



AI Studio User Manual

AI video & image creation workspace · A complete guide for creators

Version 5.1 (English) · Site: dingshengshijie.studio

First published: July 2026 · Updated: 12 August 2026

Contents

1 Introduction & quick start

2 Sign up & sign in

2.1 Email sign-up and sign-in

2.2 Sign in with Google

3 Dashboard overview & navigation

4 How credits work

5 AI video generation

5.1 Text to Video

5.2 Image to Video

5.3 Reference to Video

5.4 Video model reference

5.5 Viewing & downloading results

6 AI image generation

6.1 Text to Image

6.2 Image to Image & the asset library

6.3 Image model reference

7 Smart edit

8 Storyboard (AI Agent)

9 Product Listing Set (AI Agent)

10 Image Replica (AI Agent)

11 Garment Pattern Variations (AI Agent)

12 Pattern Transfer (AI Agent)

13 Viral Video Remake (AI Agent)

14 Video Upscale (AI Agent)

15 Motion Transfer (AI Agent)

16 Video Edit (AI Agent)

17 Text to Speech (AI Agent)

18 My Creations

19 Asset library

20 Account settings

19.1 Profile

19.2 Security & phone binding

19.3 Credits & billing

21 FAQ

1. Introduction & quick start

AI Studio is an AI video and image generation workspace for creators — generation, editing, and asset management all in one place.

Here is what you can do with it:

- **AI video generation:** turn text, a single image, or multiple reference assets into short videos (Seedance, Kling, MiniMax H3, HappyHorse, Veo, Gemini and more).
- **AI image generation:** generate images from text, or repaint uploaded images (GPT Image 2, Nano Banana Pro, Seedream 5 Pro).
- **Smart edit:** annotate a generated image by hand (markers, arrows, text) and let AI apply the changes in place.
- **Storyboard:** upload product / character / scene material and get a multi-shot storyboard, then render it into a marketing video.
- **Product Listing Set:** upload product photos and let AI plan and batch-render a full set of e-commerce listing images (white background, lifestyle, feature callouts), with 20 built-in design style templates.
- **Image Replica:** supply style reference images plus your own product photos, and AI batch-produces listing images that follow the reference layout.
- **Garment Pattern Variations:** upload 1–6 reference images (garment shots, pattern swatches, worn photos — any mix, no grouping needed); AI freely auto-identifies each image's role plus an overall style and selling positioning, and one submit after you confirm the details drafts the plan and renders every flat-lay automatically — approve them, then batch-generate the model shots with one click.
- **Pattern Transfer:** upload the same style in several colours (1–12 photos) plus one pattern (direct upload, or free extraction from a viral photo); AI migrates the pattern pixel-faithfully onto every colour and renders a flat-lay plus a model shot for each — one model throughout, a different scene for every colour.
- **Text to Speech:** type a script, pick a curated voice and a model tier, and synthesize natural spoken audio — with pause markers, emotion tags (Expressive tier) and AI script polish.
- **Motion Transfer:** upload one character photo plus one motion video and the subject in the photo (a person or a pet) performs the video's motion — optionally in the video's own setting.
- **Video Edit:** modify / add / remove elements, edit sound, extend and splice an existing video — upload it or bring it in from any result card, then describe the change in one sentence.
- **Voice Clone:** upload 1–3 clean reference recordings and clone your own voice, then pick it right inside Text to Speech (visible to you only).
- **Viral Video Remake:** upload a viral reference video; AI analyses its camera work and voice-over rhythm, swaps in your own character / product / scene, re-renders segment by segment, and stitches the final cut.
- **Video Upscale:** turn low-resolution footage into a sharp 720p–4K version.

- **Work & asset management:** every result lands in **My Creations**, and everything you upload is kept in the **asset library** for reuse.

Quick start in three steps

1. Open the site, click **Sign in** in the top-right corner and register or sign in with an **email address** or **Google** (new accounts receive **100 credits**).
2. Enter the **Studio dashboard** and pick a tool from the left menu (Text to Video, Text to Image, Storyboard, ...).
3. Add reference materials by clicking **Add**, reusing from the **Library**, or simply **dragging files from your desktop onto the material area** (the zone highlights; mismatched types are skipped with a notice).
4. Write your prompt, choose a model and parameters, then click **Generate**. Credits are charged when you submit; the result appears in the preview panel and can be downloaded.

Tip: the top-right corner switches between **light / dark theme** and **English / 中文**. All screenshots in this manual use the light theme with the English interface.

Version updates: After a new release, a notification appears in the top-right corner. Save any unsubmitted edits before choosing **Refresh now**; choosing **Later** dismisses that release for the current tab session.

2. Sign up & sign in

The sign-in page directly shows the email/password form and keeps Google sign-in available.

2.1 Email sign-up and sign-in


A screenshot of the AI Studio sign-up page. The page has a light gray background. In the top left corner, there is a small square button with a left-pointing arrow. Centered on the page is a white rectangular form. At the top of the form is the AI Studio logo (a stylized 'A' with a circular arrow) and the text "Create your AI Studio account". Below this are three input fields: "Name" with the placeholder text "name", "Email" with the placeholder text "name@example.com", and "Password" with a masked password "*****" and an eye icon to toggle visibility. A green "Sign Up" button is positioned below the password field. Underneath the button is the text "Or continue with" followed by a "Sign In with Google" button featuring the Google logo. At the bottom of the form is a link that says "Already have an account? Sign in". Below the form, centered, is a small line of text: "By clicking continue, you agree to our [Terms of Service](#) and [Privacy Policy](#)".

Figure 2-1 Sign-up page

To register:

1. Click **Sign up** at the bottom of the sign-in page.
2. Fill in **name, email and password**, then submit.
3. A **verification email** is sent — open it and click the confirmation link.

Email verification is required: unverified accounts cannot enter the dashboard. If the email is missing, check your spam folder.




Figure 2-2 Email sign-in page

To sign in:

1. Open the sign-in page; the email/password form appears directly.
2. Enter your **email** and **password**; the eye icon toggles password visibility.
3. Click **Sign in**. Forgot your password? Use **Forgot password?** above the password field.

2.2 Sign in with Google

Click **Sign in with Google** and authorise the account in the Google consent screen.

Note: Google blocks sign-in inside embedded browsers (WeChat, QQ, TikTok and similar apps) — this is Google's own policy. The page will ask you to **open the link in a system browser**; copy it into Safari or Chrome, or use email sign-in instead.

3. Dashboard overview & navigation

After signing in you land on the **Studio dashboard**: tools on the left, tool cards in the middle, recent work at the bottom.

Compact navigation: On desktop, the narrower left navigation leaves more room for the creation workspace. Collapse it to an icon rail whenever you want to focus on your work. Mobile drawer and existing collapse behaviour are unchanged.

Two-level AI Agents nav: the **AI Agents** menu group nests three collapsible media categories — **Video Agents**, **Image Agents** and **Audio Agents** — all expanded by default; click the chevron next to a category title to collapse or expand it. The category containing the page you're currently on always stays expanded so you never lose the active link. **Audio Agents** holds **Text to Speech** and **Voice Clone** (see chapter 16).

The unified “New task” button: every creation page carries a **New task** button in its top-right corner, named identically across the whole product. It clears the current form and returns you to a clean slate; if the form holds unsubmitted work you are asked to confirm first. While you are viewing a past task, the same button simply leaves view mode.

AI Studio

Home

My Creations

Assets

AI Agent

VIDEO AGENTS

Storyboard

Viral Video Rem...

Video Upscale

Motion Transfer

Video Edit

IMAGE AGENTS

Product Listing ...

Image Replica

Garment Patter...

Pattern Transfer

AUDIO AGENTS

Text to Speech

Voice Clone

AI Video Generator

Reference to Video

Image to Video

Creator

e2e-manua...500

Studio dashboard

Studio dashboard

Choose an AI video, image, or utility workflow and keep every creation tied to your studio workspace.

AI Agent

VIDEO AGENTS

Storyboard

Turn product, character, and scene references into a multi-shot storyboard image, ready to...

Open tool →

Viral Video Remake

Extract the camera work and pacing from a viral video, then rebuild it with your own product...

Open tool →

Video Upscale

Upload a video, pick a target resolution, and let the AI super-resolution operator render a...

Open tool →

Motion Transfer

Transfer motion from a reference video onto your character photo

Open tool →

Video Edit

Modify, add/remove, re-voice, extend, splice — edit a video with one sentence

Open tool →

IMAGE AGENTS

Product Listing Set

Upload product photos and let AI plan and render a full e-commerce listing set: white...

Open tool →

Image Replica

Upload style reference photos and product photos — AI renders one on-brand product her...

Open tool →

Garment Pattern Variations

Upload reference images — AI identifies the garment, pattern and style and drafts a selling...

Open tool →

Pattern Transfer

Upload the same style in several colours plus one pattern (direct upload or free extraction...

Open tool →

AUDIO AGENTS

Text to Speech

Turn scripts into natural speech with pause markers, emotion tags and premium voices.

Open tool →

Voice Clone

Upload a clean voice sample and clone it — the voice becomes yours to use in Text to Speech.

Open tool →

AI Video Generator

Reference to Video

Turn reference images into video concepts that keep characters, products, and style consistent.

Open tool →

Image to Video

Transform static images into dynamic video concepts with motion direction and prompt...

Open tool →

Text to Video

Draft short-form video concepts from a text prompt before connecting the production API.

Open tool →

AI Image Generator

Image to Image

Upload a reference and reshape it into a cleaner style frame, thumbnail, or campaign visual.

Open tool →

Text to Image

Turn prompts into thumbnails, style frames, social visuals, and campaign references.

Open tool →

Recent creations

A lightweight placeholder for generated videos, images, and drafts.

0:00 / 0:00

0:00 / 0:00

View all →




Figure 3-1 Studio dashboard home

Left menu

Group	Category	Item	Purpose
—	—	Home / My Creations	Dashboard home; all generated work
AI Video	—	Reference to Video	Multiple reference images + prompt
		Image to Video	A single image as the first frame
		Text to Video	Prompt only
AI Image	—	Image to Image	Repaint an uploaded image
		Text to Image	Prompt only
AI Agents	Video Agents	Storyboard	Storyboard-to-video pipeline for product videos
		Viral Video Remake	Analyse a viral video and re-render it with your own subjects
		Video Upscale	Upscale footage to 720p–4K
		Motion Transfer	Transfer a reference video's motion onto a character photo
	Video Edit	Modify / add / remove elements, re-voice, extend or splice an existing video	
	Image Agents	Product Listing Set	Plan and batch-render a full listing image set
		Image Replica	Batch-replicate listing images from style references
		Garment Pattern Variations	Split one pattern into several colour and motif variants
		Pattern Transfer	Migrate one pattern onto several existing colour versions
	Audio Agents	Text to Speech	Turn a script into natural spoken audio
		Clone your own voice from reference audio	
	Settings	Profile	Name, avatar and other details
		Billing	Credit balance and transaction history

		Security / Files / Notifications	Password & phone binding / asset library / notifications
--	--	----------------------------------	--

Top and bottom areas

- **Account area (bottom left):** your name and **credit balance**. Click it to switch accounts or sign out.
- **Top right:** language (EN / 中文) and theme (light / dark) switches.
- **Recent work:** the dashboard home lists recent results; **View all** opens My Creations.

Tip: click **Open tool** on any card to enter that workspace — the left menu items do the same thing. The dashboard home's **AI Agents** section mirrors the same Video/Image grouping, with a one-line note at the end for the in-development **Audio Agents** category.

4. How credits work

Generating images and videos consumes credits. Knowing the rules helps you plan your work.

