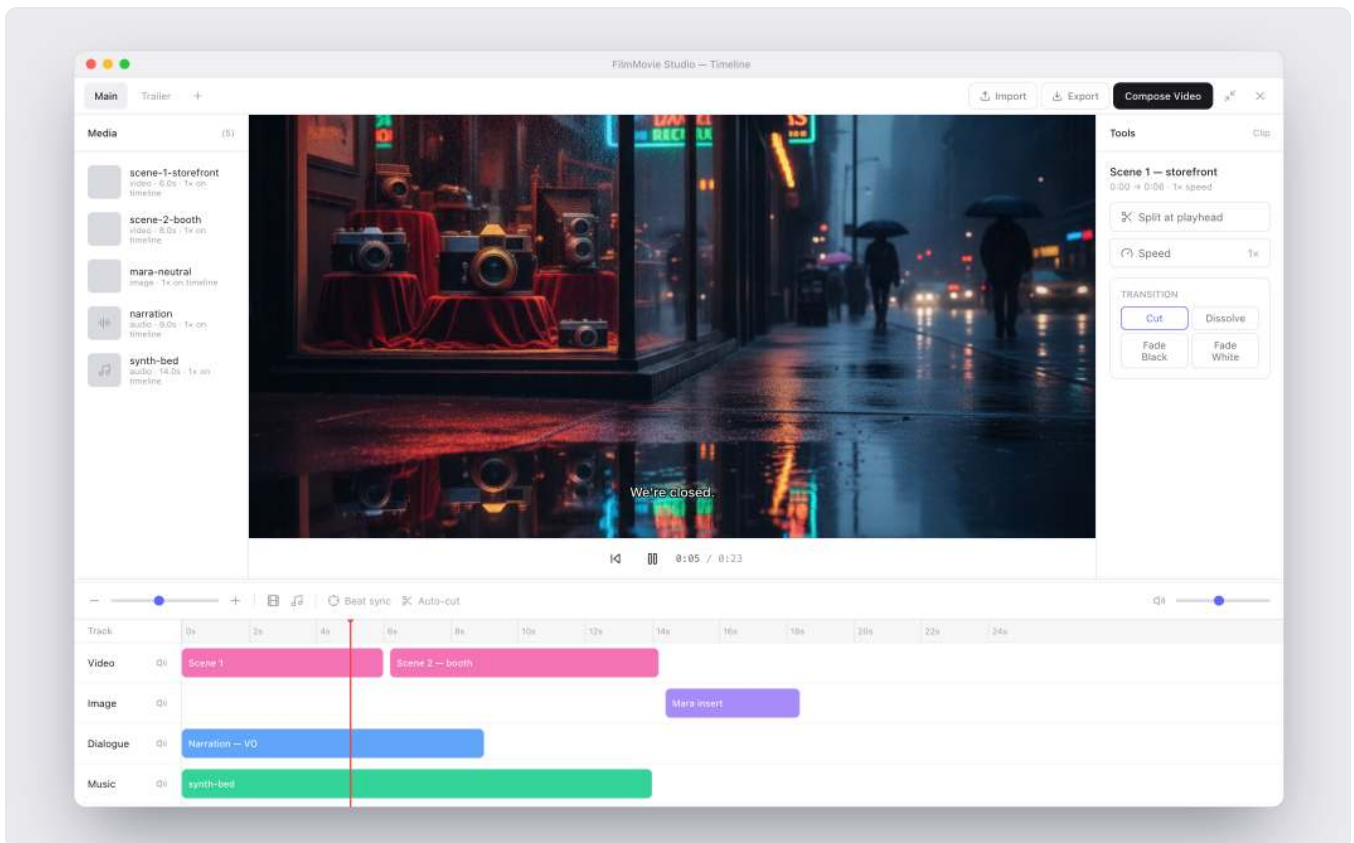
Every project workspace has a floating **Timeline** pill in the bottom-left corner of the canvas. The number in parentheses — for example **Timeline (12)** — shows how many clips are already on the active timeline, and the pill lights up in the accent color as soon as there is content ready to render.



Click the pill and the editor opens inside the workspace:



Need more room? Use the fullscreen button in the top-right of the editor (**Open timeline fullscreen**) — press **Esc** or click the button again to exit.



## Multiple timelines

One project can hold several cuts — a full film, a trailer, a vertical short. The tabs across the top of the editor are your timelines:

- Click **+** (**New timeline**) to add another timeline (they are named Timeline 1, Timeline 2, ...).
- Click a tab to switch to that timeline.
- Click the small **x** on a tab to delete it (available whenever you have more than one).

Each timeline keeps its own clips, subtitle style and playhead position, and everything is saved automatically.

## Adding content

There are three ways to get media onto the timeline:

1. **From a canvas card** — hover any video, image or audio card on the canvas and click **Add to timeline** in its toolbar. See [Canvas & Cards](#).
2. **From the Media panel** — the left column of the timeline editor lists every asset in the project (grouped by scene). Hover an item and click the crosshair icon (**Add to timeline**).
3. **Import** → **Media file...** — bring in any file from your computer; it becomes a project asset and lands on the timeline in one step.

Clips are routed to the right track automatically:

Content	Track
Video	V1
Image (stills)	V2
Voiceover / dialogue audio	Dialogue
Sound effects	SFX
Music	Music
Subtitles	Subtitle

New clips snap to the end of their track, and you can add extra video or audio tracks with the film and note icons in the track toolbar.

## Playback & view controls

Control	What it does
<b>Play / Pause</b> (▶ / ⏸)	Plays the timeline in the preview, with all unmuted audio tracks mixed live
<b>Jump to start</b> (⏮)	Moves the playhead back to 0:00
Time ruler	Click or drag anywhere on the ruler (or an empty lane) to scrub the playhead
– / + and zoom slider	Zoom the track lanes out/in
Volume slider (right end of the track toolbar)	Master preview volume
<b>Open timeline fullscreen / Exit fullscreen</b> (Esc)	Toggles fullscreen editing
Drag the divider above the tracks	Resizes the tracks panel vs. the preview

The preview also shows your subtitles live, styled exactly as they will be burned into the final video.

## Editing clips

Click a clip to select it — the **Tools** panel on the right shows everything you can do:

- **Move** — drag the body of a clip left/right along its track.
- **Trim** — drag the left or right edge of a clip to shorten or extend it.
- **Split at playhead** — park the playhead inside the clip and click to cut it in two.
- **Speed** — cycles the clip through **0.5x** → **1x** → **1.5x** → **2x**. Audio pitch-and-pace follows in the preview and the render.
- **Use asset duration** — snaps the clip back to the media file's real length.
- **Delete clip** — removes the clip (the asset itself stays in the project).

## Transitions

Select a video or image clip and pick a **Transition to next clip**:

Transition	Effect
<b>Cut</b>	Hard cut (no transition)
<b>Dissolve</b>	Crossfade into the next clip
<b>Black</b>	Fade to black between clips
<b>Fade White</b>	Fade to white between clips

A slider sets the transition length from **0.2s to 2.0s** (default 0.3s). Active transitions appear as a small marker sitting on the seam between the two clips — click it any time to edit.

## Subtitles

The **Subtitle** track carries text that gets burned into the rendered video.

- **Add subtitle** — drops a new subtitle at the playhead; edit its text in the Tools panel (changes are saved automatically).
- **Import** → **Subtitles (.srt)...** — load an existing SRT or VTT file. Subtitle files already in the project show up under **Project subtitles** in the same menu.
- **Subtitles from voice** — transcribes every Dialogue clip (and video clips containing speech) and lays the resulting cues onto the Subtitle track at the correct times. This replaces any existing subtitles.

### Subtitle style

Select the Subtitle track or any subtitle clip to open **Subtitle style (burned in)**:

- Quick presets: **Classic**, **Cinema** (yellow), **Minimal** (no outline), **Banner** (large, top of screen).
- **Size** slider (16–96), **Color** picker, **Outline** on/off.
- Position: **Top**, **Middle** or **Bottom**.

The style applies to all subtitles on that timeline, and the preview overlay updates instantly.

## Audio tools

- **Align voice** (track toolbar) — snaps every dialogue clip to its matching video clip (first voice line → first shot, and so on).
- **Align to video** (on a selected dialogue clip) — snaps just that voice clip's start to its video clip.
- **Fit video to voice** (on a selected dialogue clip) — resizes the paired video clip to the voice line's length; later clips on the track shift to follow.
- **Cut to beat** — detects the beats in your Music track and re-times every video/image cut to land on a beat. Add a music clip first; the button reports how many beats it found.
- **Track volume** — select a Dialogue, SFX or Music track to set its level from 0–200%.
- **Duck under dialogue** — on the Music track: automatically lowers the music whenever a dialogue clip is playing, both in the preview and the final render.

