Make the flamingo lagoon ad

Remake a reference shot with your own face and a new setting, then change just the float. This guide documents the workflow behind my pink flamingo Short.

Watch the original 30-second Short

1. Prepare your inputs

Use a short video with a clear subject and a readable camera move, plus a clear portrait of the person you want in the result. Use footage, likenesses and audio you have permission to reuse. Keep an untouched copy with its soundtrack.

2. Set up Genjutsu

On Higgsfield, open Video, then Create Video, and select Genjutsu. Its current help page lists 4-30 second reference clips. Upload the video and portrait. Motion transfer is the website mode for rebuilding a scene; Objects swap is for a targeted replacement. Check the quoted cost before generating.

3. Build the lagoon version

Use the first prompt on page 2. Keep the performance and camera instructions explicit: seated pose, phone gestures, bobbing, uninterrupted pullback and final fade. Describe one coherent location that can survive the full camera move.

4. Inspect the entire result

Check the face, hands, phone, water contact and coastline throughout playback. Compare the pullback against the source. Save the version you want to keep before making another change.

5. Replace only the float

Use the accepted lagoon video as the new input. Choose Objects swap on the website, retain the portrait as an identity reference if needed, and use the second prompt. Keep the flamingo neck away from the face and phone.

Original production settings: the saved 12 September run used Genjutsu through MCP, model **hf_mult_replace_object**, at 720p. The input was trimmed to 30 seconds after the first request was rejected. The website route above is the current documented alternative.

COPYABLE PROMPTS

Build, then revise

These reusable prompts are adapted from the saved lagoon and flamingo requests. Replace the portrait and source with your own material. They specify the intended result; inspect what the model actually produces.

Pass 1 / Face and location

Edit the reference video, keeping its continuous camera pullback, shot timing, framing, body actions and final fade. Replace the person with the adult in my portrait. Match their face naturally to the existing seated pose, clothing and light. Preserve the phone, hand gestures, head turns and floating motion.

Rebuild the setting as a secluded volcanic lagoon: calm deep teal water, warm late-afternoon sunlight, a black-sand shore, basalt rocks and mossy green cliffs. As the camera pulls back, reveal one small cedar cabin near the shore. Keep the geography consistent through the whole shot. No extra cuts, crowds, boats, text or logos. The mood is peaceful solitude, enjoying being exactly where you want to be. Keep the existing donut float for this first pass.

Pass 2 / Donut to flamingo

Make one targeted change to this accepted lagoon video: replace the orange sprinkle donut float supporting the person with a pink inflatable flamingo. Give it a ring-shaped body, curved neck, small head with a black-tipped beak, subtle wings, glossy vinyl and realistic water contact.

The ring must support the person in the same seated position and follow the existing bobbing and trajectory. Put the neck to the side so it never blocks the face or phone. Remove the donut sprinkles. Preserve the person, clothes, hands, phone, volcanic lagoon, cliffs, cabin, lighting, camera pullback, timing and final fade. Do not redesign the background or add cuts, people, text or logos.

Keep the accepted version as your revision input. Starting again from the original shot gives the model more to rebuild and makes it harder to preserve the lagoon you already liked.

FINISH AND CHECK

Keep the sound.
Check the whole shot.

The original workflow restored the source soundtrack after generation. A request to preserve music inside a visual prompt is not proof that the audio survived.

Restore the soundtrack

With FFmpeg installed, put the generated picture in **generated.mp4** and the permitted original soundtrack/video in **source.mp4**. Run this in that folder; backslashes continue the command across lines.

```
ffmpeg -i generated.mp4 -i source.mp4 \
-map 0:v:0 -map 1:a:0 \
-c copy -movflags +faststart final.mp4
```

This takes the picture from the first file and audio from the second. It does not fix a duration or timing mismatch. Compare both durations first, then play the finished file from start to finish. If they differ, align the audio in your editor before exporting.

Before exporting

- The face and glasses stay recognisable throughout the pullback.

- Hands and phone remain coherent; the float supports the body.

- The flamingo neck stays clear of the face and phone.

- The coastline and cabin stay consistent as the camera retreats.

- The soundtrack begins correctly, stays in sync and reaches the intended ending.

About the $7 / 8.5-minute headline

Those were the figures stated in the Short, not a guaranteed budget or turnaround. The saved follow-up quoted 195 credits per 720p render and approximately 390 for the first version plus the flamingo revision; it was an estimate, not an inspected billing receipt. Check current pricing and count revisions.

Sources and useful links

Official Genjutsu setup guide (checked 2 November 2026).
Original AI Luke Short.
Workflow and prompts adapted from the saved 12 September 2026 production.