Where credits come from

- **Sign-up bonus:** every new account receives **100 credits** once.
- **Top-up purchase (launched 2026-08):** open **Pricing** from the top navbar or the sidebar. Four one-time tiers are available: ¥250 for 5,000 credits, ¥500 for 10,500 (+500 bonus), ¥2,500 for 55,000 (+5,000 bonus), ¥5,000 for 115,000 (+15,000 bonus). Payment is currently by **corporate bank transfer**: click "Top up now" and contact **路易电商 (Louis E-commerce)** as prompted, including your account email in the transfer note. **Top-up credits never expire**, and top-ups appear in the billing history.

How credits are spent

1. Every task shows its exact price on the submit button **before** you commit (for example "Generate · 200 credits").
2. Credits are **charged up front**. If your known balance is too low, the paid action stays clickable and shows a dedicated "Insufficient generation credits" message without starting a task. The server checks again when you submit, so the same message also covers a balance changed in another tab.
3. The price depends on the **model, resolution / quality, duration and number of reference images** — higher settings cost more. **The number on the button is always authoritative.**

Automatic refunds: if a task fails or times out, the charge is **refunded automatically** and the refund appears in your billing history.

Processing time and timeouts

Type	Maximum wait	Notes
Image	5 minutes	GPT Image 2 Flex is slower and gets 10 minutes
Video	~17 minutes	Tasks past the limit fail and are refunded
Viral Video Remake · analysis	8 minutes	The AI shot-analysis stage
Video Upscale	Minutes to hours	Scales with duration / frame rate / target resolution; the UI shows a live estimate

Note: you can leave the page while a task runs — it keeps going in the background and shows its latest state when you come back.

5. AI video generation

Three modes share one interface: Text to Video, Image to Video and Reference to Video. Switch between them from the left menu.

5.1 Text to Video

Generate a short video from a written description — ideal for drafting ideas quickly.




Figure 5-1 Text to Video workspace

1. Open **Text to Video** from the left menu.
2. **Model**: pick a video model from the dropdown (see 5.4). Each entry carries its brand icon.
3. **Resolution**: 480p / 720p and so on, depending on the model.
4. **Duration**: drag the slider or click a duration button.
5. **Aspect ratio**: 16:9, 9:16, 1:1 and others.
6. **Advanced**: expand for the audio toggle and, on some models, a random seed.
7. **Prompt**: describe the shot, camera movement, lighting and mood.
8. Check the price on **Generate · N credits** and submit.




Figure 5-2 Video model dropdown

5.2 Image to Video

Upload one image as the **first frame** and let AI bring it to life.




Figure 5-3 Image to Video workspace

1. Open **Image to Video**.
2. Click **Add** in the **first frame** slot to upload, or **Library** to reuse something you uploaded before.
3. Some models also accept a **last frame**; single-frame models only show the first slot.
4. Set model, resolution, duration and aspect ratio.
5. Describe the motion you want ("slow push-in", "the subject turns their head slightly").
6. Click **Generate**.

5.3 Reference to Video

Supply several reference images — plus reference video and audio — to keep a character, product or style consistent across a series.




Figure 5-4 Reference to Video workspace

1. Open **Reference to Video**.
2. Upload material into the **reference image / video / audio** groups. Each group header shows its quota (for example "0/9"), which varies by model.
3. Type **@** in the prompt to open the list of uploaded material; picking one inserts a **mention chip** that refers to that exact asset.
4. Set the model and parameters. For Seedance the service provider is scheduled automatically by the platform (Kie / DeepWL / X Deal, identical prices) — nothing to choose. Then click **Generate**.

Tip: @ mentions let you write things like "@image 1 wears the jacket from @image 3", so the model knows precisely which asset you mean.

Shot picker (pick reference segments by shot): uploading a reference video opens a selection flow right away: “Detecting shots...” first (skippable — just use the whole video), then choose the **whole video** (its row shows a three-frame strip, total length and shot count; over-long sources are auto-trimmed to the model’s first 15/30s, called out with a highlighted note) or **check several shots** — each selected shot becomes an **independent reference clip** (they are not spliced together), priced by its own seconds. Two caps guard the selection: the total selected duration cannot exceed the model’s per-clip cap (15/30s), and the count cannot exceed the remaining reference-video slots — you’ll be warned when no more can be added. Shots under 2s (the vendor minimum) are merged into their neighbors automatically. After confirming, the material group shows “N shots selected” with **Re-pick shots** and a X to **restore the whole video**; if detection fails or the video is a single shot, the flow continues with the whole video automatically.




Figure 5-5 Choosing the reference segment: whole video or several shots

Duration follows the reference: in Reference to Video the **Duration** slider automatically follows the total seconds of the reference video (or the selected shots) — no dragging needed; once you move it by hand it stops following. Re-picking shots or restoring the whole video recalculates from the new total.

View and “Adjust & retry” keep the shots: opening a past task shows the reference-video group as “Shot N · start-end s” entries exactly as submitted; after **Adjust & retry** the shots stay — they never fall back to the whole video. Shot detection does not re-run on its own: only **Re-pick shots** detects once, the dialog’s whole-video row shows the source’s real length, and the shot list starts **unchecked** (a clean start; confirming replaces the set, cancelling keeps the current shots).

5.3+ Prompt wizard (Seedance family)

Not sure how to write a video prompt? In **Reference to Video** with a Seedance-family model or MiniMax H3 selected, a **Prompt wizard** button appears beside the prompt box (for H3 the draft follows its official English six-section format) (not offered in Image/Text to Video for now). The wizard supports category switching (Apparel and General at launch, Apparel default; media analysis detects the category automatically and never overrides a manual pick) — detail-mining vocabulary and suggested actions adapt per category, with more categories to come. It is a single-page form: an editable analysis card plus six sections (asset roles → goal → scene & mood → subject & action → camera & rhythm → audio) on one screen, always visible and editable. A progress rail on top lights up completed sections, highlights the suggested next one, and jumps on click; fill them in and generate a structured prompt following the selected model's official prompt guide. Closing and reopening keeps your input; Reset clears it.

The screenshot shows the 'Prompt wizard' interface. At the top, there's a progress bar with six sections: 1. Assets (checked), 2. Goal (active), 3. Scene, 4. Action, 5. Camera, and 6. Audio. Below the progress bar, there's a 'Category' selector with 'Apparel' and 'General' options. The main content area is divided into sections. The first section, 'Assets', has a title 'What is each material for?' and a subtitle 'Assign each asset a duty, or keep auto-assign for AI to decide from the real image'. It shows two video assets, @Video1 and @Video2, each with a 'Motion reference' and 'Camera reference' button. The second section, 'Goal', has a title 'What is this video for?' and a subtitle 'Pick one goal tag — every later plan builds on it'. It lists various goal tags: Product promo, Product showcase, Scenario demo, Brand tone, New launch, Promo push, Unboxing review, Try-on showcase, Styling demo, Outfit transition, and Fabric & texture. The third section, 'Scene', has a title 'Where and what mood?'. At the bottom, there's a 'Generate - 10 credits' button.

Figure 5-6 Prompt wizard: asset roles and the one-click analysis card

1. Every step offers **suggested options plus free-form input**. When images are uploaded, the wizard first runs a free AI media analysis (product, selling points, audience); suggestions come from it, and the analysed product/selling points/audience prefill the editable card on top — one edit makes it the authoritative fact, and edits ride into the drafting call.
2. The asset-roles section assigns a purpose to each image/video: single-select for images (product reference, subject reference, ...; the default "Auto-assign" lets the AI pick roles from what each image actually shows — the reference images are attached to the drafting call). Videos are multi-select and come **pre-checked as "Motion reference + Camera reference"** (leave it and the source is mirrored shot by shot); a video marked style-reference only is not shot-mirrored — only its mood and pacing are

borrowed. You can also explicitly mark an asset as “Not used”. Every role actually takes effect in the final prompt: first frame = shot 1 starts from that image, last frame = the final shot settles on it, props are reproduced as pictured and actually used in the action, identity photos contribute face/hair/build only (wardrobe from product refs), scene images strictly define the environment.

3. **With a reference video: one-click analysis.** Once the roles are confirmed, a **“Analyze the reference video to draft the plan”** card sits between the assets and goal sections (it waits while the media analysis is still running, then becomes clickable). One click breaks the video down by the confirmed roles and **fills in at once** the goal tag, the scene & mood (the source environment is restored when no scene image is given; with a “Scene & style reference” image it describes that image instead), the per-shot action plan and the camera plan — every section stays editable, and you can regenerate. With several selected shots (up to 3), the wizard joins them in order into one reference timeline and analyses segment by segment with source labels such as “4-8s (video 2): ...”. The breakdown only restores the source; drafting then swaps the source’s people and objects for your subject and product — the source cast’s wardrobe is never copied, and every shot restates the product’s colours and structure. A static display shot of the source product is not mirrored — only its framing is borrowed and the displayed items are rewritten from your product images, without forcing a product image as the first frame (the shot picker flags such shots and suggests leaving them unchecked). With a reference video, assign at least one image as Product reference; the wizard reminds you when none is set. For brand lettering, upload a lettering image and assign it the Logo-reference role (the wizard prompts you automatically; lightweight models are unreliable at letters — use 2.5 or 2.0 Standard). Below the Audio section, a free-form Extra-requirements box takes anything else and the draft fulfils every item. **Product-swap mode:** when the source is detected as a same-category product demo, the shot picker opens in swap mode — pick one contiguous span (4-30s, Seedance 2.5) and the result keeps the source people, motion, camera and cuts while swapping the objects for your product; the wizard maps source objects one by one (replace / remove / keep) instead of reskinning them.
4. **Without a reference video: serial cascade.** Pick a goal, generate scene & tone ideas (matched to the goal and product), confirm the scene, then the AI drafts three complete timestamped action plans designed inside that scene (choose one, editable), and finally derives three camera schemes from the chosen action — each step builds on the previous confirmation. Assigning an image the “Scene & style reference” role makes the shooting environment strictly follow that image. The Scene / Action / Camera and Audio sections each carry a **notes-then-generate bar** (“Section notes” input with the generate button beside it; Enter also generates): state your constraints before generating (e.g. “must include an unboxing moment”) and the plans are designed around them from the start; edit the notes and hit Regenerate (or Enter) to redraft.
5. The **audio** section decides the clip’s sound design: pick **Voice-over** to generate three per-shot, time-segmented VO scripts (lines in curly braces, pacing matched to each shot’s length, product facts drawn only from your confirmed product info; the VO language is selectable — 中文, English, Español, 日本語, ไทย, Bahasa Melayu — defaulting to the site language); or **Music only** for three per-segment music plans; with a reference video you can also pick **Follow source audio**. Choose one of three and keep editing; the final prompt embeds the lines per shot in the official curly-brace syntax, or writes the

music into each shot's audio slot. Leave it unselected to add no audio constraint. Note: the sound design only takes effect when the Generate audio toggle is on.

6. Click **Generate prompt** (10 credits per generation, fully auto-refunded when you apply it and submit a video render within 24 hours — effectively free when used here; the questions, media analysis and video analysis are free). The result is editable, and **Apply** fills it into the prompt box — asset references become mention chips, and everything downstream works exactly like a hand-written prompt.
7. **Prompt health check and Fix all.** A "Prompt health check" panel appears under the draft: every asset referenced, no missing or "not used" assets cited, product shots spelling out colour/material, no "toddler/blogger"-style aliases replacing the defined subject, time segments covering the full duration, voice-over length fitting the duration, and no render parameters such as 4K/16:9 in the body. All-green means it passed; otherwise each finding is listed — edit the text directly (re-checked as you type) or click "Fix all · 10 credits" to regenerate once with the findings as revision requirements (rebated once applied and rendered). The server runs the same check plus one automatic revision pass during generation.
8. **Video feedback.** Under every successful result: "Did this video meet your expectations? 👍 👎 " — 👍 submits at once, 👎 lets you tick what went wrong (product looks wrong / person changes / wrong action, camera or scene / voice-over or audio / subtitles or watermark / quality or distortion / pacing) and add a note, editable any time. It never affects the task or credits; feedback is linked to the wizard recipe you used and drives the next round of wizard and prompt improvements.

