



C U R A T E D S T R E A M 4

AI for Photography

How modern AI tools are reshaping the way photographers shoot, edit and share.

AI is now part of every photographer's toolset

AI is no longer a niche plug-in.

By 2026, the leading photo and video apps ship with generative models built straight into the timeline. Photoshop, Lightroom, Capture One, Affinity, Final Cut and DaVinci all carry some flavour of AI-assisted retouching, masking or upscaling.

What used to take an hour of clone-stamp and dodge-and-burn now takes a single prompt or a click. The skill shift is moving from “how do I do this in Photoshop?” to “what do I want this picture to say?”

4

AI features in Photoshop 2026 (Fill, Remove, Expand, Upscale)

3

Generative models you can switch between in one app (Firefly Image 4, Nano Banana Pro, FLUX.2 Pro)

15+

Video models in Higgsfield AI's aggregator (Sora 2, Veo 3.1, Kling 3.0, Seedance 2.0, Wan 2.7 ...)

SECTION 1 OF 5

Where AI helps

From the shoot to the print — six places it earns its keep.

Six places AI helps a photographer



Cull & cull again

Pick the keepers from a 2,000-frame shoot. AI flags closed eyes, soft focus and near-duplicates.



Mask & select

One-click subject, sky, hair, eye and clothing masks — the boring half of retouching, automated.



Retouch & fill

Generative Fill, Remove and Expand handle bins, power lines, photobombers and tight crops.



Upscale & sharpen

Turn 24 MP into print-ready 100 MP without the plastic look of older upscalers.



Restyle & look-dev

Try a film stock, a season or a palette without re-shooting. Useful for moodboards and pitches.



Stills to motion

Turn a single frame into a 4–8 second clip for Reels, TikTok or a website hero.

Where did the 2026 leap come from?

2022

Stable Diffusion 1.5

Free image generation goes mainstream. Output is rough.

2023

Generative Fill in Photoshop

First built-in AI fill. Adobe Sensei retired in favour of Firefly.

2024–2025

Firefly Image 3 / FLUX.1

Photorealism becomes routine. Text inside images becomes readable.

2026

Photoshop 2026 multi-model

Pick Firefly Image 4, Nano Banana Pro or FLUX.2 Pro inside a single editor.

The big shift: you no longer pick a tool, you pick a **model** — photorealism, character consistency or stylised — inside the same editor.

SECTION 2 OF 5

Editing inside Photoshop 2026

Generative Fill, Remove, Expand, Upscale — and a multi-model picker.

Firefly Image Editor (2026) — the four generative tools

Adobe's Firefly Image Editor ships four generative tools inside Photoshop 2026. Each replaces a slow manual technique with a prompt.

Generative Fill

Add or replace anything inside a selection. The 2026 version is geometry-aware — added objects match the scale, rotation and lighting of the scene.

Generative Remove

One-click removal of distractions: power lines, tourists, dustbins, lens flares. Cleans up edges where the old Spot Healing Brush would smudge.

Generative Expand

Drag the crop handle outward and Firefly invents the rest of the scene. Save a portrait that was framed too tight without re-shooting.

Generative Upscale

Turn a 24 MP frame into a print-ready 100 MP file with believable fine detail — fabric weave, eye texture, hair.

One editor, three generative models

Photoshop 2026 lets you switch models for each fill or generation. Pick the model that fits the shot.

Firefly Image 4

Adobe

Best for

Default. Commercially-safe training data — the model you reach for when the image will be paid work.

Watch out

Slightly less photoreal than FLUX.2 Pro on close-up skin.

Nano Banana Pro

Google DeepMind

Best for

Character consistency across a series — same subject, different shots. Strong with multimodal references.

Watch out

Slower than Firefly. Less control over fine retouching.

FLUX.2 Pro

Black Forest Labs

Best for

Photorealism. Skin, hair, fabric and lighting that hold up on a big print.

Watch out

Less restrictive on training data — check rights for commercial release.

Firefly AI Assistant — the editor talks back

Firefly AI Assistant

I have 40 raw portraits from yesterday's shoot. Cull the keepers, even the white balance, then export web-ready JPEGs at 2000 px.

Done. 12 keepers selected, white balance neutralised against the grey card in frame 003, JPEGs exported to /Portraits/web. Want me to add a soft contrast tweak before posting?

Yes please — and remove the lanyard from frame 18.

Soft contrast applied. Generative Remove ran on frame 18; the shirt collar now reads cleanly. New version saved alongside the original.

Released April 2026

Conversation, not menus.

The Firefly AI Assistant orchestrates multi-step jobs across Creative Cloud. Speak the outcome, and it picks the tools.

What it can chain together:

- Cull, rate and rename a shoot in Lightroom
- Apply edits across selected images
- Run Generative Remove on a specific frame
- Export to social-ready dimensions
- Build a contact sheet or InDesign layout

It is not a replacement for taste. It is a faster way to execute a decision you have already made.

Train a custom model on your own work

Firefly Custom Models

Upload 10–20 of your own photographs. Firefly learns your colour grade, lighting and subject style — then applies that look to new prompts and generative fills.

Why photographers care:

- Keep visual continuity across a series
- Match new generated elements to your existing portfolio
- Bake a signature look into Generative Fill output
- Train on a character, location or product for a campaign

Your training images stay private to your account. Custom Models are not used to train the public Firefly model.



Pick a look

Choose 10–20 images that share a style — colour palette, lighting, mood.