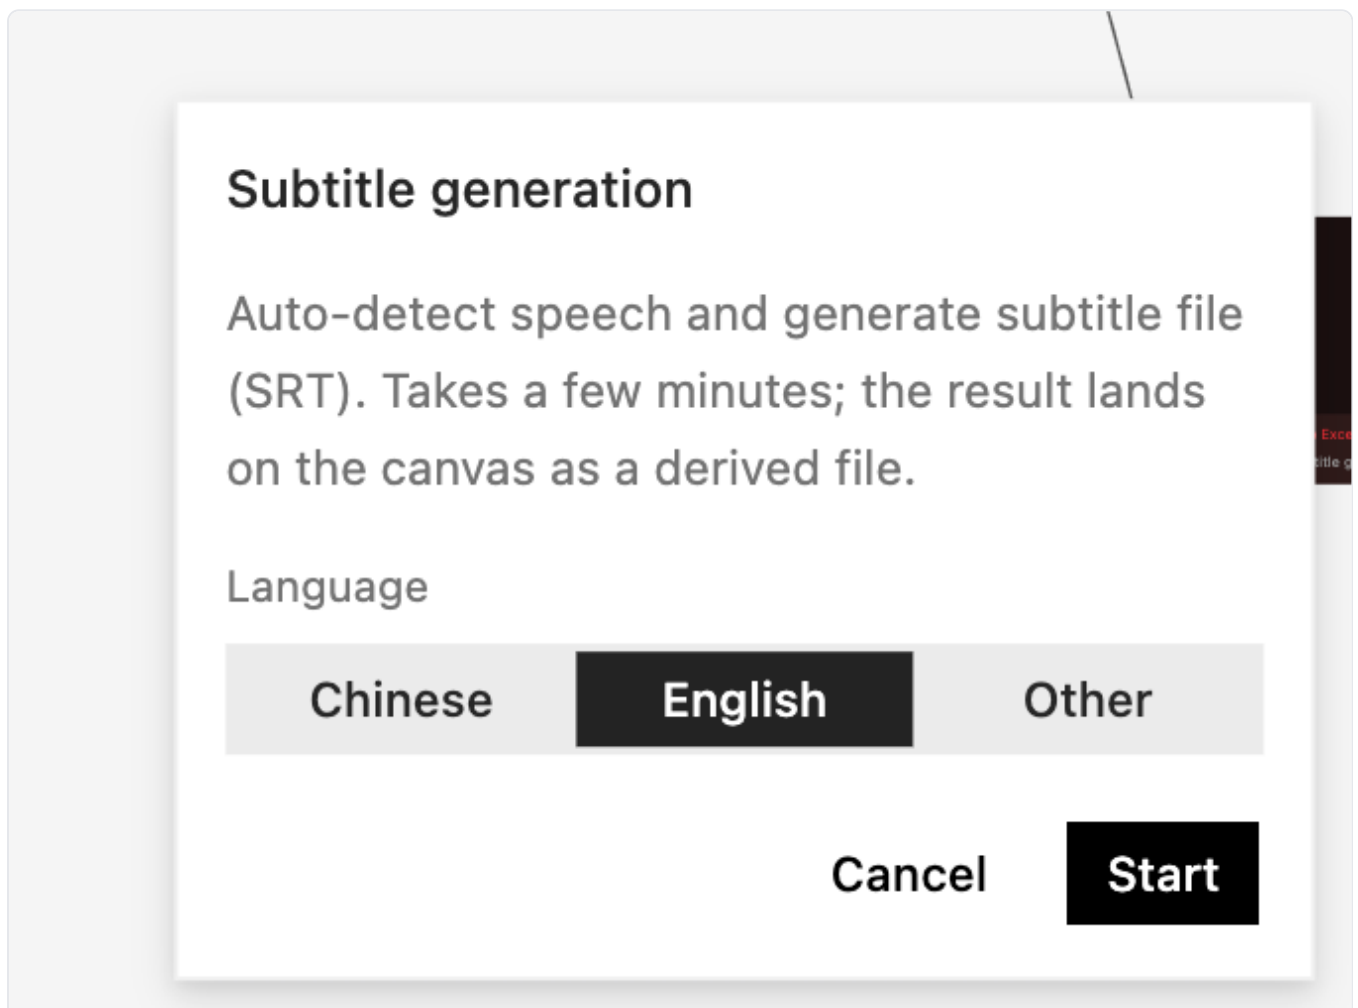
- **Mute** — every track has a speaker icon to mute/unmute it.

## Export & render

The top-right of the editor has everything you need to get your work out:

- **Compose Video** — renders the active timeline into a real MP4. You pick where to save the file; the app then composes all tracks (cuts, speeds, transitions, audio mix, music ducking and burned-in subtitles included) and shows the finished duration and file size when done.
- **Export** → **Bundle for editors (CapCut...)** — writes a folder with your ordered media files, an SRT and a manifest, ready to drop into an external editor.
- **Export** → **Subtitles (.srt)...** — saves just the subtitle track as an SRT file.
- **Export** → **Project JSON...** — saves the project data as a JSON file.

You can also render from chat: just ask the assistant to "*render the timeline*" and it composes the video for you and puts the result on the canvas. See [Chat & the Agent](#).



**Tip:** You rarely have to build a timeline by hand. Skills like `/promo-video` and `/beat-sync-editor` generate the media *and* assemble the timeline for you — then you just polish and press **Compose Video**. Browse them all in the [Skills Reference](#).

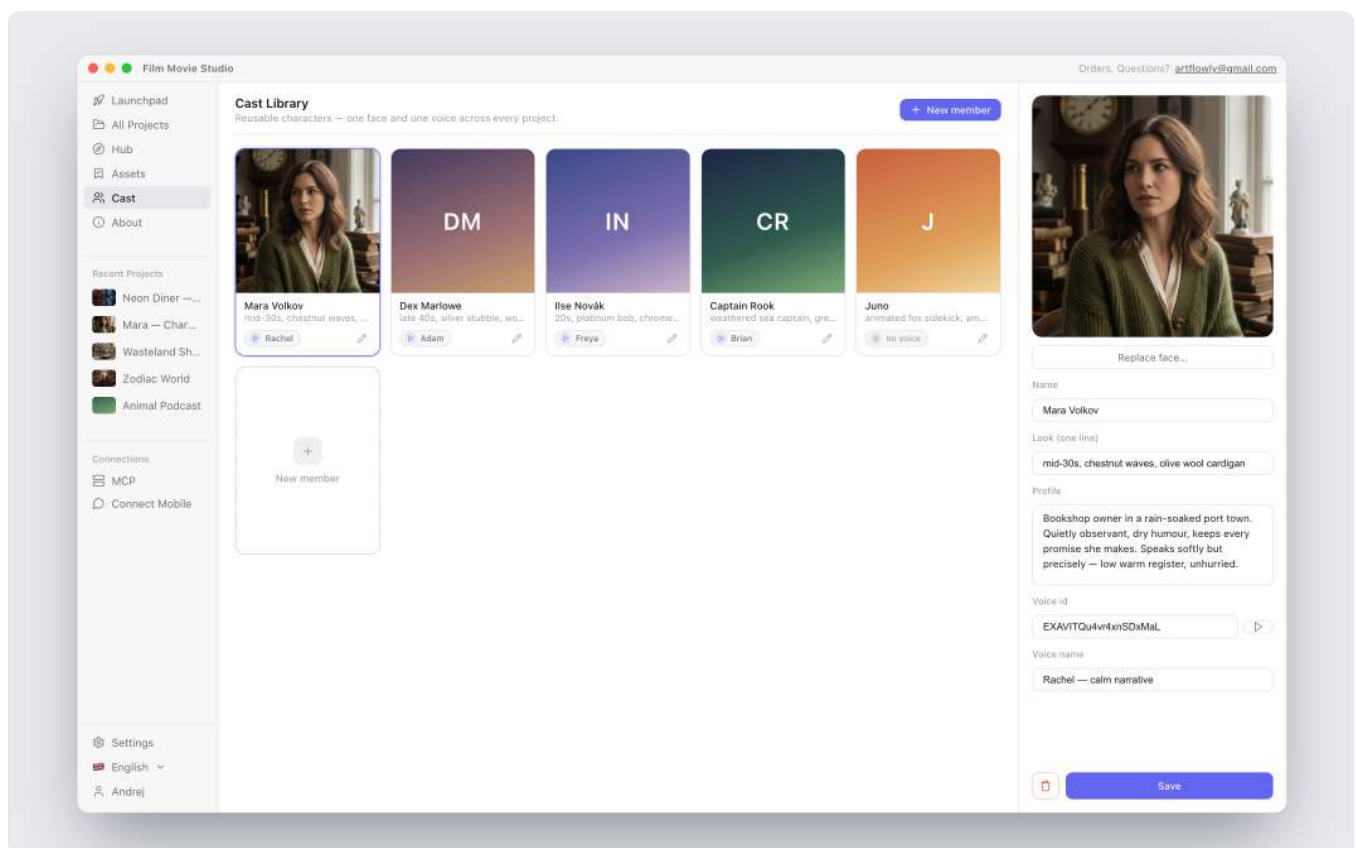
# Cast & Asset Libraries

Keep characters consistent across projects with the Cast library and reuse media via the Assets library.

Two global libraries follow you across every project: the **Cast Library** keeps your characters' faces and voices consistent, and the **Assets Library** bookmarks your favorite media so you can reuse it anywhere. Both live in the left sidebar — **Cast** and **Assets**.

## Cast Library

The problem the Cast Library solves: AI generations tend to reinvent a character's face and voice every time. A cast member locks both down — one face, one voice, in every project and every skill run.



Open **Cast** in the sidebar to see the grid of members — each card shows the face photo, name, @handle and voice.

## Creating a member

There are three ways:

- **From a canvas card** — hover a character image on the canvas and click **Save to Cast** in its toolbar (it changes to **In Cast** once saved). See [Canvas & Cards](#).
- **From chat** — just ask the assistant to save a character to the cast.
- **On the Cast page** — click **New member** and fill in the editor panel:

- **Choose face photo...** — pick the reference image that defines the character's face.
- **Name and Look (one line)** — a short appearance line, e.g. *"mid-20s, red curls, green wool coat"*.
- **Profile** — identity, appearance, personality, motivation and voice notes.
- **Voice id and Voice name** — bind a voice (an ElevenLabs voice\_id, a designed/cloned id, or a built-in one like `af_bella`). Click the play button (**Preview voice**) to hear a sample before saving.

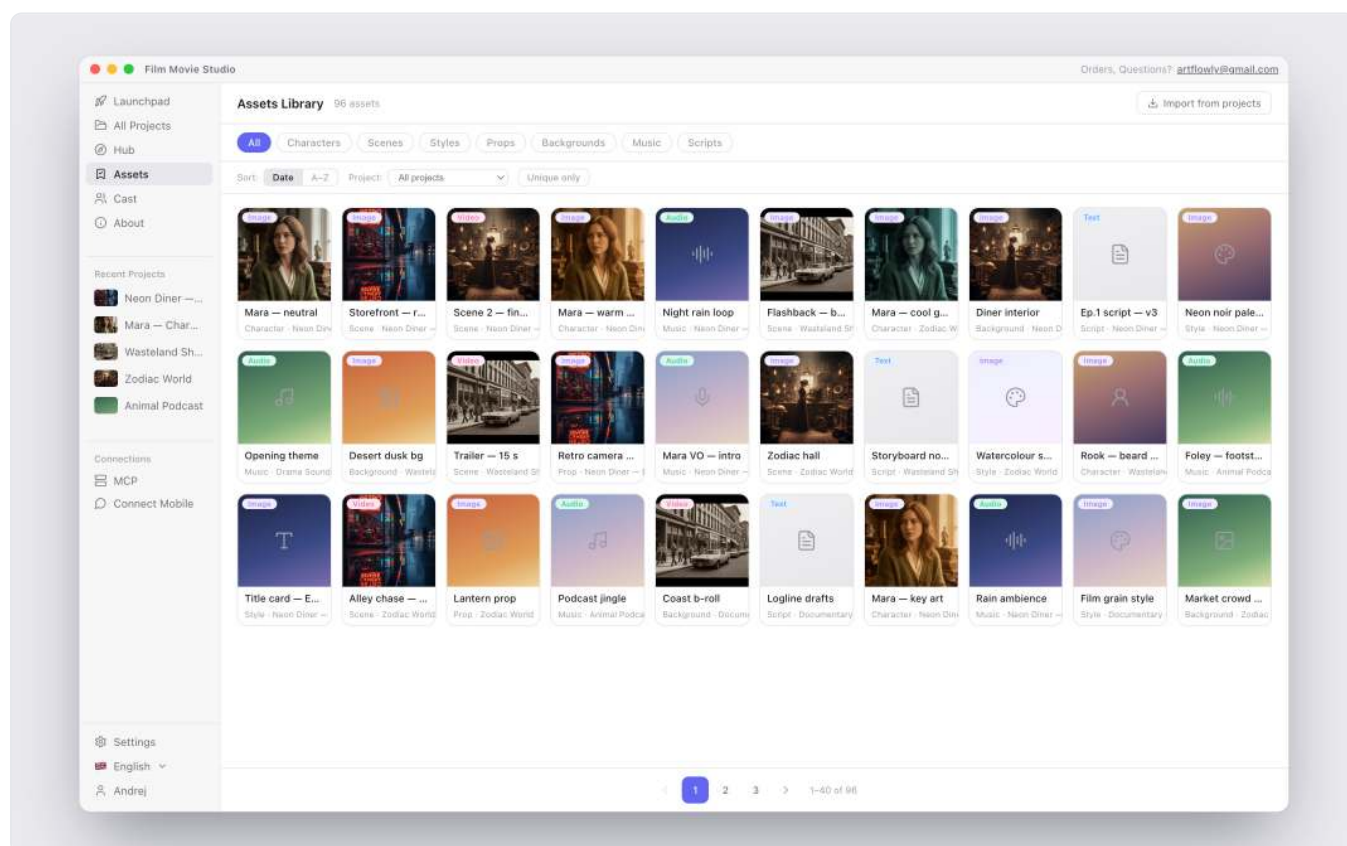
To edit a member, click their card, change anything and press **Save**. The trash button removes a member — past generations are not affected.