Note: the wizard writes per model dialect — Seedance 2.0 uses "Shot 1/2/3" organisation with the official constraint tail; Seedance 2.5 only uses timestamps when you supplied a time allocation yourself (e.g. "0-3s close-up") or a reference video is attached. The fee appears in Billing as its own "Video prompt wizard" line.

5.4 Video model reference

Model	Strengths	Modes
Seedance 2.0 / Fast / Mini	Best all-round value, fast	Text / Image / Reference
Seedance 2.5	Newer Seedance generation, 480p/720p	Text / Image / Reference
Kling 3.0	Optional audio, supports element references	Text / Image / Reference
Kling 3.0 Turbo	The value Kling tier, faster turnaround	Text / Image
MiniMax H3	High-quality motion and multimodal reference control	Text / Image / Reference
Veo 3.1 Quality / Fast / Lite	High fidelity, from Google	Text / Image / Reference
Gemini Omni	Multimodal reference generation	Reference
HappyHorse 1.1	From Alibaba, with nine selectable aspect ratios	Text / Image / Reference

Seedance 2.5 parameters: resolution is 480p or 720p only (no 1080p/4K), with a 4–30 second duration and the same 9-image / 3-video / 3-audio reference budget as the 2.0 trio. Text to Video always generates audio for this model — the audio switch is locked on and cannot be turned off; Image to Video and Reference to Video keep the normal toggle.

Seedance 2.5 pricing: without a reference video, 28 credits/second at 480p or 63 at 720p. Attaching a reference video drops the rate to 17 or 38 credits/second, plus that reference video's own whole seconds. Since 2026-08-16 the Seedance family's provider (Kie / DeepWL / X Deal, identical prices) is scheduled automatically by the platform. Seedance 2.5 is available in the standard Video workspace, Storyboard and Viral Video Remake.

Kling 3.0 parameters: 720p / 1080p / 4K, ratios 16:9 / 9:16 / 1:1, and a 3–15 second duration (5 by default). Audio can be toggled. Image to Video takes a first frame plus an optional last frame. **Reference to Video uses named "elements"** rather than a flat reference pool: define up to 3 elements, give each an English name and upload 2–4 photos, then **reference every one of them in the prompt with @name** — an element you never mention has no effect, and submission is blocked with a reminder.

Kling 3.0 pricing: with audio off, 14 credits/second at 720p and 18 at 1080p; with audio on, 20 credits/second at 720p and 27 at 1080p; 4K is a flat 67 credits/second regardless of audio. **Kling 3.0 Turbo** is the value tier: 720p / 1080p only, no audio toggle, first frame only for Image to Video, at 18 credits/second (720p) and 22.5 (1080p). Both are available in the standard Video workspace only.

MiniMax H3 parameters: 768P by default with optional 2K, and a 4–15 second duration. Text to Video supports 21:9 / 16:9 / 4:3 / 1:1 / 3:4 / 9:16; Image to Video derives its ratio from the first frame (and optional last frame); Reference to Video also offers **Auto (adaptive)**. Reference mode accepts up to 9 images, 3 videos and 3 audio files. Audio cannot be the only input — include at least one image or video.

MiniMax H3 pricing: 16 credits/second at 768P and 26 credits/second at 2K (lowered 2026-08-17 with the upstream price cut). Output duration and every reference video's whole seconds are charged at the selected resolution. The first five reference images add no input fee; images 6–9 add 8 credits each; reference audio adds no input fee. The platform chooses and freezes the backend provider for each task as an operational detail, without changing your model controls or price. MiniMax H3 is **available in the standard Video workspace only** — it left Storyboard and Viral Video Remake in August 2026; use the Seedance family there instead.

Seedance provider: since 2026-08-16 the provider is scheduled automatically by the platform (Kie / DeepWL / X Deal, identical prices) — there is nothing to choose. Small capability differences between routes (an always-on audio route shows no audio switch; ratio options may vary slightly) render automatically in the form. A submitted task freezes its route; later scheduling changes never affect it.

Note: switching model filters the resolution, duration and aspect-ratio options down to what that model supports. Prices always follow the number on the submit button.

5.5 Viewing & downloading results

1. After submitting, the preview panel shows “Waiting / Generating” and then the finished video.
2. Click to play; use **Download** to save it locally.
3. Every task also appears in **My Creations** (chapter 17) where you can reopen or retry it.

Adjust while it runs: as soon as a task is submitted, **Adjust & retry** appears in the current workspace. It keeps the same materials and parameters in an editable new-task form while the original keeps running. The adjustment itself is free; credits are charged only when you submit the new task.

6. AI image generation

Two modes share one interface: Text to Image (prompt only) and Image to Image (repaint an uploaded reference).

6.1 Text to Image

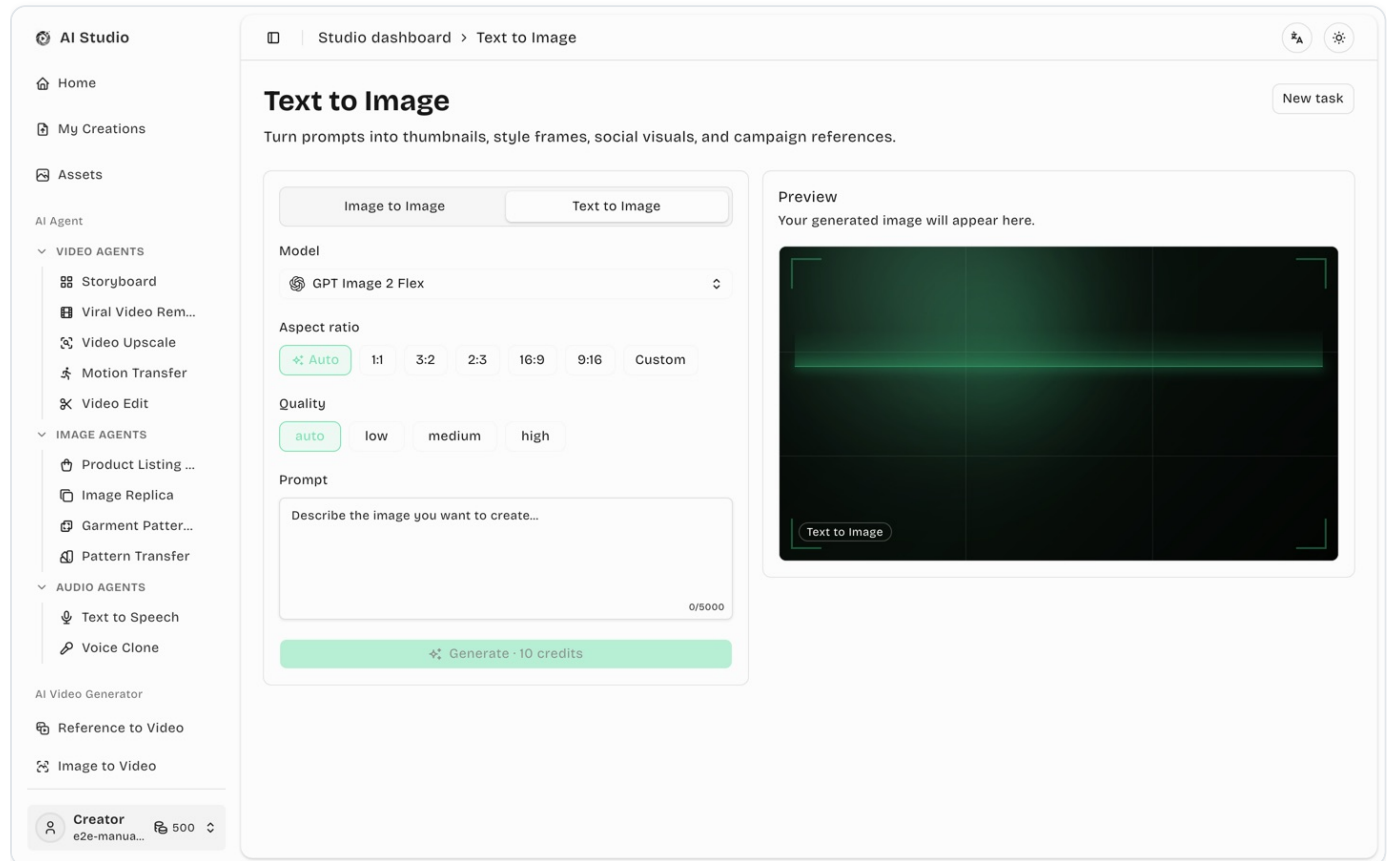


Figure 6-1 Text to Image workspace

1. Open **Text to Image** from the left menu.
2. **Model:** pick an image model (see 6.3).
3. **Aspect ratio:** Auto / 1:1 / 3:2 / 16:9 / 9:16 and more; some models add a **Custom** size.
4. **Quality / resolution:** GPT Image 2 Flex offers auto / low / medium / high quality tiers; other models offer 1K / 2K / 4K.
5. **Prompt:** describe the image.
6. Check **Generate • N credits** and submit.




Figure 6-2 Image model dropdown (each model carries its brand icon)

6.2 Image to Image & the asset library

Upload one or more reference images and let AI repaint them.




Figure 6-3 Image to Image workspace

1. Open **Image to Image**.
2. Click **Add** to upload, or **Library** to reuse an earlier asset.
3. Describe the change you want ("replace the background with a clean white studio").
4. Set the model and parameters, then click **Generate**.




Figure 6-4 Asset picker (multi-select; click a selected item again to deselect)

Working with the picker: use **All assets / folder** buttons at the top to narrow the list; scrolling to the bottom **loads more** assets. Clicking a thumbnail selects it and adds it to the upload area (selected items get a coloured outline), and clicking it again removes it without changing the source asset's folder.

Adjust while it runs: after an image task is submitted, **Adjust & retry** appears immediately. Click it to keep the same prompt, images and parameters in an editable new-task form while the original generation continues. No credits are charged until you submit the new task.

6.3 Image model reference

Model	Character	Precision options
GPT Image 2 Flex	Official route, quality tiers and custom sizes	Quality auto / low / medium / high
GPT Image 2 HD	High-definition route	1K / 2K / 4K
Nano Banana Pro	Gemini-family image model	1K / 2K / 4K
Seedream 5 Pro	ByteDance flagship, 7 fixed ratios	1K / 2K

7. Smart edit

Make targeted changes to a finished image: draw markers, arrows and notes on it, describe what should change, and AI repaints it **in place**. Works on Text to Image and Image to Image results, storyboard panels, and every card in Product Listing Set, Image Replica, Garment Pattern Variations and Pattern Transfer.

Where to find it

Once an image finishes, a **Smart edit** button appears next to **Download** in the preview panel. Storyboard boards have the same button in their action row, and listing / replica / garment-pattern / pattern-transfer cards carry their own entry point.




Figure 7-1 Image result · the “Smart edit” entry point

Editor tools




Figure 7-2 Smart edit editor

Tool	What it does
Marker	Click to drop a numbered pin, or drag a dashed box around the area to change (numbers increment automatically and can be referenced in the prompt)
Pen	Free drawing in a colour of your choice (red / yellow / green / blue / black / white)
Arrow / line	Drag to draw an arrow or line indicating direction
Text	Click to place a text box, type, then press Enter
Eraser	Click any annotation to remove it (object-level, not pixel-level)
Undo / redo	Step back and forward through your annotations
Zoom	Use “-/+” or scroll the mouse wheel over the canvas (50%–300%)

Steps

1. Click **Smart edit** to open the editor.
2. Annotate the areas you want changed.

3. **Describe the change** in the box at the bottom; you can reference marker numbers ("replace the red apple at marker 1 with a green one").
4. Pick the **model** and **resolution** for this edit (the original image's model is preselected) and check the price.
5. Click **Edit • N credits**. The image shows "Editing..." and the **new version replaces the original** when it finishes.