Train

Firefly fine-tunes a model in 20–40 minutes. No coding.



Use it

Pick your custom model in the Generate dropdown alongside Firefly Image 4.

SECTION 3 OF 5

From stills to motion

Turn a photo into 4–8 seconds of usable footage.

Turning stills into short films

Photographers are asked for motion as often as stills. Two model families dominate in 2026.

Sora 2

OpenAI

- 720p standard render
- Lengths: 4 / 8 / 12 seconds
- Object permanence + physics: weight, inertia, fabric, liquid
- API: US\$0.10 per second of video

Veo 3.1

Google

- 1080p output
- Up to 8 seconds per clip
- Strong scene control + reference-image input
- Best for branded campaign work

Higgsfield AI — the aggregator

One subscription, 15+ video models.

Higgsfield AI is a marketplace-style aggregator. You write the prompt once and run it against multiple video models side by side — then pick the best take.

Why it matters for photographers:

- Skip the cost of separate Sora / Veo / Kling accounts
- Compare physics, motion and style in one view
- Useful for client work where you need options
- Native support for image-to-video starting from your photograph

Models available in 2026

Sora 2

Veo 3.1

Kling 3.0

Seedance 2.0

Wan 2.7

+ 10 more


Tip: run the same prompt across three models in parallel — you will spot the winner in seconds.

SECTION 4 OF 5

A repeatable workflow

Shoot → cull → retouch → motion → publish.

A repeatable shoot-to-publish workflow



Where AI saves you the most time

Steps 2 and 3 — culling and routine retouching — used to consume more than half of a working photographer's editing time. In 2026 they are largely automated. The hours you free up go back into Steps 1 and 5: the parts only you can do.

Prompting for photographers — anatomy of a shot prompt

Prompt example

A **weathered fisherman** standing on a **Nelson harbour jetty at golden hour**, shot on **85mm f/1.8, shallow depth of field, Kodak Portra 400 grain**, **soft side-lighting from the west**, **documentary feel**.

Subject

Who or what — with specific, vivid detail.

Setting

Where, when, weather, time of day.

Camera spec

Lens, aperture, film stock, depth of field.

Light

Direction, quality, colour temperature.

Mood / style

Genre cue — documentary, editorial, cinematic.

Refining results — the second pass

Adjust, don't redo

- If 80% works, change only the failing part.
- Use Generative Fill on the broken region.
- Lock the seed where the tool allows.

Layer references

- Drag in 1–3 reference images.
- Specify which aspect each reference covers (pose, light, palette).
- Most 2026 tools support multi-reference input.

Iterate small

- Change one variable at a time.
- Keep a saved prompt history — you will reuse phrasing.
- When stuck, switch models (Firefly ↔ FLUX ↔ Gemini).

SECTION 5 OF 5

Ethics, rights & safety

What changes when the camera is also an editor.

Disclosure, copyright & consent in 2026

Disclosure

When an image is materially altered by AI, say so. Major stock libraries, news outlets and competitions now require an “AI used” flag at submission. Adobe embeds Content Credentials (C2PA) in every Firefly export.

Copyright

In NZ, AI-generated work has unsettled status. The safer path: substantially edit AI output before commercial release, keep your prompt history, and credit Firefly Image 4 / FLUX.2 Pro / Nano Banana Pro where it does serious image work.

Consent

Likeness rights still apply. Do not generate a recognisable real person without permission. Use Nano Banana Pro's image consistency on yourself or paid talent — not on someone in a crowd shot.

EU AI Act

Fully applicable from 2 August 2026, with the AI Act Omnibus (political agreement July 2026) adding rules on AI-generated intimate content. If you sell into the EU, know the deepfake and disclosure rules.

WRAP-UP

What to keep on your radar

A short list of tools and habits worth your attention.

Tools on your 2026 radar

Photoshop 2026

The hub

adobe.com/photoshop

Multi-model Firefly Image 4 / Nano Banana Pro / FLUX.2 Pro picker. Default for any working photographer.

Firefly AI Assistant

Conversational editor

firefly.adobe.com

April 2026 release. Drive Creative Cloud by conversation — cull, edit, export end-to-end.

Sora 2

Photo → short film

sora.com

Text or image to video. Object permanence, physics. 4 / 8 / 12 s clips.

Veo 3.1

Cinematic motion

deepmind.google

1080p clips up to 8 s. Strongest scene control. Reference-image input.

Higgsfield AI

Video aggregator

higgsfield.ai

One subscription, 15+ video models including Sora 2, Veo 3.1, Kling 3.0 and more.

Midjourney v7

Concepting & moodboards

midjourney.com

Still the go-to for mood, concept and editorial-style imagery.

Try it this week

Three quick experiments

1

Clean a single photo

Open one of your own photos in Photoshop 2026. Use Generative Remove to clean a distraction. Note how long it takes.

2

Expand a tight crop

Pick a portrait that was framed too tight. Drag the crop outward and let Generative Expand fill in the room.

3

Make a 4-second clip

Take one keeper and upload it to Sora 2 or Higgsfield. Generate a 4-second motion clip. Post it.

Three things to take away

1. Photoshop 2026 is multi-model.

Pick Firefly Image 4 (default), Nano Banana Pro (consistency) or FLUX.2 Pro (realism).

2. The four big edits are generative.

Fill, Remove, Expand, Upscale — each replaces an hour of manual work.

3. Stills can move.

Sora 2, Veo 3.1 or Higgsfield turn a keeper into 4–8 seconds of motion. Pair them with your stills.