## Using cast members

- **@mention them in chat** — type `@` in the chat input and pick a member from the menu (they appear alongside your saved assets, with their `@handle`). The generation then uses their locked face as a reference and their bound voice for speech — automatically, with no extra prompting. See [Chat & the Agent](#).
- **Skills use the roster too** — the assistant always sees your cast. Workflow skills like **Short Drama** save your approved lead characters to the library automatically and reuse them for every keyframe and dialogue line, so episode 5 stars the same person as episode 1. See the [Skills Reference](#).
- **Cross-project continuity** — the library is global, so a character created in one project can star in the next one.

## Assets Library

The Assets Library is a cross-project bookmark shelf: any card you like — a character sheet, a style frame, a music track, a script — can be saved once and reused everywhere.



## Saving assets

Hover any canvas card and click **Add to Assets** in its toolbar; the icon changes to a filled bookmark (**In Assets**) once saved. On the Assets page you can also click **Import from projects** to bookmark everything from your existing projects in one go.

## Browsing

Open **Assets** in the sidebar. The page gives you:

- **Category chips** — **All, Characters, Scenes, Styles, Props, Backgrounds, Music, Scripts**.
- **Sort** by **Date** or **A–Z**, and a **Project** dropdown to see only assets from one project.
- **Unique only** — hides duplicate copies of the same file.

Click any asset to preview it (videos and audio play right in the popup) and choose:

- **Add to Canvas** — drops the asset onto the canvas of the project you have open.
- **Send to Chat** — attaches it to the chat as a reference for your next generation.

You can also reach your saved assets from the chat itself: type @ and pick from the **Assets** tab of the mention menu.

Removing an asset from the library (the bookmark button on its card) only removes the bookmark — the original file stays in its project.

**Tip:** Save your style frames and hero characters to the libraries early. Then every new project can start with "@mention" instead of regenerating from scratch.

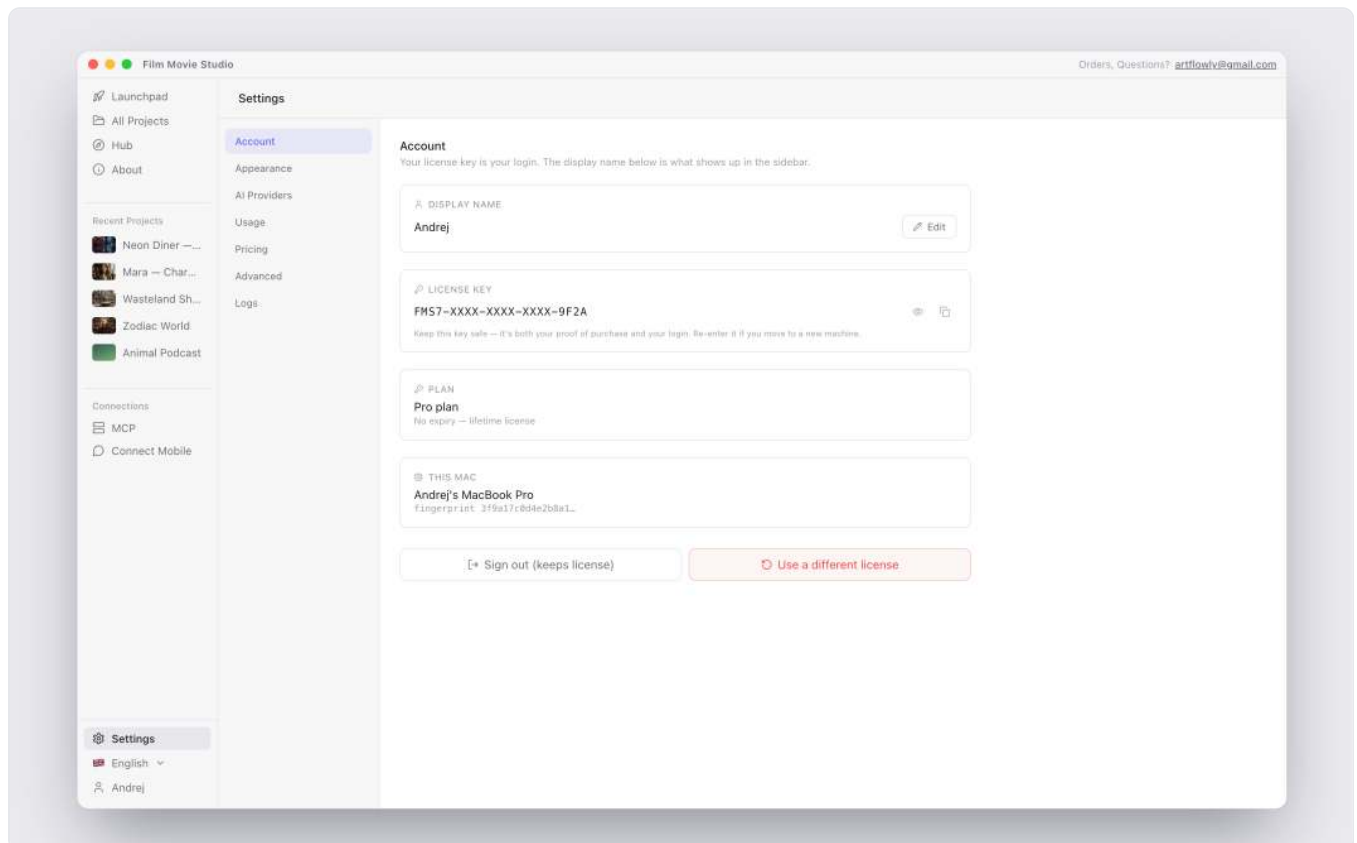
## Where your files live

Everything — projects, generated media, cast faces and bookmarked assets — is stored locally on your computer inside the app's data folder. Nothing is uploaded anywhere unless a generation needs it. Keep in mind that deleting a project removes its media files from disk (Assets Library bookmarks point to those original files), so export anything you want to keep before deleting the project it came from.

# Settings & AI Providers

API keys, model selection, the managed subscription with credits, free local models and usage.

Everything that shapes how FilmMovieStudio works for you lives in one place: **Settings**, reached from the sidebar or the account menu. This chapter walks through every section — your account, the look of the app, the AI providers and models that power generation, your subscription credits, and the advanced switches. If you haven't set up a provider yet, this is the chapter that gets your first generation working.



The left column lists the sections: **Account**, **Appearance**, **AI Providers**, **Usage**, **Pricing**, and **Advanced**.

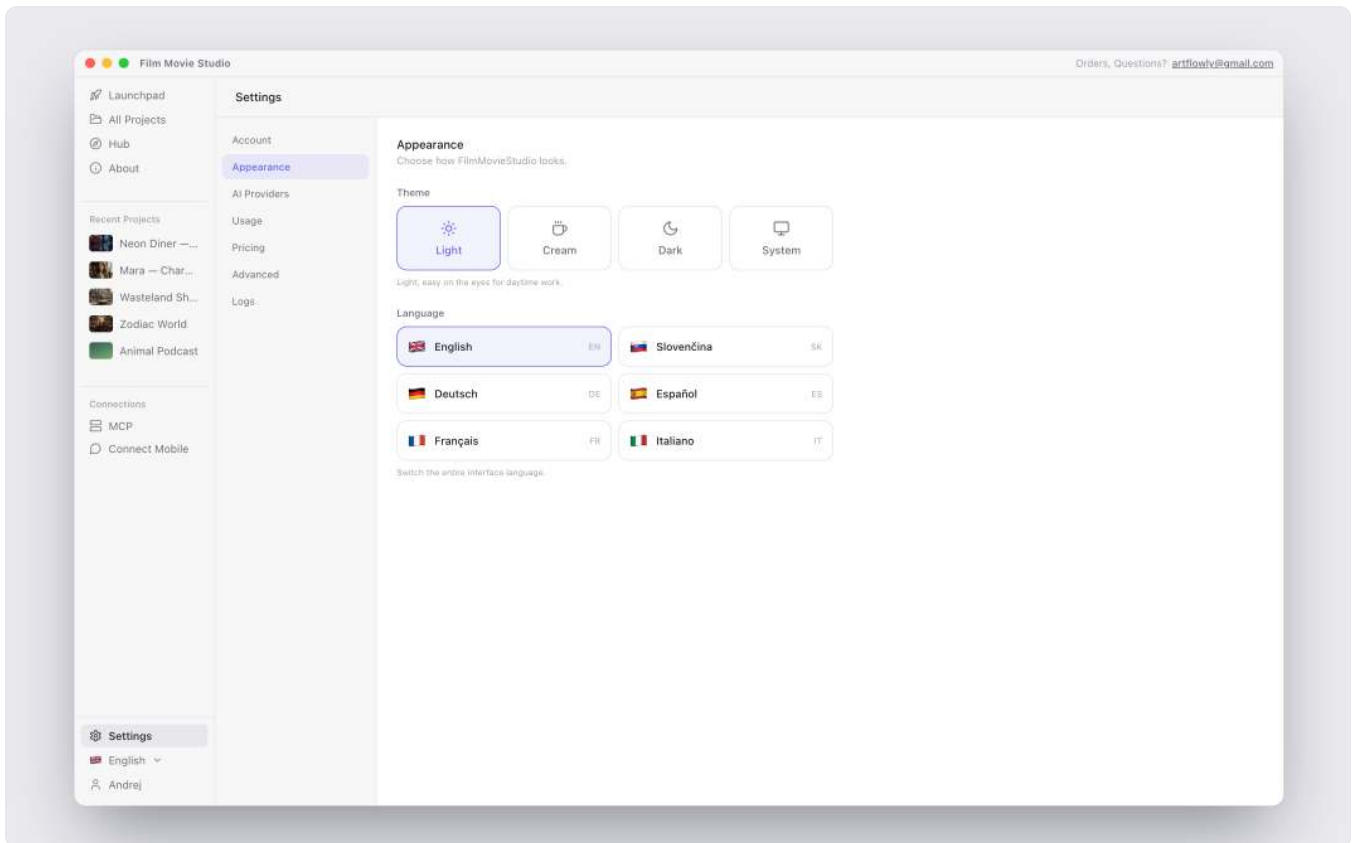
## Account

The Account section shows the identity you activated with in [Getting Started](#):

- **Display name** — what appears in the sidebar. Click **Edit** to change it.
- **License key** — shown masked, with reveal and copy buttons. Keep it safe: it's both your proof of purchase and your login, and you'll re-enter it if you move to a new machine.
- **Plan** — your license plan and expiry (lifetime licenses show "No expiry — lifetime license").
- **This Mac** — the machine name and fingerprint this activation is bound to.