The replacement cannot be undone. A successful edit overwrites and destroys the original. If the edit fails, credits are refunded automatically and the original is untouched.

8. Storyboard (AI Agent)

A complete short-video pipeline for e-commerce: upload product / character / scene material → an AI scriptwriter writes the board prompt → an image model renders a multi-shot board → the AI writes the video prompt → a video model renders the final cut.




Figure 8-1 Storyboard workspace

Steps

1. **Upload material:** add references to the **product** (required), **character** and **scene** groups. The three groups **share one quota** (8 images on most models, 6 on Gemini Omni), so you can split it however you like.
2. **Video parameters:** target aspect ratio, voice-over language (English / 中文 / Español / 日本語 / Bahasa Melayu / ไทย), prompt model, board model, video model, Seedance service provider (shown only for Seedance), resolution, duration and shot count.

3. **Extra requirements (optional):** selling points, style, target audience — the AI gives these priority.
4. **Generate:** the scriptwriter writes the board prompt, then the board model renders the panels. The prompt box on the right fills in live and stays editable.
5. **Video prompt:** once the board is ready, click **AI Assist** in the video-prompt block to have the AI write the video prompt from the board.
6. **Render:** confirm and submit — the video model renders the final cut.

Adjust without stopping the Storyboard: after submission, the left-side materials and parameters stay frozen until you click **Adjust & retry**. That action detaches an editable copy for a new Storyboard; it does not re-render, stop or change the old board, which continues in the background.

Parameters

Parameter	Description
Board model	GPT Image 2 Flex / HD
Video model	Seedance 2.0 / Fast / Mini, Seedance 2.5, Gemini Omni
Seedance provider	Scheduled automatically by the platform (Kie / DeepWL / X Deal, identical prices); frozen with the Storyboard task and reused when the final video is rendered
Voice-over language	Controls spoken lines and subtitles; text rendered inside the image is always English (image models are unreliable with non-Latin scripts)
Duration / shot count	Duration suggests a shot count (roughly one shot per 1.7 s); the count stays editable

Seedance 2.5 in Storyboard: appears right after Seedance 2.0 Mini, shares the same 8-image reference pool, and offers 480p/720p only (no 1080p/4K). Its provider is scheduled automatically by the platform, and the rendered video is priced at the no-reference-video rate (chapter 5.4).

Combine with Smart edit: if a board panel needs a tweak, use **Smart edit** on the board (chapter 7) before rendering the video.

Note: each step (script, board, video) is charged separately and the button always shows the current price. A finished storyboard appears as a single combined card in My Creations.

9. Product Listing Set (AI Agent)

Upload product photos and let AI plan and batch-render a complete set of listing images: white-background hero shots, lifestyle scenes, feature callouts and more. Built for Amazon, TikTok Shop, AliExpress and similar marketplaces.

The screenshot shows the 'Product Listing Set' workspace in the AI Studio interface. On the left is a sidebar with navigation options: Home, My Creations, Assets, and AI Agent. Under 'AI Agent', there are sections for VIDEO AGENTS (Storyboard, Viral Video Rem..., Video Upscale, Motion Transfer, Video Edit) and IMAGE AGENTS (Product Listing..., Image Replica, Garment Patter..., Pattern Transfer). Below these are AUDIO AGENTS (Text to Speech, Voice Clone) and AI Video Generator (Reference to Video, Image to Video). At the bottom of the sidebar is a 'Creator' profile for 'e2e-manua...' with 500 credits. The main workspace has a breadcrumb 'Studio dashboard > Product Listing Set' and a 'New task' button. The title 'Product Listing Set' is followed by a description: 'Upload product photos and let AI plan and render a full e-commerce listing set: white-background mains, lifestyle scenes, selling-point shots and more.' The workspace is divided into two main panels. The left panel contains: 'Product photos' (0/8) with an 'Add' button and a 'Library' link; 'Target platform' (TikTok Shop) and 'Target market' (United States) dropdowns; 'Copy language' (English) dropdown; 'Selling points' (AI fill) with a text area containing a skeleton: 'Product name:', 'Key selling points:', 'Target audience:', 'Usage scenarios:', 'Dimensions & specs:'; 'Design style' with 'AI pick' (selected), 'Style templates', and 'Custom' buttons, and an 'Analyze & recommend a style' button; and 'Set structure' with 'Smart match' (AI analyzes the product and picks the best mix) and 'Custom' (Adjust each category (4-12 images total, at least 1 white-background)) options. A note says 'Upload at least one product photo first' and a green button at the bottom says 'Generate plan - 2 credits'. The right panel is titled 'Set plan' and contains the instruction: 'Click "Generate plan" — AI drafts each image's role and prompt; review, then render in one batch.'

Figure 9-1 Product Listing Set workspace

Steps

1. **Product photos:** up to 8, uploaded or picked from the library.
2. **Platform / market / copy language:** choose the marketplace (TikTok Shop, Amazon, AliExpress, own store, ...), the target market and the language of on-image copy. The AI plans against each platform's real image conventions (Amazon's pure-white background and ~85% frame fill, for instance).
3. **Selling points:** fill in the five-line skeleton (product name / key selling points / target audience / intended scenes / size & specs), or click **AI generate** (2 credits) to have the AI draft it from your photos and then edit it.
4. **Design style** — pick one of three routes:

- **AI pick:** click **AI style analysis** and the AI proposes a style from your photos (2 credits); the text stays editable.
 - **Style templates:** choose one of 20 curated presets (Figure 9-2) across four groups — premium, natural, vibrant and category-specific — such as “Premium editorial”, “Creamy pastel”, “Dark luxury” or “UGC casual”.
 - **Custom brief:** type your own requirements: palette, layout, creative direction.
5. **Set structure:** **Smart match** lets the AI decide how many images of each type; **Custom** lets you set the counts yourself (4–12 images in total, at least one white-background shot).
6. **Generate the plan** (2 credits): the AI plans a category, title and English render prompt for every image and lists them as cards. Cards can be selected, deselected and edited.
7. **Batch render:** choose the image model (GPT Image 2 / Seedream 5 Pro / Nano Banana Pro), quality tier and aspect ratio on the right, then click **Generate selected · N credits**. Each image is charged and rendered independently and appears as soon as it is ready.

New task

Upload product photos and let AI plan and render a full e-commerce listing set: white-background mains, lifestyle scenes, selling-point shots and more.

Product name:
Key selling points:
Target audience:
Usage scenarios:
Dimensions & specs:

Mar...
Bright
marble
glass
and
champ
gold,
airy
high-
key

Vint...
Film-
grain
warmth
faded
sun-
washed
nostalg

- Natural & fresh**
-  **Sca...**
Functi
calm,
pale
oak,
bri
diffu
dayli
-  **Web...**
Quiet
imper
earthe
tones,
handc
texture
-  **Gol...**
Honey
-toned
wood,
golden
kitcher
light,
cozy
homeli

Bot...
Sage
greene
dapple
dayligr
organit
vitality

Coa...
Sunlit
sands,
sea-
glass
blue,
palm
shadow

Vibrant & trendy

- ✦ **Vibr...**
Bold
color
block
pop-
art
energy
shadow
light
- 🏠 **Urb...**
Concre
harsh
sun,
neon
accent
candid
attitud
- ♡ **Cre...**
Dream
high-
key
softne
satin
and
pearls

Fest...
Bokah
lights,
red-
and-
gold
celebr
glow

UGC...
Authen
smartp
realiser
everyd
setting

Set plan

Click "Generate plan" — AI drafts each image's role and prompt; review, then render in one batch.




Figure 9-2 Design style · Style templates (20 presets with thematic icons)

Per-image actions

Action	Description
Smart edit	Opens the annotation editor (chapter 7); the edit replaces that image
Zoom / download	View the full-size image or save it
Rewrite & regenerate	Edit that card's prompt and render it again (normal charge; the old image is retired)

Tip: planning costs only 2 credits — review the whole set and its prompts first, then batch-render. Not happy? **Adjust & retry** reloads your inputs into a fresh task.

Note: prompts sent to image models are always in English (models render other scripts unreliably). If you edit a prompt in Chinese, it is translated automatically before rendering. The whole set shows as one card in My Creations, and all its charges are grouped into a single billing entry.

10. Image Replica (AI Agent)

Found a listing image whose style you love? Upload it as a style reference, add your own product photos, and AI batch-produces listing images in that style — one output per reference image.




Figure 10-1 Image Replica workspace

Steps

1. **Style references:** up to 12. Every reference produces one output in its style.
2. **Product photos:** your own product. The cap adapts to the model and the number of character photos and is shown live in the UI. Only **your** product is ever painted — never the one in the reference image.
3. **Character photos (optional):** up to 2. Only the face, hair and build are taken from them — **the garment always comes from your product photos**, which is what makes this work for apparel.

4. **Market / copy language:** drives on-image copy and the visual sensibility.
5. **Fidelity** — pick one: **Layout reference** keeps the layout, background structure and subject relationships but designs the palette around your product; **Strict replica** reproduces composition, layout, palette and detailing, swapping only the product and the selling points.
6. **Key selling points (required):** type them or click **AI generate** (2 credits).
7. **Unified brief (optional):** constraints for the whole set, e.g. "all copy in English", "keep the model's pose unchanged". Each requirement is written into every card's prompt: numbers (sizes / thickness / angles) verbatim, style names and structural selling words in explicit English; numbers are also machine-checked and auto-repaired once if missing.
8. **Generate the plan** (2 credits): the AI writes one render prompt per reference; cards are selectable and editable.
9. **Batch render:** choose model, quality tier and ratio on the right, then click **Generate selected**. Per-card actions (smart edit / download / rewrite & regenerate) match Product Listing Set.

Copyright reminder: reference images are only used to convey composition and style. Do not commercialise someone else's assets — the product and people in your outputs come from the material you upload.

11. Garment Pattern Variations (AI Agent)

Upload 1–6 garment, pattern, worn or detail references into a local draft. Click **Start analysis (free)** to create a durable task, review the saved AI interpretation, then click **Confirm generation** to pay for the plan and flat-lays. Once those finish, batch-generate the model shots with one click.




Figure 11-1 Garment Pattern Variations workspace

Steps

1. **Build the local draft (1–6 images, any mix, no grouping needed):** garment shot (the source of truth for cut and silhouette), pattern swatch (style DNA), worn photo (cut plus pattern), or detail reference (context only). Add, remove and reorder freely. Uploading alone makes no AI call and creates no task.
2. **Start analysis (free):** this deliberate click creates the task first, then analyses every source in the durable background workflow. The task immediately appears in **My Creations**; you can leave, close the browser or switch devices without interrupting it. No credits or billing entry are created.
3. **Review the saved analysis:** each source is tagged as garment, pattern, worn or detail and receives a summary. A degraded source is marked **Needs manual confirmation**. Source images and order stay frozen, while the role, summary and all analysis fields remain editable.
4. **Category / target wearer / base colour:** AI fills these in from the uploads; all editable. Base colour is optional free text (e.g. "black", "off-white").
5. **Style DNA / selling positioning / suggested scene / suggested model persona:** four AI-drafted notes (colour & motif, target audience, shoot setting, model styling) — edit any of them directly.
6. **Variant count:** 10–20 variants, default 10, adjustable with the stepper.
7. **Model / tier / both ratios:** Garment Pattern Variations supports **GPT Image 2 HD (default)**, Nano Banana Pro, and Seedream 5 Pro; choose the quality tier plus a separate **flat-lay ratio** and **model shot ratio** (both default to 1:1). These settings freeze only when paid generation is confirmed.
8. **Confirm generation** (2 credits for the plan, plus every variant's flat-lay priced by the chosen model — the exact total is shown on the button): this is the first paid action. AI drafts the plan and renders every flat-lay immediately.
9. **English Logo and typography:** each new variant receives a distinct fictional English brand name, optional slogan and matching typography/layout. Prompt regeneration preserves its frozen design; model shots copy the latest flat-lay's exact visible spelling and layout.
10. **Results (three-column workflow):** each variant is one row, showing **Flat-lay / Model shot / Pattern vector file** in generation order — three columns side by side on desktop, the same order stacked vertically on phones. Each card header has a checkbox — checkable once that card's own flat-lay succeeded, checked by default.
11. **Multi-select batch generation:** once every flat-lay has finished (success or fail) and at least one succeeded, a toolbar appears at the top of the results — a **select-all** checkbox (toggles between select-all and clear-all) plus two independent buttons, **Generate selected model shots** and **Generate selected pattern vectors**, each priced live from the cards that are both checked and don't already have that render. Uncheck a card to leave it out of this batch; a card that already has a model shot or pattern vector is never re-charged or re-rendered by either button, even while still checked — the two buttons' eligibility is independent. You can also click **Generate model shot** on a single card.
12. **Pattern vector file (print-ready, the third result column):** once a variant's flat-lay renders successfully, that column shows a **Generate pattern vector** button (16 credits, regardless of whether a model shot exists) — AI extracts only the printed graphic from that flat-lay and redraws it as a standalone file: white background made transparent, roughly 4096px wide (about 34.7cm at 300dpi,

inside the 32–36cm print-size range). The preview renders over a checkerboard backdrop so the transparency is visible; **Download** saves a PNG with a `-300dpi` filename suffix. A failed render can be retried in place at the quoted price.

13. **Unified edit and regeneration:** the separate regeneration buttons below each image are removed. Use **Edit & regenerate** at the top-right of the prompt area to edit the flat-lay prompt, model persona, shooting scene and lighting. Changing only model settings regenerates only the model shot. Changing the flat-lay prompt first regenerates the flat-lay, then automatically renders a new model shot from it. The submit button shows the actual scope and credits; the durable two-stage job continues if you leave the page.
14. **Other per-card actions:** Smart edit (chapter 7 — annotate and overwrite in place), zoom and download.

Two safe stages: **Start analysis** is free and durable; **Confirm generation** is the explicit billing boundary. **Adjust & retry** is available as soon as the first task exists, including while analysis is running. It opens a prefilled new draft without changing or stopping the original task.

Commercial-use reminder: Logo names are creative examples, not trademark-cleared marks. Search target-market availability before commercial use.

Note: free analysis creates no ledger entry. Paid confirmation charges the flat 2-credit plan fee plus every flat-lay; model shots are charged separately. The whole task shows as one card in My Creations, with analysis / awaiting-confirmation copy before real child renders exist and live child progress afterwards. Historical Flex tasks remain viewable; **Adjust & retry** switches a fresh form to GPT Image 2 HD.

12. Pattern Transfer (AI Agent)

The same style in several colours shares one pattern: upload one garment photo per colour (1–12) plus a pattern (direct upload, or free extraction from a viral photo). Click **Start analysis (free)** to create a durable task; review the colour list and pattern preview, then click **Confirm generation** to pay for the plan — the pattern is migrated pixel-faithfully onto every colour and every flat-lay renders automatically. Once a flat-lay is approved, generate its model shot (one model throughout, a different scene for every colour).




Figure 12-1 Pattern Transfer workspace

Steps

1. **Same-style colour garments (1–12 photos):** one photo per colour version of the same style; add, remove and reorder freely.
2. **Pattern source:** choose either **Upload pattern** (a standalone swatch or a close-up of the print), or **Extract from a viral photo:** upload a viral garment photo carrying the target pattern and the system immediately starts a free background extraction — the card shows **Extracting...**, then switches to **Extraction ready** with a preview once done. Click **Re-extract** for another attempt, or switch back to **Upload pattern** to use a different image instead. Re-using the same viral photo hits a cache — instant and free.
3. **Model reference photos (optional, 0–2):** upload to anchor that exact model's identity throughout; leave empty and AI designs one model matching the garment style and target market, kept the same across every colour.
4. **Pattern colour strategy: Preserve original colours** (default — pixel-faithful, no recolouring) or **Adaptive colour** (the motif and layout stay identical; colours are nudged to complement each garment colour).
5. **Flat-lay style / model / tier / target market / both ratios:** Pattern Transfer supports **GPT Image 2 HD (default)**, Nano Banana Pro, and Seedream 5 Pro, with a separate **flat-lay ratio** and **model shot ratio** (both default to 1:1). These settings freeze only when paid generation is confirmed.
6. **Start analysis (free):** this deliberate click creates the task first, then reads every colour's name (English + local), category, gender, and checks whether every garment photo is genuinely the same style. The task immediately appears in **My Creations**; leaving the page never interrupts it. No credits are charged at this stage.
7. **Review analysis_ready:** the colour list (a thumbnail plus editable English/local names for each colour), the pattern preview, and editable category/gender. A garment photo that looks like a different silhouette surfaces a **non-blocking consistency notice** — continuing is your call.
8. **Confirm generation** (2 credits for the plan, plus every colour's flat-lay priced by the chosen model — the exact total is shown on the button): this is the first paid action. AI migrates the pattern onto every colour and renders every flat-lay immediately.
9. **Results (colour card grid):** one card per colour, headed by a colour-name badge, with **Flat-lay / Model shot / Pattern vector file** below it in generation order — three columns side by side on desktop, the same order stacked vertically on phones. Each card header has a checkbox — checkable once that colour's own flat-lay succeeded, checked by default. Once a colour's flat-lay is a terminal success with an archived thumbnail, the model-shot column's **Generate model shot** button unlocks — it renders that colour's own scene and pose with the same model identity, priced independently.
10. **Multi-select batch generation:** once every flat-lay has finished (success or fail) and at least one succeeded, a toolbar appears at the top of the results — a **select-all** checkbox (toggles between select-all and clear-all) plus two independent buttons, **Generate selected model shots** and **Generate selected pattern vectors**, each priced live from the colours that are both checked and don't already have that render. Uncheck a colour to leave it out of this batch; a colour that already has a model shot or pattern vector is never re-charged or re-rendered by either button, even while still checked.

1. **Pattern vector file (print-ready, the third result column):** once a colour's flat-lay renders successfully, that column shows a **Generate pattern vector** button (16 credits, regardless of whether a model shot exists) — AI extracts only the printed graphic from that flat-lay and redraws it as a standalone file: white background made transparent, roughly 4096px wide (about 34.7cm at 300dpi, inside the 32–36cm print-size range). The preview renders over a checkerboard backdrop so the transparency is visible; **Download** saves a PNG with a `-300dpi` filename suffix. A failed render can be retried in place at the quoted price.
2. **Unified edit and regeneration:** use **Edit & regenerate** at the top-right of a card's prompt area to edit the flat-lay prompt, scene and pose together. Changing only scene/pose regenerates only the model shot. Changing the flat-lay prompt first regenerates the flat-lay, then automatically renders a new model shot from it.
3. **Other per-card actions:** Smart edit (chapter 7 — annotate and overwrite in place), zoom and download.

Two safe stages: pattern extraction and **Start analysis** are both free and durable; **Confirm generation** is the explicit billing boundary. **Adjust & retry** is available as soon as a task exists — it opens a prefilled new draft without changing or stopping the original.

Note: pattern extraction is always free; its result is content-addressed to the source photo, so re-extracting the same viral photo costs nothing extra. Paid confirmation charges the flat 2-credit plan fee plus every colour's flat-lay; model shots are charged separately, and character reference photos never count toward the flat-lay's own input images. The whole task shows as one card in My Creations, grouped in billing as one "Pattern Transfer" entry (tracked and filtered separately from Garment Pattern Variations).

13. Viral Video Remake (AI Agent)

Upload a viral reference video. AI analyses its camera work, pacing and voice-over, swaps in your own character / product / scene, re-renders it segment by segment, and stitches the final cut with one click. Eligible test accounts can also compare the stable Original-video path with an experimental Depth-reference path.




Figure 13-1 Viral Video Remake workspace

Steps

1. **Reference video:** MP4/MOV, 3–60 seconds, up to 100 MB. After upload the system **detects the voice-over for free** and transcribes it word for word. The same analysis always creates a publish title plus 3–

5 directly copyable `#tags` . Pasting the original title and tags is optional and only helps match their structure and tone.

2. **Reference mode (controlled test accounts only): Original video · Stable** is the default and gives the video model the trimmed source clip. **Depth reference · Experimental** converts each segment into a depth-motion video before rendering, preserving body/object geometry and motion while suppressing the source clip's colour, texture and identity. The selector is hidden when the experiment is not enabled for your account; saved historical tasks still show the mode they used.
3. **Replacement images:** a product image is required; character / product / scene share 8 slots. Character photos receive a free automatic **white-background pass** ("White ✓"). In Original mode an empty group can retain or infer an equivalent from the source, and an empty scene group may use the cleanest source frame. In Depth mode source appearance is deliberately unavailable: only your uploaded replacement images anchor identity, product and scene, so upload every appearance you need.
4. **Selling points and unified requirements:** add factual product selling points (required; AI fill is available) and optional whole-video requirements for camera, styling or messaging. These inputs are shared by every segment.
5. **Voice-over** — pick one: **AI rewrite** writes a new script from the visible product facts and your selling points; **Keep original (editable)** reuses the detected transcript after your edits. Keeping an empty transcript preserves the reference audio; rewrite can add narration to a silent reference. **Subtitles** are off by default, and the language picker covers English, Chinese, Spanish, Japanese, Malay and Thai.
6. **Choose video-generation parameters:** before pressing Analyse, choose the model, resolution and ratio. The whole Seedance family's provider is scheduled automatically by the platform; **Seedance 2.5** appears right after Seedance Mini and offers 480p / 720p only. This complete parameter set freezes when analysis is submitted, and later top-ups and re-renders reuse it.
7. **Analyse shots** (2 credits): the AI returns a segment-by-segment script — each 4–15 seconds with camera movement, action and its voice-over line. The same analysis call writes directly for the final model, with no extra request or fee: Seedance uses its media-role syntax. Voice-over lines are individually editable.
8. **Render:** select the segments to render (or select all), then click **Generate**. Segments render as a **sequential chain**: the previous segment's final frame is supplied to the next one.
9. **Compose the full video:** once every segment succeeds, **Compose full video** appears (free) and stitches the cut. Re-rendering a segment lets you recompose.

Depth render-requirements preview: unrendered Depth segments show their complete render requirements in the current site language and update live as you add directions, so you can review them before submitting. The generation request still uses a model-optimised English internal prompt; spoken lines are not translated to English and always remain in the voice-over language you selected.

Seedance 2.5 in Viral Video Remake: every segment render always attaches a materialized reference clip, so pricing always uses the cheaper with-reference-video rate from chapter 5.4 (never the plain flat rate). The same automatic recovery that retries a moderation-rejected or copyright-flagged segment (prompt rewrite, then muted clip + generated music) applies to Seedance 2.5 segments too.