Two buttons at the bottom are deliberately separate: **Sign out (keeps license)** logs you out without touching the activation, while **Use a different license** removes the key from this machine so you can enter another one.

# Appearance

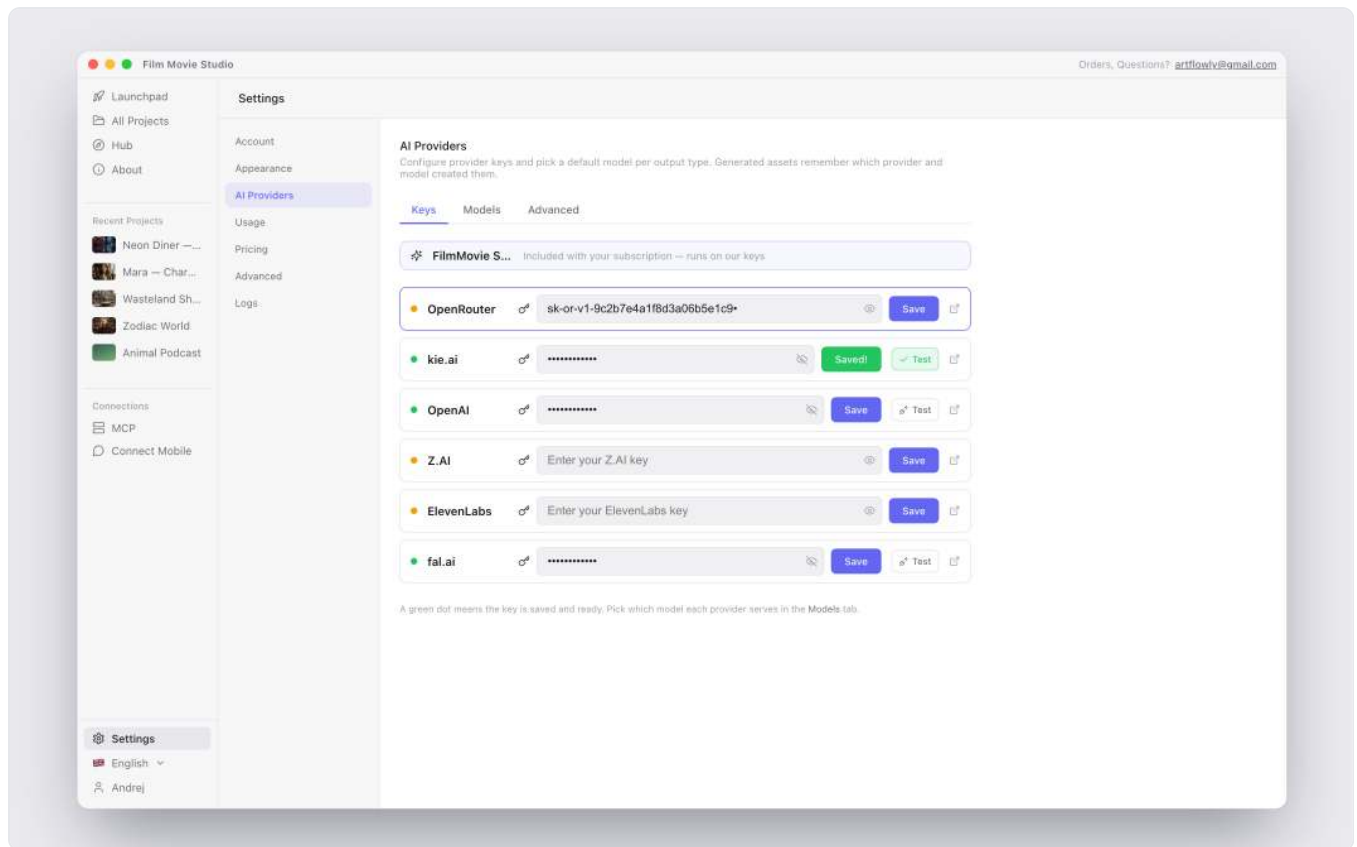


Pick a **Theme** — **Light**, **Cream** (a warm ivory look), **Dark**, or **System** (follows your OS setting) — and the interface **Language**. Six languages are available (English, Slovak, German, Spanish, French, Italian); switching applies instantly, no restart needed.

## AI Providers

This is the engine room. FilmMovieStudio orchestrates external AI providers rather than generating media itself, and this section has four tabs: **Keys**, **Models**, **Advanced**, and **Local**.

## Keys tab — connect your providers



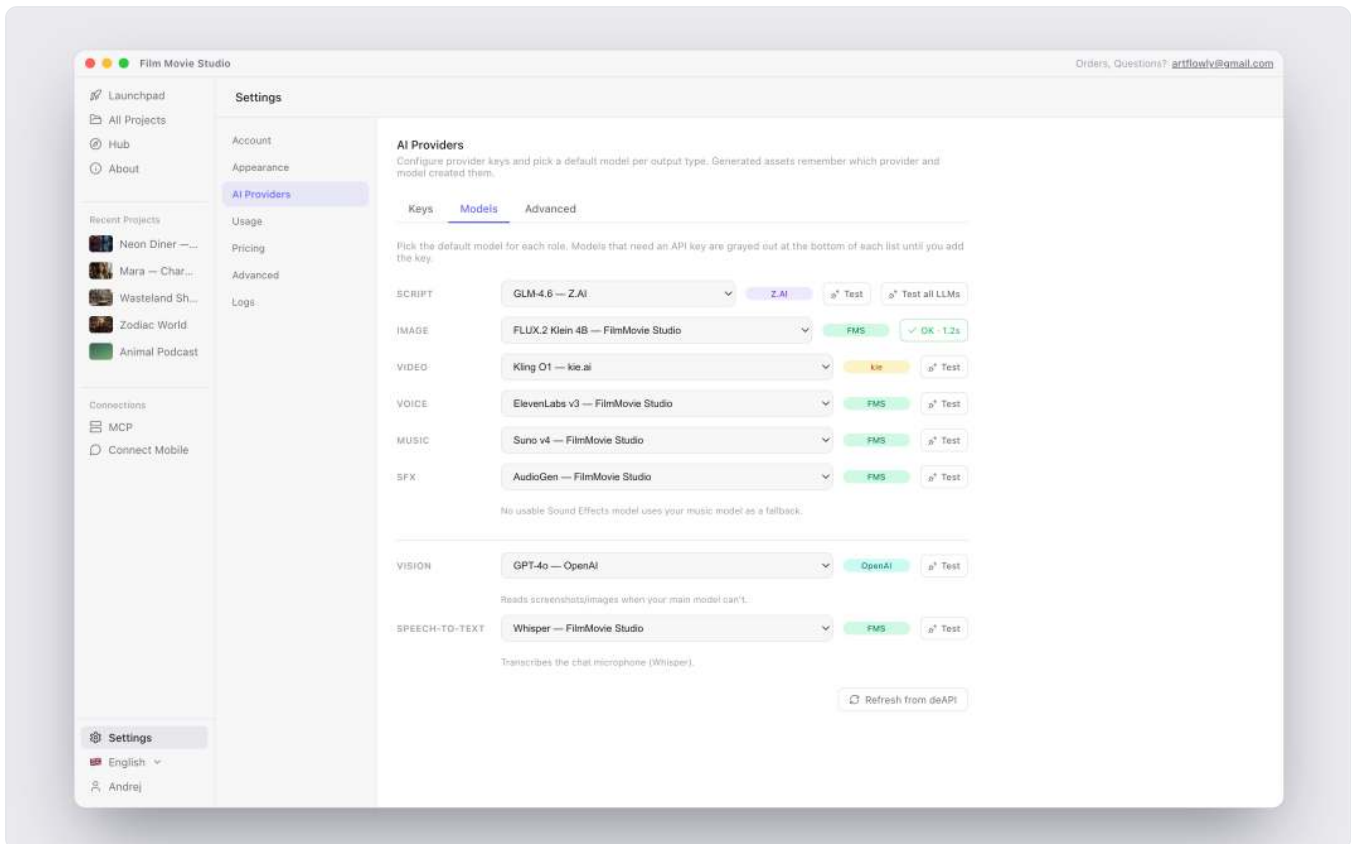
Each provider gets one row: a status dot (green = key saved and working), the key field, a **Test** button, and a link to the provider's website. Paste a key, click **Save**, then **Test** to verify it.

Provider	What it powers	Where to get a key
<b>FilmMovie Studio</b> (managed)	Image, video, voice & music on our keys — no key to paste, unlocked by your subscription	<a href="https://filmotvor.ai/studio">filmotvor.ai/studio</a>
<b>Local Models</b>	Free on-device voice + script models — no key, works offline	No key needed — see the Local tab
<b>OpenRouter</b>	Script / text LLMs (Claude, Gemini, Llama, ...)	<a href="https://openrouter.ai">openrouter.ai</a>
<b>kie.ai</b>	Premium media — Nano Banana images, Seedance video, Suno music	<a href="https://kie.ai">kie.ai</a>
<b>OpenAI</b>	GPT models for scripts + GPT Image / DALL-E 3 for images	<a href="https://platform.openai.com/api-keys">platform.openai.com/api-keys</a>
<b>Z.AI</b>	Script / text generation (GLM) — a great default writing model	<a href="https://z.ai">z.ai</a>
<b>ElevenLabs</b>	Premium text-to-speech voices	<a href="https://elevenlabs.io">elevenlabs.io</a>
<b>fal.ai</b>	Identity-lock image generation (locks a face from a reference photo)	<a href="https://fal.ai/dashboard/keys">fal.ai/dashboard/keys</a>

**Tip:** You can mix freely — for example the subscription for images and video, your own Z.AI key for scripts, and the free local voice model. If the top "FilmMovie Studio" row looks blurred, you just haven't

activated a subscription code yet (see Usage below).