Staying in control

Action	Description
Stop generating	Stop at any time: unfinished segments are marked failed and refunded, finished ones are kept and can be re-rendered individually later
Retry failed / rewrite & regenerate	Re-render a single segment (normal charge); it chains from the nearest earlier successful segment
Adjust & retry	Loads the model, provider, resolution, ratio and other inputs from this submission (or recovers them from a historical first child) into an editable brand-new task

Choosing a mode: use Original for the established path and when the source scene/appearance is useful. If Depth is visible on your test account, use it to evaluate motion transfer with stronger appearance isolation. Depth preparation can add processing time; if it cannot be prepared, the segment fails safely and is refunded — switch back to Original and retry.

Billing: mode selection adds no separate fee. Analysis costs 2 credits; every segment is charged separately by model / resolution / duration; a failed predecessor cancels and refunds the segments that follow; composing the final cut is free. All charges for one task are grouped into a single billing entry.

Content moderation: occasionally a reference frame is rejected by the vendor's moderation (the system first retries with a sanitised prompt and muted audio). If a segment reports that its reference frame was blocked, swap the reference video or adjust that segment.

14. Video Upscale (AI Agent)

Turn soft, low-resolution footage into a sharp 720p / 1080p / 2K / 4K version — an AI super-resolution operator rebuilds detail frame by frame. The empty workspace shows a real before/after demo: drag the divider to compare.




Figure 14-1 Video Upscale workspace (slider comparison demo on the right)

Steps

1. **Pick the upscale model** (one of two): **Volcengine Video Upscale** — choose a target resolution from 720p to 4K, any source size; or **Aliyun Video Upscale** — a fixed 2× upscale with quality enhancement, requiring a source above 360×360 and below 1920×1080, up to 1 GB, priced by duration at a flat 2 credits/second.
2. **Upload the source video**: up to 600 seconds (10 minutes); the library works here too.
3. **Pick the target resolution** (Volcengine mode): 720p / 1080p / 2K / 4K. Portrait footage is measured on its short edge (portrait 1080p outputs 1080×1920). Aliyun mode skips this step — output is always 2× the source.
4. **Check the quote**: the button shows the exact price — it scales with duration, frame rate and target resolution — alongside an **estimated processing time**. Duration and frame rate are probed server-side, which is what the charge is based on.
5. Click **Upscale**. The task runs in the background with a live “time remaining” estimate; you can leave the page.

6. When it finishes, the result appears on the right — click **Download** to save it.

Price guide: 720p is the cheapest tier; 1080p costs about 3×, 2K about 5× and 4K about 11×, and footage above 30 fps carries a frame-rate multiplier. A one-minute 30 fps clip upscaled to 720p costs roughly 53 credits.

Note: long footage at a high target resolution can take hours; the system waits up to 24 hours and refunds automatically on timeout or failure. Deleting a running task refunds it too.

15. Motion Transfer (AI Agent)

Transfer the motion from a reference video onto your own character photo: upload one character photo plus one motion video, and AI makes the subject in the photo perform the motion from the video. The result runs as long as the motion video and keeps its audio.




Figure 15-1 Motion Transfer workspace (materials and parameters on the left, preview on the right)

Choose a model first

The page offers two complementary models; the choice changes both the available parameters and the price:

	Wan 2.2 Animate	Kling 3.0 Motion Control (default)
Subject	Not limited to people — pets and toys work too	Real people only (the provider runs person detection and rejects pets and objects outright; such failures cost no credits)
Resolution	480p / 720p	720p / 1080p
Video length	5–120 seconds	3–30 seconds
Price	8 credits/s at 480p, 16 credits/s at 720p (billed as 5 seconds minimum)	20 credits/s at 720p, 27 credits/s at 1080p
Keeping the video's scene	Native — pick "Into the video"	Needs our "Match the video's scene" preprocessing

In short: **pick Kling when you need 1080p, otherwise pick Wan** — it is cheaper, takes longer videos, and is the only one that handles non-human subjects.

Steps

1. **Upload the character photo** (1 image): it must clearly show the subject's head and torso. JPG / PNG, up to 10MB, min 340px on the short side, aspect ratio between 2:5 and 5:2.
2. **Upload the motion reference video** (1 clip): MP4 / MOV, up to 100MB, min 340px, and its length must fall inside the selected model's range (see the table above).
3. **Transfer mode (Wan only): Into the video** (default) swaps the subject in the video for your character and **keeps the video's own scene and camera**; **Onto the photo** applies the motion to your character photo and keeps that photo's own background.
4. **Match the video's scene** (Kling, and Wan's "Animate the photo"): before submitting, the system composites both the setting **and the subject's pose** from the motion video's first frame into your character photo for free — the result then shares the video's background and starts from a matching pose, so the motion picks up naturally. Turn it off to keep the photo's own background and pose. Compositing takes about a minute, during which Generate stays disabled. **Once it finishes, the character photo is replaced by the composite**, so you can confirm exactly what will be rendered before submitting. If you don't like the result, add a line of extra direction in the box below (for example "stand closer to the camera" or "no glasses") and press **Regenerate the scene image** — the extra direction applies only when you press that button. Wan's "Into the video" already uses the video's setting natively, so the switch is hidden there.
5. **Pick a resolution**: follows the model (Wan 480p/720p, Kling 720p/1080p).
6. **Pick the character orientation (Kling only)**: "Match video" reproduces complex motion more faithfully; "Match photo" supports more camera movement.

7. **Prompt (optional)**: can be left empty. Add one only when you want to emphasise a specific detail of the motion.
8. Click **Generate**. The button shows the exact credit cost.

Pricing

Both models are charged by the **length of the motion video** (the result runs exactly as long; partial seconds are not counted). Wan 2.2 Animate: 8 credits/second at 480p and 16 at 720p, **billed as at least 5 seconds**. Kling 3.0 Motion Control: 20 credits/second at 720p and 27 at 1080p. An 8-second motion video at 720p therefore costs 128 credits on Wan and 160 on Kling. The server re-reads the video's real duration before charging, so the final amount always follows the true length.

Shooting tips: a front-facing, evenly lit photo with the whole subject in frame works best; the clearer the subject and the more definite the movement in the motion video, the more accurate the transfer.

About the scene: apart from Wan's "Into the video", the result's background comes from the **character photo**. That is what the "Match the video's scene" switch is for — it composites the video's setting and opening pose into the photo first, so the result looks like it was shot in the video's environment and the motion picks up cleanly. If compositing fails, rendering continues with the photo's original background instead of aborting, and the workspace says so.

16. Video Edit (AI Agent)

Change a finished video with one sentence: modify / add / remove visual elements, edit the sound, extend it, or splice 2–3 clips with seamless transitions. Upload a video directly, or jump in from the “Video Edit” button on any AI result card.




Figure 16-1 Video Edit workspace (canvas + filmstrip + instruction bar, parameters on the right)

Six task types

Task	What it does
Modify element	Swap something in the frame for something else (color / material / subject / background...)
Add element	Insert a new element at a given position and moment
Remove element	Erase watermarks, passers-by, clutter — everything unmentioned stays as-is
Sound edit	Remove background music / change the voice / add sound effects; the picture stays unchanged
Extend	Continue the video with seamless style and narrative
Transition splice	Bridge 2–3 clips into one seamless video

Selecting a task prefills the officially recommended sentence template — just replace the angle-bracket placeholders. Text you have edited yourself is never overwritten when switching tasks.

Steps

1. **Add the source video:** upload, pick from the library, or bring in a result. Per-clip length caps follow the model: 30s on Seedance 2.5, 15s on the 2.0 family; splice accepts 2–3 clips.
2. **Browse the filmstrip:** a CapCut-style timeline under the player (evenly sampled to fill the width, no scrolling) — tap or drag to move the playhead.
3. **Scope the edit range (Seedance 2.5):** press and drag directly across the strip to select the range (the outside dims; edge handles fine-tune, a single click still seeks); “within Ns–Ms” is written into the instruction automatically so the edit only touches that window. Drag back to full width to clear it.
4. **Pick a task and finish the instruction:** replace the placeholders, or write your own.
5. **(Optional) auxiliary reference images:** for combo tasks like swapping products or prints, attach a product photo and reference it as “referencing Image 1’s product...”.
6. **Choose model and resolution:** all four Seedance models. **Seedance 2.5 edit tasks automatically match the source’s length and aspect ratio** (locked by the official spec); the 2.0 family and extend tasks set an explicit output length.
7. Click **Generate** (the button quotes the price live). When done, use the **Original / Edited** toggle on the player for instant A/B comparison; continue editing the result for chained iterations, or send it to **Video Upscale**.

Pricing

The existing reference-video rate applies: **(output seconds + floored source seconds) × credits per second**. Example: editing a 13s source on Seedance 2.5 at 480p = $17 \times (13+13) = 442$ credits; the same task on 2.0 Mini at 480p with a 5s output = $2.4 \times (5+13) \approx 43.2$ credits. Splice output is the sum of the clips plus ≈ 1 s per transition. Failures refund automatically.

Three tips for a good edit instruction: ① open with a clear verb (edit / modify / remove / add); ② describe the target with recognizable traits (color, position, material); ③ keep the closing “everything else stays unchanged” — the official guide stresses that stating what to keep markedly improves fidelity of the untouched parts (the template includes it, and the system restores it if deleted).

About results: an edit is a guided re-generation, not pixel surgery — untouched parts are highly faithful but not bit-identical. If unsatisfied, adjust the instruction and retry (a new task; the original is never altered).

17. Text to Speech (AI Agent)

Turn a script into natural spoken audio: type or paste the copy, pick a voice and a model tier, and synthesize an mp3 in one click. Pause markers, emotion tags (Expressive tier) and AI script polish are built in; history and examples live on the right side of the workspace.