## Models tab — pick a default model per role



Every generation role gets one dropdown: **Script**, **Image**, **Video**, **Voice**, **Music**, and **SFX**, plus three helper roles below the divider — **Agent** (runs skill/tool rounds; pick a fast model here), **Vision** (reads screenshots and images), and **Speech-to-text** (transcribes the chat microphone). Whatever you select becomes the default the chat and the agent use — see [Chat & the Agent](#).

- Models are grouped by provider inside each dropdown. Models whose provider has no key yet appear grayed out at the bottom under "Needs a key — add in API Keys".
- The **Test** button next to each dropdown runs a real sample generation to verify the selected model (media tests spend credits and save the result to a separate "TEST — Model probes" project).
- **Test all LLMs** on the Script row pings every text model at once and shows response times — handy for choosing the fastest one.
- If no SFX model is available, sound effects automatically fall back to your Music model.

The managed catalog also updates itself: in the chat's model picker you'll find a **Refresh** button ("Fetch latest models from deAPI") that pulls newly released managed models without an app update.

st × clip-export × skill-creator 2 ×

Models Select all · Deselect all

Image (5/5) Video (6/6) Audio (2/2)

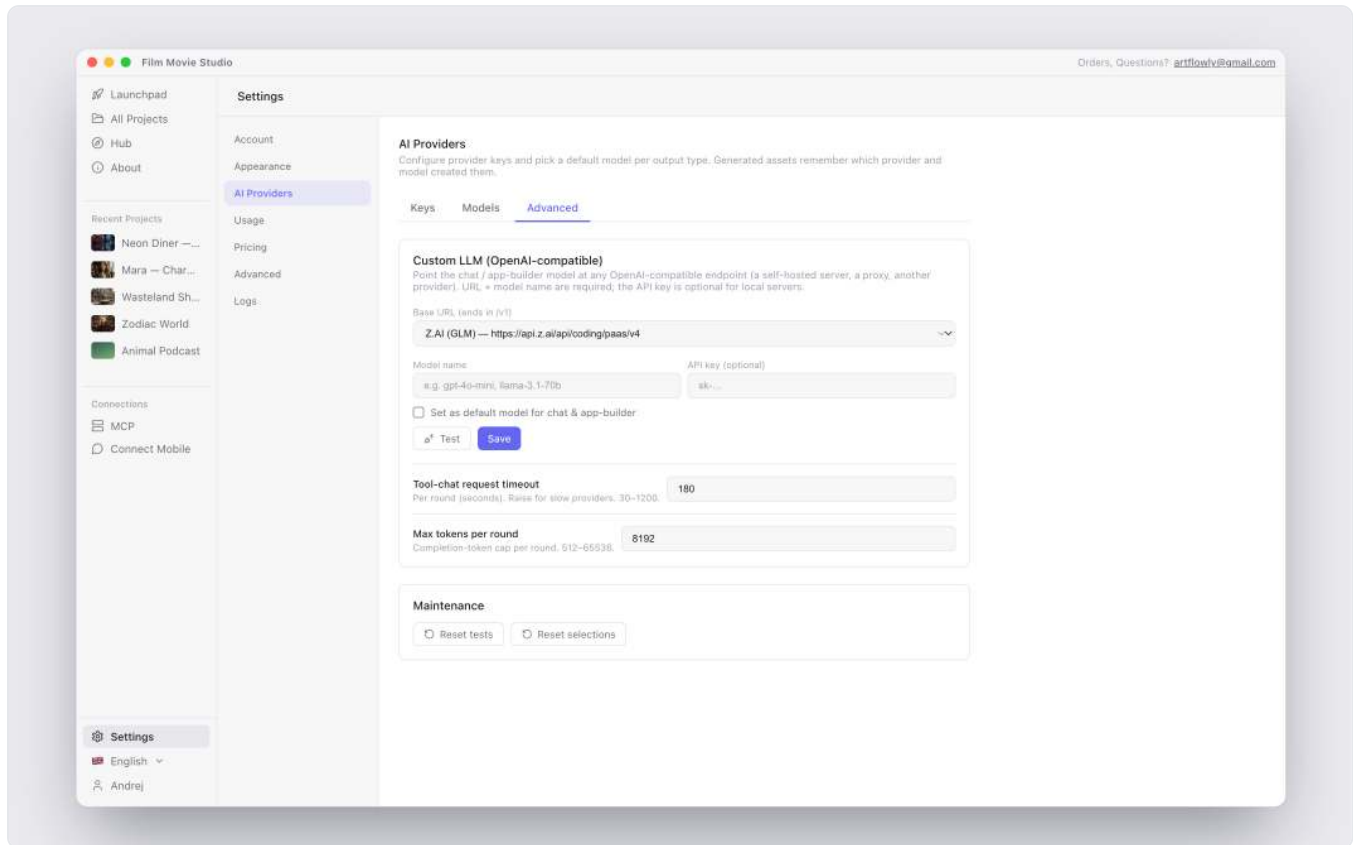
- Nano Banana Series(Pro/2)
- GPT Image 2
- Seedream Series(4.5/5.0 Lite)
- Midjourney Series(V8.1/V7/Niji7)
- Kling Series(v3/o1)

Type / to explore  
Tap here to view

Agent can use all models in this category Esc to discard

+ Skills Models

## Advanced tab



Power-user options:

- **Custom LLM (OpenAI-compatible)** — point the chat at any OpenAI-compatible endpoint. Presets are included for Z.AI, OpenRouter, OpenAI, kie.ai, **Ollama (local)** and **LM Studio (local)**; enter the model name, an optional API key, tick **Set as default model for chat & app-builder** if you want it everywhere, then **Test** and **Save**.
- **Tool-chat request timeout** — seconds per agent round (30–1200). Raise it for slow providers.
- **Max tokens per round** — lower it if your provider reports "out of credits", raise it for big builds.
- **Maintenance** — **Reset tests** re-arms every model's Test button (useful after rotating a key); **Reset selections** re-enables every model if some got switched off.

## Local tab — free, private, offline

Two models run entirely on your computer: no API key, no credits, and they keep working without internet once downloaded.

Model	Role	Download size	What it does
<b>Supertonic 3</b>	Voice (TTS)	~130 MB	Narration voices in 31 languages, including Slovak
<b>Qwen3 4B</b>	Script	~2.5 GB	Drafts scripts and chats offline — a solid free draft tier

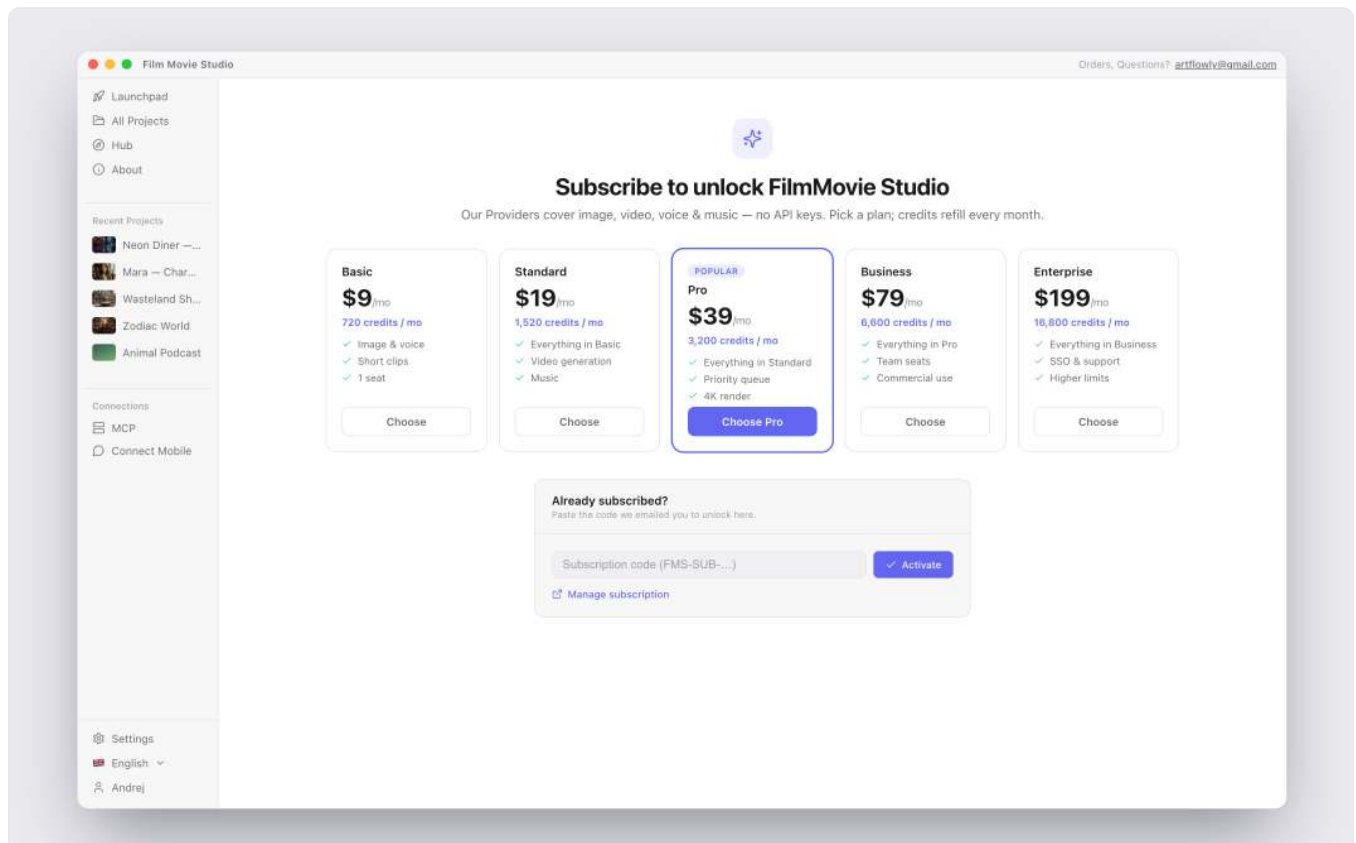
Click **Download** on a card; a progress bar shows downloading → extracting → verifying, and **Set as default** appears once it's ready. Supertonic 3 offers ten voices — **F1–F5** (female) and **M1–M5** (male) — selectable

when you generate audio. The script model runs a small local engine that starts automatically when used and shuts itself down when idle; delete a model anytime with the trash button to reclaim disk space.

**Tip:** Local models are fast on Apple Silicon and modern PCs, and a little slower on older Intel machines — fine for narration, but be patient with long scripts.

## Usage — subscription & credits

The Usage section holds two cards: **Subscription** and **Credit history**.



**Redeeming your code:** after you subscribe on [filmotvor.ai/studio](https://filmotvor.ai/studio), your subscription code arrives by email. Paste it into the **Subscription code** field and click **Activate** — the "FilmMovie Studio" provider row unblurs immediately and your credits appear. **Manage subscription** opens your billing page on the website; **Remove** detaches the code from this app.

Once active, the card shows your remaining **Credits** and the **Renews** date. Credits refill monthly with your plan:

Plan	Credits / month
Basic	720
Standard	1 520
Pro	3 200
Business	6 600

Plan	Credits / month
Enterprise	16 800

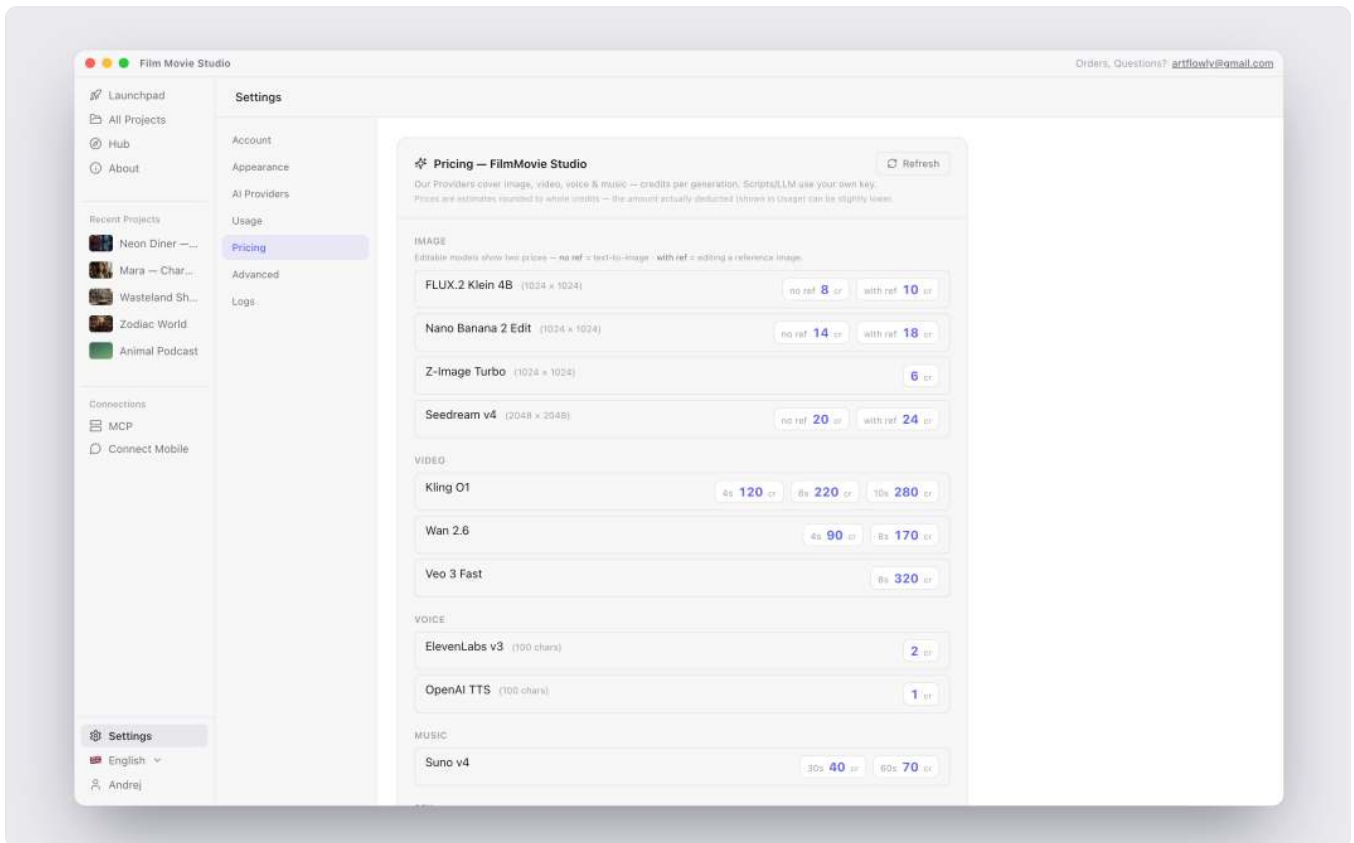
(200 credits correspond to \$1 of generation value.)

The screenshot shows the 'Usage' section of the Film Movie Studio settings. The subscription is active and 'SUBSCRIBED'. The current credit balance is 1,412 cr, with a renewal date of July 29, 2026. Below this is a 'Credit history' table showing the most recent generations.

WHEN	PROJECT	TYPE	MODEL	PARAMS	PRICE
2h ago	Neon Diner — Ep.1	Video	Kling 01	1280x720 · 8s	220 cr
2h ago	Neon Diner — Ep.1	Voice	ElevenLabs v3	210 chars · narration	5 cr
2h ago	Neon Diner — Ep.1	Image	FLUX.2 Klein 4B	1024x1024 · with ref.	10 cr
3h ago	Neon Diner — Ep.1	Image	FLUX.2 Klein 4B	1024x1024	8 cr
yesterday	Wasteland Short Film	Video	Wan 2.6	1280x720 · 4s	90 cr
yesterday	Wasteland Short Film	Music	Suno v4	60s · synth bed.	70 cr
yesterday	Mara — Character Set	Image	Seedream v4	2048x2048	20 cr
2d ago	Animal Podcast	Voice	OpenAI TTS	320 chars	3 cr

**Credit history** lists every generation charged to your subscription — newest first — with the project, type, model, parameters, and the exact price in credits. It's the best place to see where your credits actually go.

# Pricing



The Pricing section lists every managed model and its live per-generation cost in credits: images by size (with separate "no ref" / "with ref" prices for editing models), video by duration, music by length, and voice per characters. Prices are estimates rounded to whole credits — the amount actually deducted (shown in Usage) can be slightly lower. As reference points from the canvas card toolbar: **Upscale** costs about **100 credits** and **Remove background** about **20 credits** per image (see [Canvas & Cards](#)). Click **Refresh** to re-fetch live prices.

## Advanced

The last section holds the **Enable Automation** switch — it lets the AI run code on your machine for file and data tasks, and you approve every action before it runs. Turning it on shows a consent dialog explaining exactly what it means; a **Reset session** button appears for clearing a stuck automation session. Scripts, spreadsheets and documents produced this way are saved into your project's assets folder.

Some installs also show a **Logs** section with the app version, data folder, an **Open Logs** button and an update check — see [Troubleshooting & FAQ](#) for where logs and data live on disk.

## Advanced: MCP, Telegram & More

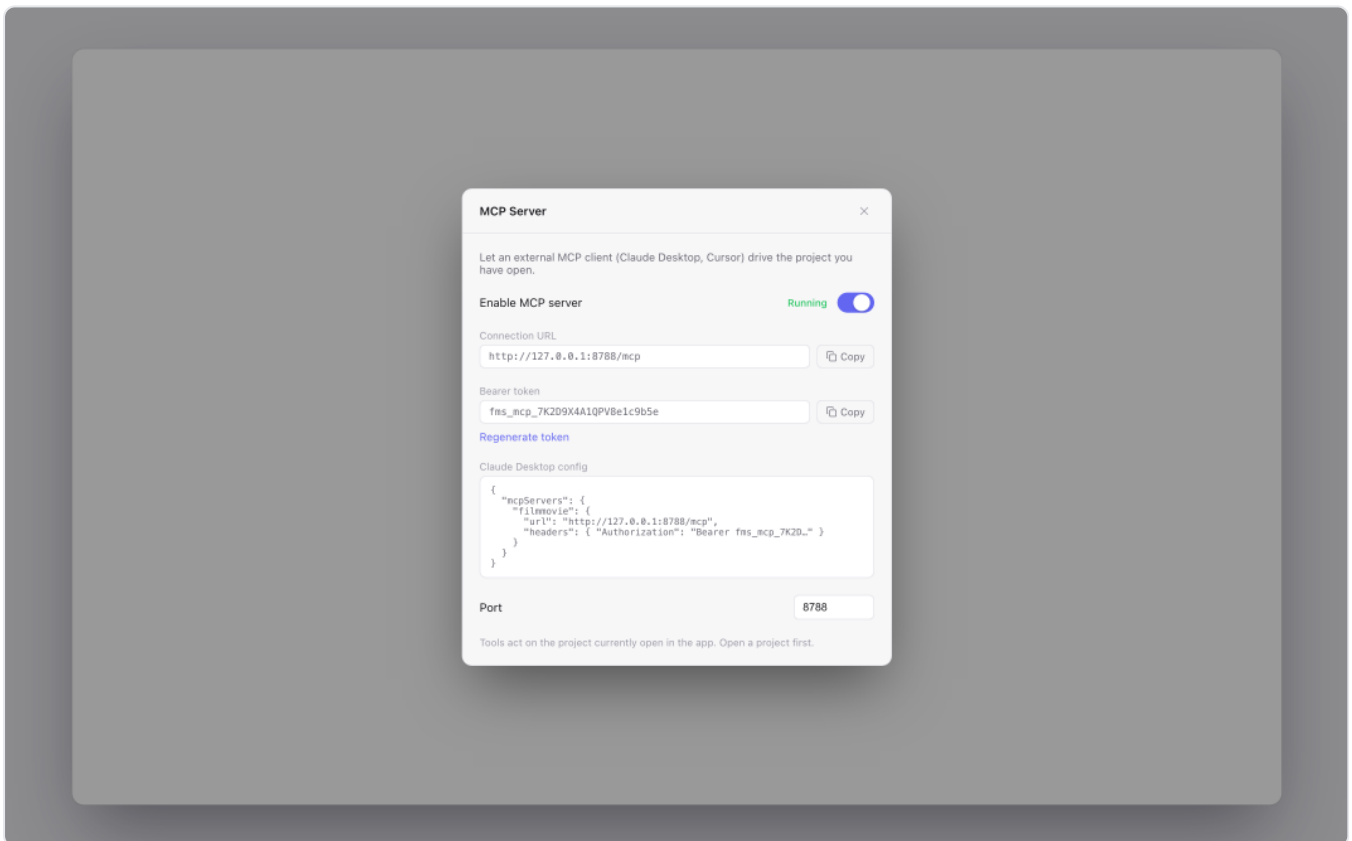
Remote-control the studio via MCP or Telegram; App Builder, games, documents, automation and export.

FilmMovieStudio doesn't stop at the canvas and the timeline. This chapter covers the power features: driving the studio from an external AI assistant over MCP, controlling it from your phone via Telegram, and the four "beyond video" skill families — apps, games, documents and automation — plus how to export a project.