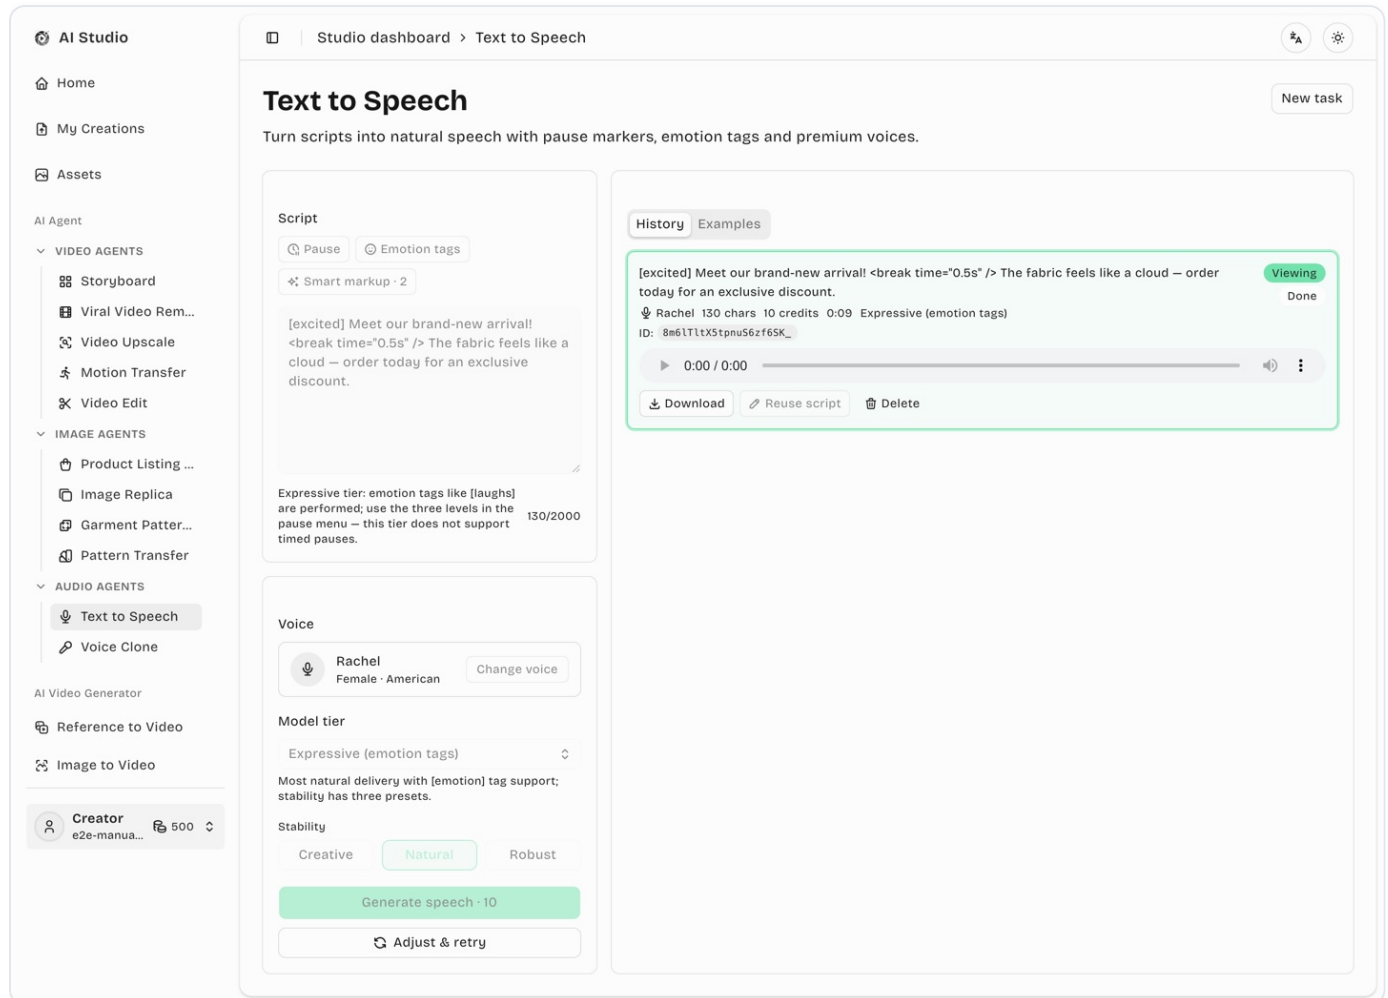


Figure 17-1 Text to Speech workspace (script and parameters on the left, history/examples on the right)

Steps

1. **Enter the script:** up to 2,000 characters (billing counts characters, inserted markers included). The toolbar inserts **pause markers** — the menu follows the tier: the **Standard** tier offers exact durations (0.3 / 0.5 / 0.8 / 1 / 1.5 / 2 / 3 s, read as silence of that length), while the **Expressive** tier offers short / medium / long semantic pauses (about 0.5 / 1.5 / 2.3 s) because that model does not support timed pause markers; **Smart markup** (2 credits) analyzes the script and inserts fitting pause markers and emotion tags (pauses only on the Standard tier) without changing a single word.
2. **Pick a voice:** "Change voice" opens the catalog — 16 curated female/male voices with gender filter, name search, previews and favorites; the "Custom voices" tab lists the voices you cloned (see the "Voice Clone" section below). **Your voice is remembered** — the workspace preselects the one you used last time (falling back to the built-in default if that voice was deleted).

3. **Pick the model tier: Standard** (3 credits / 100 chars) is fast and economical with full speed / stability / similarity / style sliders; **Expressive** (5 credits / 100 chars) delivers the most natural read and unlocks **emotion tags** — the toolbar groups 50 common tags under Laughter / Emotion / Breath & mouth / Delivery / Other (e.g. [laughs], [whispers], [professional], [warmly]); the model accepts open-ended descriptors, so you can also type your own, such as [frustrated sigh]. Stability comes as a three-preset control (Creative / Natural / Robust).
4. Click **Generate speech** (the button shows the live price; under 100 characters bills as 100). The task appears under **History** and refreshes to completion automatically.
5. Play the result right on the card, **Download** the mp3, **Reuse script** to edit again, or **Delete** (which also destroys the result file).

Examples: the Examples tab ships three tagged sample scripts with preview audio — "Try it" loads one straight into the form.

Note: Chinese is driven by the multilingual model and accent varies slightly per voice — preview before you pick. Failed or timed-out tasks refund automatically. My Creations and Credits & Billing are fully wired: the polish fee and the synthesis fee group into one "Text to Speech" billing entry.

Voice Clone

Upload your own reference audio and clone a personal voice — once ready it appears in the Text to Speech voice catalog under "Custom voices", and every history card offers one-click "Use in Text to Speech".




Figure 17-2 Voice Clone workspace (name and reference audio on the left, clone history on the right)

1. **Name the voice** and upload **1–3 reference recordings** (30 s–3 min of clean speech each works best — no background music or noise; the asset library works too).
2. Click **Clone voice (100 credits each)**. Cloning usually finishes within a minute, and the finished card carries an auto-generated preview.
3. Hit **"Use in Text to Speech"** to jump straight into synthesis with that voice — or pick and favorite it any time from the voice catalog's "Custom voices" tab.

Note: cloned voices are **visible and usable by you only**; each account keeps up to 2 cloned voices, and **deleting one frees the slot** (audio you already generated is unaffected). Failed or timed-out clones refund automatically. Reference quality drives similarity — use clear speech without reverb or background music.

18. My Creations

All of your work in one place — videos, images, storyboards, listing sets, replicas, garment pattern variations, pattern transfers, remakes and upscales — as cards you can filter, open, download and delete.




Figure 18-1 My Creations (task-type filter row at the top)

Filtering

- **Task type:** a row of buttons across the top (All / Reference to Video / Image to Video / Text to Video / Image to Image / Text to Image / Storyboard / Product Listing Set / Viral Video Remake / Image Replica / Garment Pattern Variations / Pattern Transfer / Video Upscale) — click to switch.
- **Status + time:** the second row filters by status (generating / done / failed) and by time (today / this month / custom range).
- **Filters are remembered:** your chosen type, status and time range are **restored automatically** the next time you open the page.

Card contents

- **Thumbnail and status badge:** generating / done / failed at a glance; video results carry a play icon.
- **Prompt, model, parameters and task ID** below the thumbnail; agent tasks also show render progress and a strip of source material.
- **Credits used:** the total for that task (agent tasks show the net total for the whole pipeline).

Common actions

1. **Open / retry:** click the card body to open that workspace in view mode; **Adjust & retry** there loads the parameters into a new task.
2. **Download:** save the result with the **Download** button.
3. **Delete:** the trash icon in the **bottom-right corner** of the card, with a confirmation dialog. Deleting removes the work and its result files, and refunds anything still running.

Note: a storyboard with a linked video shows as **one combined card** rather than two; Product Listing Set, Image Replica, Garment Pattern Variations, Pattern Transfer and Viral Video Remake likewise appear as a single card per task.

19. Asset library

Materials is one searchable library for your uploaded images, videos and audio. One-level folders keep a large library manageable without making extra file copies.




Figure 19-1 Materials · All assets and folders

18.1 Views, search and folders

- **All assets** shows the complete library. **Recent** is ordered by assets you used most recently, **Favorites** keeps marked assets together, and **Unfiled** contains assets that have no folder.
- Folder names appear in the Materials page's own rail on desktop and in the view selector on mobile. Create, rename or delete one-level folders from the same navigation.
- Use the search field to find a filename or folder. Type, reference-status and sort controls narrow the results; pages are loaded from the server, so older assets remain discoverable as the library grows.

18.2 Organize without copying

1. Select one or more rows, then **Move to folder**, **Move to Unfiled**, favorite or delete them in one batch.
2. Each asset has at most one primary folder. Moving it only changes organization metadata; its Cloudinary or R2 file remains stored once.
3. Deleting a folder never deletes its materials. They remain in **All assets** and move to **Unfiled**.

18.3 Use folders from a workspace

- Uploads made inside video, image and agent workspaces enter **Unfiled** by default. Organize them later from Materials.
- The **Library** picker (Figure 6-4) can filter All assets by folder. Picking records Recent use only and never moves the source asset.

18.4 Deletion rules

1. Use a row action or the batch bar to delete the underlying asset. Batch deletion reports deleted and skipped rows separately.
2. If an asset is **still referenced by a task**, deletion is blocked and the protected row is skipped. Moving or favoriting it is still safe.

Tip: folder membership and task references are separate. Deleting a folder never breaks previous work.

Deployment status (August 7, 2026): the folder replacement is live. Production migration `0039_asset_folder_management.sql` has been applied and the replacement Worker has been deployed.

20. Account settings

The Settings group in the left menu holds your profile, security and billing pages.

19.1 Profile




Figure 20-1 Profile page

- 1. Open **Settings** → **Profile**.
- 2. Update your **name**, **avatar** and other details, then save.

19.2 Security & phone binding

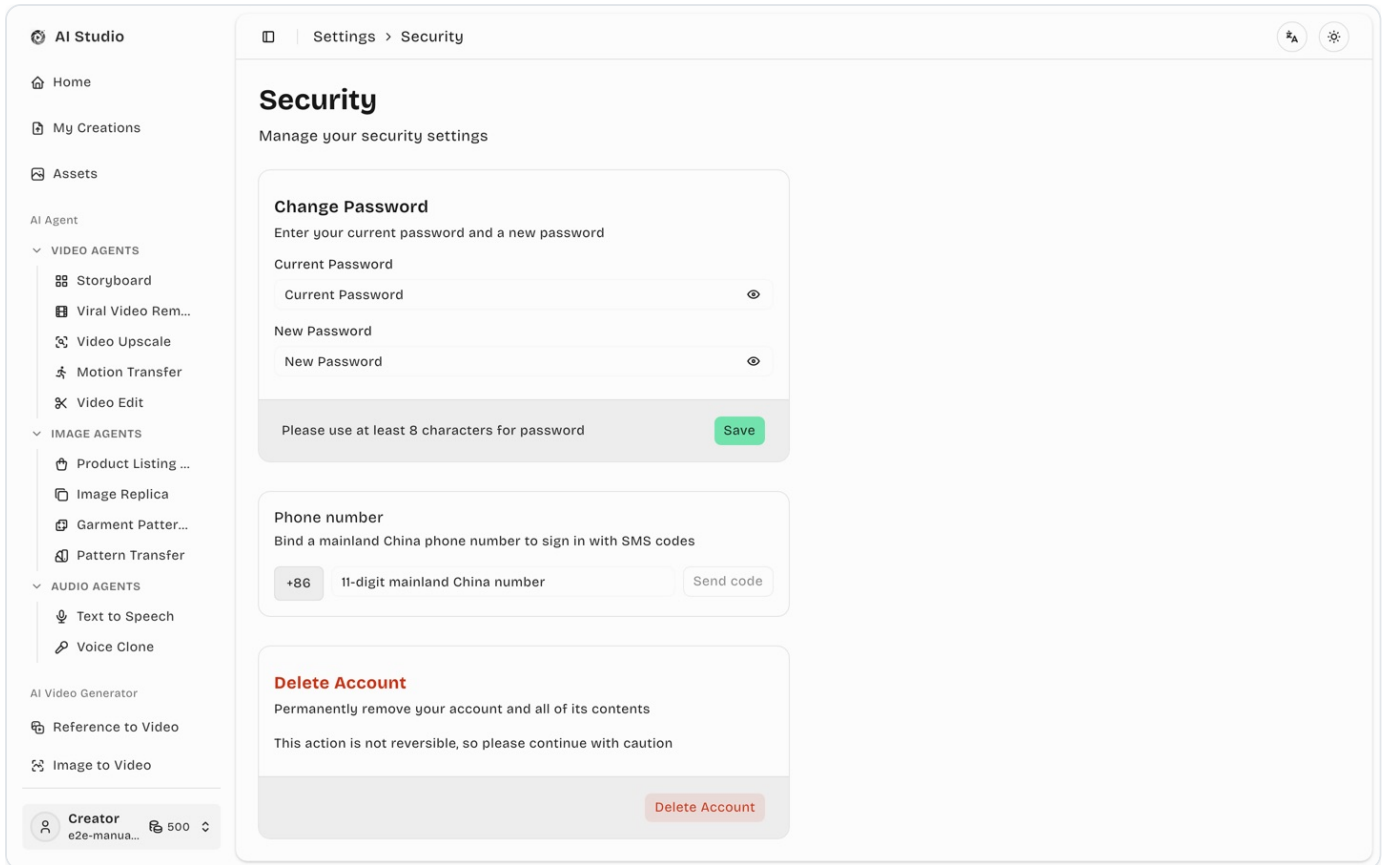


Figure 20-2 Security page (with phone binding)

Phone binding remains available as an account-security capability; the product sign-in page itself offers email/password and Google:

1. Open **Settings** → **Security** and find the **Phone** card.
2. Enter your 11-digit number and click **Send code**.
3. Enter the code and click **Bind**. The card then shows the masked number.