### MCP server — drive the studio from another AI

FilmMovieStudio can act as an **MCP server** (Model Context Protocol), so an external MCP client — Claude Desktop, Cursor, or any other MCP-capable assistant — can drive the project you have open: create projects, generate images, video, voice and scripts, place and edit cards on the canvas, read your assets, run skills, and even open and render the timeline.

Open it from the sidebar: under the **Connections** header, click **MCP**.



To set it up:

1. Flip **Enable MCP server** — the status changes to **Running**.
2. The panel reveals the **Connection URL** (by default `http://127.0.0.1:8765/mcp`) and a **Bearer token**. Both have Copy buttons.

3. Copy the ready-made **Claude Desktop config** snippet straight into your MCP client's configuration — it already contains the URL and the `Authorization: Bearer ...` header.
4. If the default port clashes with something else on your machine, change **Port** and re-enable.

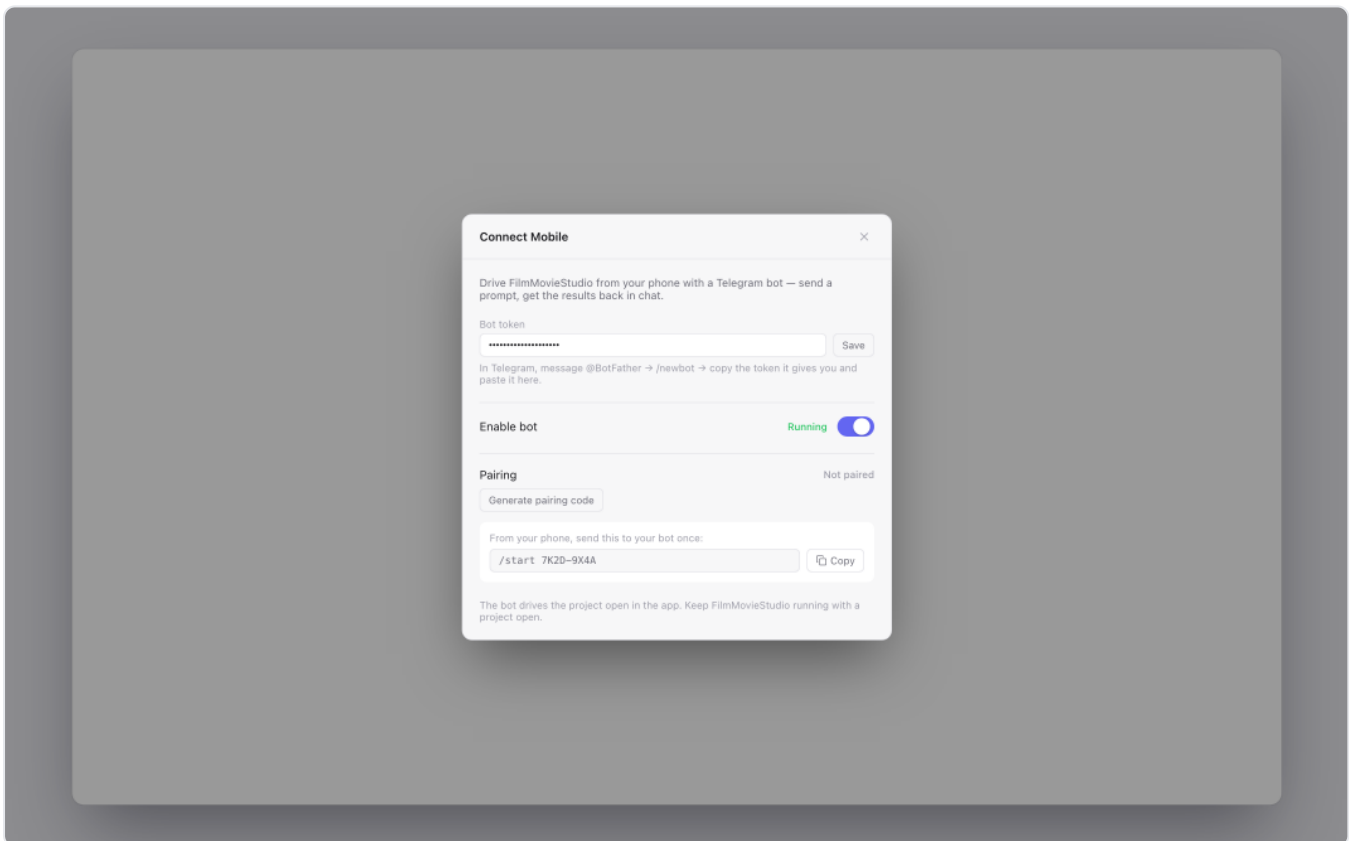
Use **Regenerate token** any time you want to invalidate old clients. The server only listens on your own computer (127.0.0.1), never the network.

One important rule, quoted from the panel itself: *"Tools act on the project currently open in the app. Open a project first."* Keep FilmMovieStudio running with a project open, and everything the external assistant generates appears live on that project's canvas — exactly as if you'd asked in the chat yourself.

**Tip:** This is a great way to script bulk work — for example, asking Claude Desktop to "create 10 product images from this list and put them on the canvas" while you keep working in another app.

## Connect Mobile — drive the studio from your phone

**Connect Mobile** pairs FilmMovieStudio with a Telegram bot so you can send prompts from your phone and get the results back in the chat — images, clips, and voiceovers included. Open it from the sidebar under **Connections** → **Connect Mobile**.



Pairing takes three steps:

1. **Create a bot:** in Telegram, message **@BotFather**, send `/newbot`, and copy the token it gives you. Paste it into the **Bot token** field and click **Save**.
2. **Enable the bot:** flip **Enable bot** — the status changes to **Running**.

3. **Pair your phone:** click **Generate pairing code**, then from your phone send the shown `/start <code>` message to your bot once. The panel switches to **Paired**.

From then on, anything you send the bot is handled exactly like a typed chat message: plain prompts go to the general assistant, `/skill` commands launch workflows, and when the agent asks a question mid-workflow you can answer it right from Telegram. Generated media is sent back to your phone as photos, videos and audio files.

As with MCP: *"The bot drives the project open in the app. Keep FilmMovieStudio running with a project open."* Use **Unpair this device** to disconnect a phone.

## App Builder

Type `/app-builder` in the chat and describe an app — "a habit tracker with streaks", "a recipe box with categories" — and the agent designs a working mobile app: screens, navigation, real data collections, lists bound to that data, forms, and actions. The result appears as an interactive preview card on the canvas that you can click through, and the agent iterates on your feedback ("make the header purple", "add a detail screen"). When you're happy, you can export clean React code from the card's toolbar. It's a genuine no-code builder driven entirely by conversation — see the [Skills Reference](#) for the full command list.

## Games

Two skills build playable games right on the canvas. `/games` creates 2D browser games — puzzles, arcade shooters, platformers — while `/threejs-game` is the 3D specialist, complete with generated skyboxes, textures, sprites, sound effects and a music bed, so the result looks and sounds finished rather than programmer-art. Games appear as a card with a **Play** button, run offline once built, and can be exported as a single HTML file to share with anyone. If something breaks, the card shows the error and the agent fixes its own code. Start simple, then iterate: "add power-ups", "make the enemies faster". More in the [Skills Reference](#).

## Documents

The `/documents` skill reads and writes office files. Attach a PDF, Word, Excel or CSV file and ask questions about it, have it summarized, or pull tables out of it. In the other direction, ask for output as a real file — "turn this research into a formatted PDF report", "make a spreadsheet of these results" — and the agent generates a proper `.docx`, `.pdf` or `.xlsx` that lands on the canvas as a document card you can open or reveal in your file manager. Everything runs natively inside the app; no external office software needed. See the [Skills Reference](#).

## Automation

`/automation` lets the AI run Python on your computer to automate real work: renaming and organizing files, batch-converting images, crunching CSVs, filling spreadsheets, scraping data into reports. It's off by default — enable it in **Settings** → **Advanced** → **Enable Automation** — and every block of code is shown to you on an approval card **before** it runs; nothing executes without your click. Produced files are saved into the current project's assets folder so they show up alongside your media. Treat it like a very capable intern: powerful, but read what you approve. Details in the [Skills Reference](#).

## Exporting a project

Your main output is the rendered **MP4** from the timeline — but projects themselves can be exported too:

- **Export Project** (from the project's menu) saves a **Project Bundle (.json)** — project data, workflow, timeline and asset references — useful for backups.
- The timeline's export menu offers **Bundle for editors (CapCut...)** — a folder with your ordered media files, an `.srt` subtitle file and a manifest, ready to drop into an external editor — plus **Subtitles (.srt)** and **Project JSON** on their own.
- Individual cards export from their toolbar: images as PNG, games as HTML, app-builder cards as React code — see [Canvas & Cards](#).

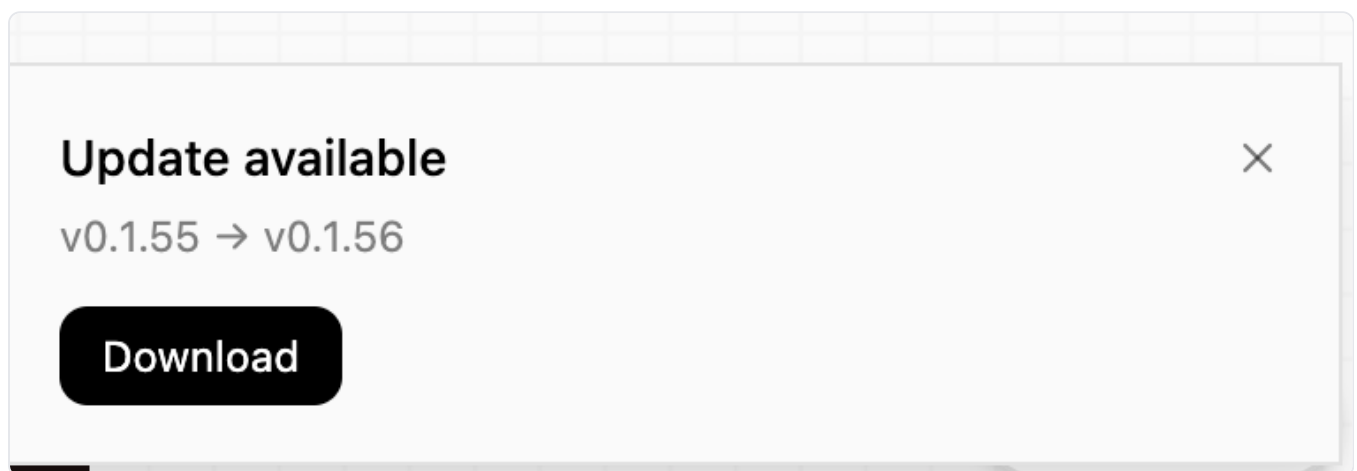
# Troubleshooting & FAQ

Common errors and fixes, where to find logs, updates and how to get support.