The same page lets you **change your password** — enter the current and new password, then save.

Note: this release supports **binding** only; rebinding and unbinding are not available yet.

19.3 Credits & billing

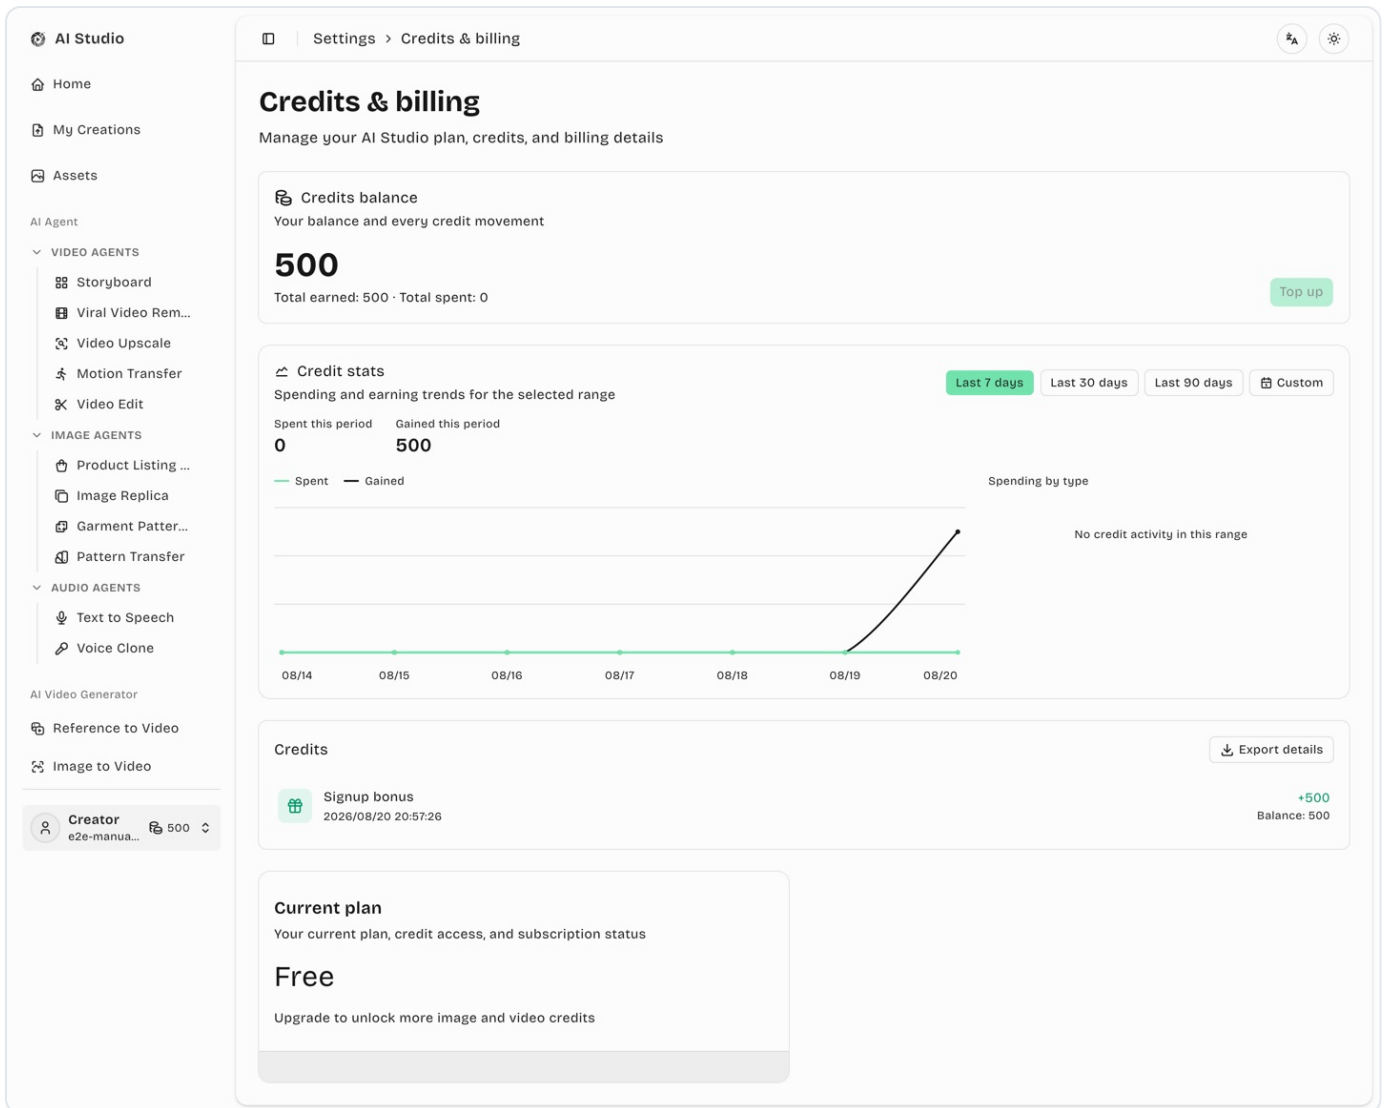


Figure 20-3 Credits & billing page

- **Balance:** your current balance plus lifetime earned and spent totals.
- **History:** sign-up bonuses, task charges/refunds, top-ups and adjustments with their time and linked task; related task movements are grouped for readability.
- **Grouped agent tasks:** the sub-charges of one storyboard, listing set, replica, garment pattern variation, pattern transfer or remake are **merged into a single row** ("Storyboard · 3 entries"); **Details** expands them and **View task** opens the work.
- **Export your usage:** click **Export details** on the Credit history card and choose the last 7 / 30 / 90 days, a custom range (up to 366 days), or all records. The UTF-8 CSV keeps one row per logical task and lists total deducted, total refunded, net consumed, balance movement and the balance after the latest movement; sign-up bonuses, top-ups and manual adjustments remain separate rows. An all-time export includes the newest 10,000 grouped records and shows a warning if older records were omitted.
- **Current plan:** shows your plan (Free). Online credit purchase is coming in a later phase.

21. FAQ

Q1 It still says "Generating" — how long should I wait?

Images usually take tens of seconds to a few minutes and videos a few minutes. The limits are 5 minutes for images (10 for GPT Image 2 Flex), about 17 minutes for videos, and up to several hours for video upscaling depending on the job (the UI shows a live estimate). Anything past its limit fails and is refunded. You can leave the page while a task runs.

Q2 My task failed — do I get the credits back?

Yes. Tasks that fail or time out are **refunded in full** automatically, and the refund appears in your billing history. A failed smart edit leaves the original image untouched.

Q3 I'm not receiving the phone-binding SMS code.

Make sure you are using a mainland-China (+86) number and that international SMS is not blocked. Retry later or contact support if it persists; you can still sign in directly with your email and password.

Q4 Can I recover the original after a smart edit?

No. A successful edit overwrites and destroys the original. Download the original first if you want a backup.

Q5 Why do some older images have no "Smart edit" button?

Smart edit needs a high-resolution working copy of the image. Very early results predate that copy, so the button is hidden for them. Everything generated since then supports it.

Q6 Text inside my images comes out wrong in Chinese.

Image models are unreliable with non-Latin scripts, so text rendered inside images is always English in the storyboard and agent flows. Spoken voice-over and subtitles can be in any supported language.

Q7 I've run out of credits.

Clicking a paid generation action shows an "Insufficient generation credits" message and does not create a task or charge your account. Top up from the Pricing page (top navbar or sidebar): four one-time tiers, currently paid by corporate bank transfer via the contact shown in the dialog. The top-up shows in your billing history.

Q8 How is Viral Video Remake charged, and can I stop midway?

Shot analysis costs 2 credits; after that **each segment is charged separately** (by model, resolution and duration) and failed segments are refunded automatically; composing the final cut is free. You can hit **Stop generating** at any time — unfinished segments are refunded, finished ones are kept and can be topped up later.

Q9 Are "planning" and "rendering" charged separately in Product Listing Set / Image Replica?

Yes. Planning (the AI writing the set plan and prompts) costs only 2 credits. Review the plan, then batch-render — each image is charged independently by model. If you don't like the plan, **Adjust & retry** starts

over; the small planning fee is not refunded.

Q10 What does Video Upscale cost and how long does it take?

The price scales with duration, frame rate and target resolution (720p is the cheapest tier, 4K about 11× that). The exact quote and an estimated processing time are shown before you submit. Long footage at a high resolution can take hours — you can leave the page — and failures or timeouts are refunded automatically.

Q11 Why does the Garment Pattern Variations workflow render the flat-lay before the model shot?

Upload your images and AI freely auto-identifies each one's role plus an overall style; confirm or edit it and submit once, and the system drafts the plan and renders every flat-lay automatically (2 credits for the plan plus every flat-lay priced by the chosen model — the total is shown right on the submit button), with no separate plan-review or per-card render-selection step. The place you still approve before spending more is the model shot: a card's **Generate model shot** button only appears once that flat-lay has rendered successfully, so you never pay for a model shot on a pattern you don't like — once every flat-lay is done, a batch toolbar appears at the top of the results, pre-checking every card whose flat-lay succeeded; uncheck any card you don't want, then click **Generate selected model shots** or **Generate selected pattern vectors**, each priced live from the checked-and-not-yet-rendered set, never re-charging a card that already has that render. For later changes, use **Edit & regenerate** in the prompt area: model-only edits charge one model shot, while a flat-lay prompt edit clearly prices and durably runs the flat-lay then model-shot chain, keeping the model on the latest flat-lay. The whole task (plan + every flat-lay + every model shot) is grouped into one "Garment Pattern Variations" billing entry.

Q12 Does pattern extraction in Pattern Transfer cost credits? How is it different from Garment Pattern Variations?

Pattern extraction is always **free**: upload a viral garment photo and the system starts a free background extraction immediately; the result is content-addressed to the source photo, so re-extracting the same photo again is instant and costs nothing. A failed extraction can be retried with **Re-extract**, or you can switch to **Upload pattern** instead. The difference from Garment Pattern Variations is the input shape and the goal: Garment Pattern Variations turns "1 garment photo + a pattern reference" into 10–20 different pattern variants, while Pattern Transfer turns "1 pattern + 1–12 existing colour versions" into that SAME pattern migrated pixel-faithfully onto every colour — better suited to when you already know the exact pattern and just want to cover every colour you sell it in. Both share the same flat-lay-then-model-shot two-stage billing and the "approve the flat-lay before spending on the model shot" gate, and are tracked and billed as separate entries in My Creations and on your ledger.

Need more help? If your question isn't covered here, get in touch through the contact details on the site and we'll help you out.