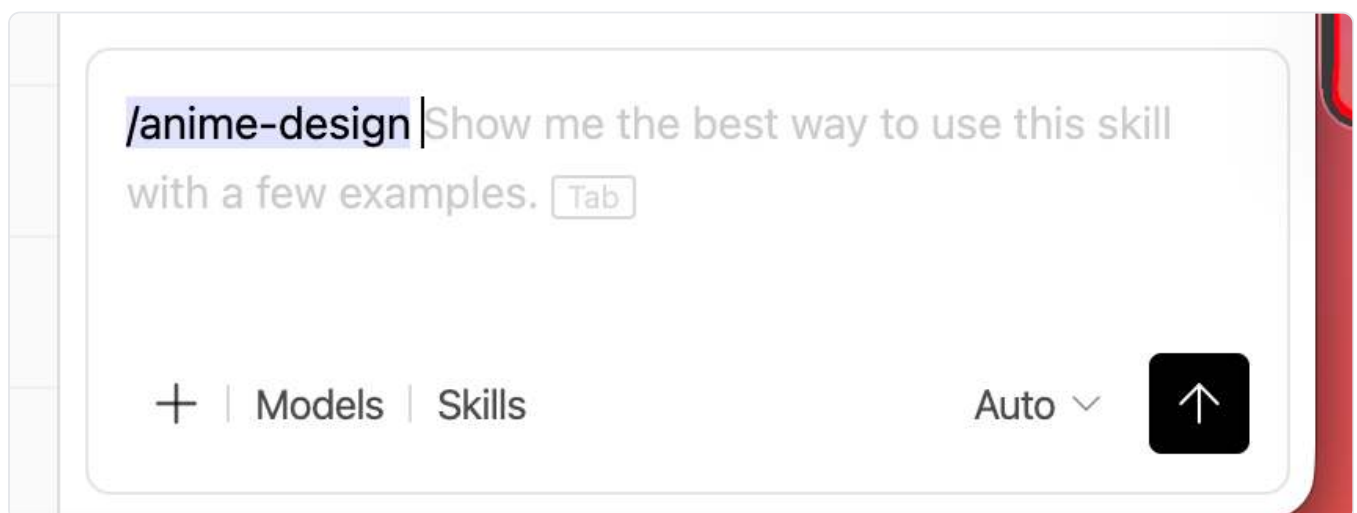
Most problems in FilmMovieStudio come down to one of three things: a provider key or subscription issue, a model that isn't configured yet, or a workflow that paused and is waiting for you. This chapter shows you where to look — updates, your data, the logs — and gives a quick-fix table for the most common errors.

## Updates

When a new version is released, a banner appears at the top of the app: **Update available**, with a **Download** button and a **Dismiss** button.



**Download** opens the official download page at [filmotvor.ai/studio](https://filmotvor.ai/studio) — always get installers there. Install the new version over the old one; your projects, settings, keys and license are untouched (they live in your user data folder, not inside the app). You can also check manually anytime from the Help/Guide tab or the download page.



## Where your data lives

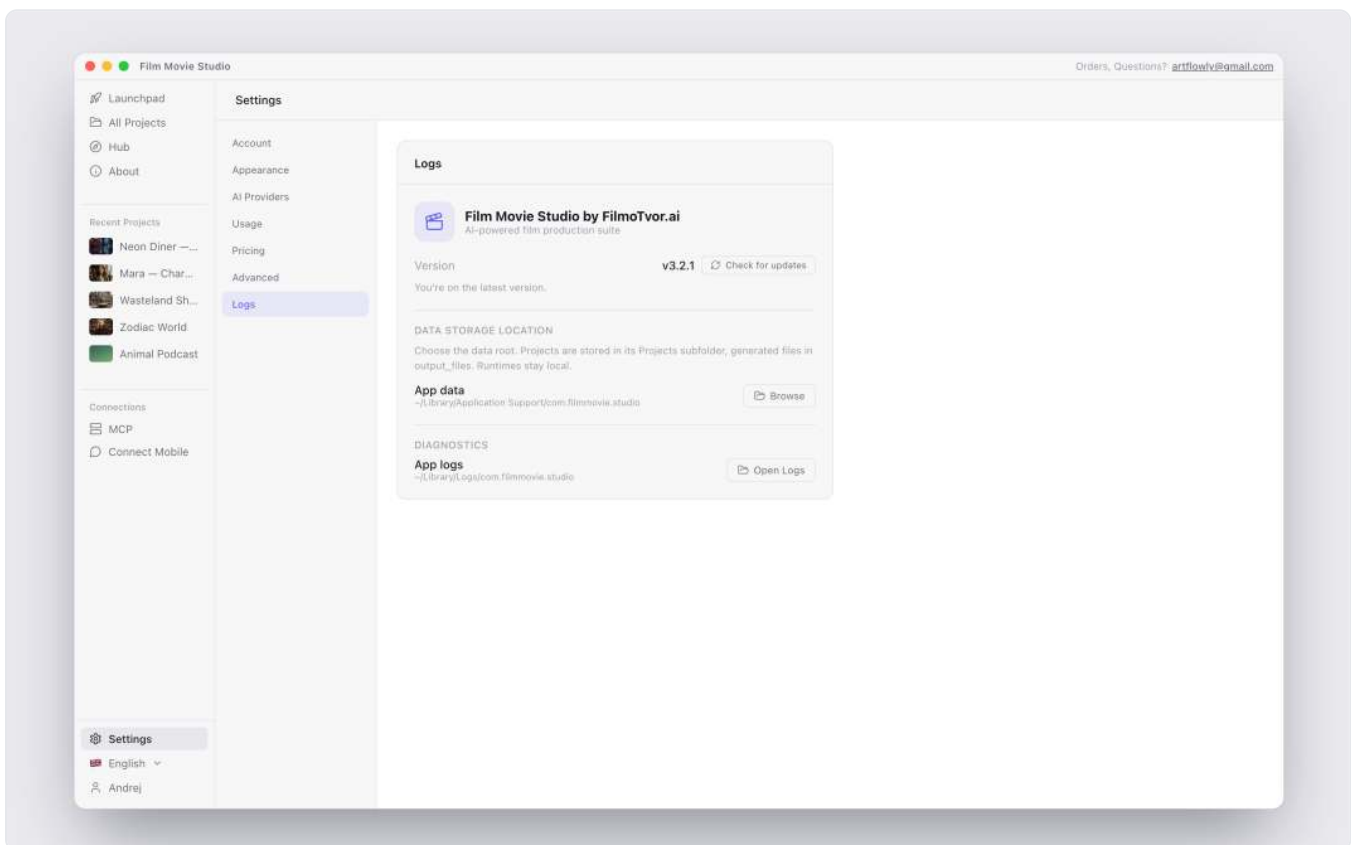
Everything you make stays on your computer. On macOS:

- **App data (projects, media, database):** `~/Library/Application Support/com.filmmovie.studio/`  
— each project keeps its generated media under `projects/<id>/assets/`.
- **Logs:** `~/Library/Logs/com.filmmovie.studio/`

Back up the app data folder and you've backed up everything. On Windows and Linux the equivalents live under your user profile's application-data directories.

## Logs

The app writes one log file per category per day (for example `agents-2026-07-05.log`, `media-...`, `render-...`) into the logs folder above. When something fails, the matching log almost always contains the real error message from the provider — which is exactly what support needs.



Some installs also show a **Logs** page inside Settings with the app version, the data storage location, an **Open Logs** button, and a **Check for updates** action.

## Common problems

Problem	Cause	Fix
"Set up a Script model" dialog blocks the chat	No text/LLM model is configured	Click <b>Open AI Providers</b> and add a key for a Script provider (Z.AI, OpenRouter, OpenAI or kie.ai), activate your subscription, or download the free local Qwen3 4B model in the <b>Local</b> tab
Generation fails with 401 / 403	The provider key is invalid/expired, or your credits ran out	Re- <b>Test</b> the key in Settings → AI Providers → Keys; check your credit balance (Settings → Usage, or the provider's billing page)
"This model supports a maximum of N input image(s)"	Too many reference images attached to one generation	Remove extra references — keep only the most important ones (character/face refs first) or pick a model that accepts more inputs
The "FilmMovie Studio" provider row is blurred	No active subscription code	Subscribe at <a href="https://filmotvor.ai/studio">filmotvor.ai/studio</a> , then paste the emailed code in Settings → Usage → Subscription → <b>Activate</b>
"kie.ai credits exhausted" toast	Your kie.ai balance hit zero mid-generation	Top up at kie.ai → Billing, then type <code>/continue</code> in the chat — the workflow resumes from the failed step
Workflow paused with "provider rate-limited the request"	Too many requests in a short time (HTTP 429)	Wait about 60 seconds, then type <code>/continue</code>
A model in the Models tab is grayed out	Its provider has no key, or (for Local models) it isn't downloaded	Add the key in the <b>Keys</b> tab, or click <b>Download</b> in the <b>Local</b> tab
Local models feel slow	Older Intel Macs run them slower than real time	Normal — use them for short narration/drafts, or switch that role to a cloud model
A skill stopped mid-run / app was closed	The session paused — progress is saved	Type <code>/continue</code> to resume, <code>/status</code> to see where it stands, <code>/redo &lt;what&gt;</code> to regenerate one part, <code>/start-over</code> to clear it
No sound effects model available	No SFX provider configured	SFX automatically falls back to your Music model — or add an ElevenLabs key
Chat microphone doesn't transcribe	No Speech-to-text model set	Settings → AI Providers → Models → <b>Speech-to-text</b> : add an OpenAI key or use the subscription's Whisper
Non-English prompts produce garbled text/speech	Affordable models work best in English	Pick a higher-quality model for other languages, or keep prompts simple
A model's Test button is stuck on the green tick	Test results are cached	Settings → AI Providers → Advanced → <b>Reset tests</b>
MCP server won't start	The port is already in use	Change <b>Port</b> in the MCP panel and enable again
Telegram bot won't start	Wrong or revoked bot token	Get a fresh token from @BotFather and save it again in Connect Mobile

## Getting help

- **Ask the app first** — type "how do I...?" in any chat; the assistant knows the studio and will walk you through it.
- **Email support:** [artflowly@gmail.com](mailto:artflowly@gmail.com). For bugs, attach the relevant log file from the logs folder above and mention your app version — that usually turns a back-and-forth into a one-reply fix.
- For licensing questions (moving machines, "license registered to another machine"), include your license key — you can copy it from Settings → Account.

Still stuck on setup? Re-read [Getting Started](#), the provider walkthrough in Settings & AI Providers, and how the agent behaves in [Chat & the Agent](#) — most first-week issues are covered there.