No Rehearsals

From character sheet to cinematic scene

A practical tutorial for Krea 2 and MiniMax H3





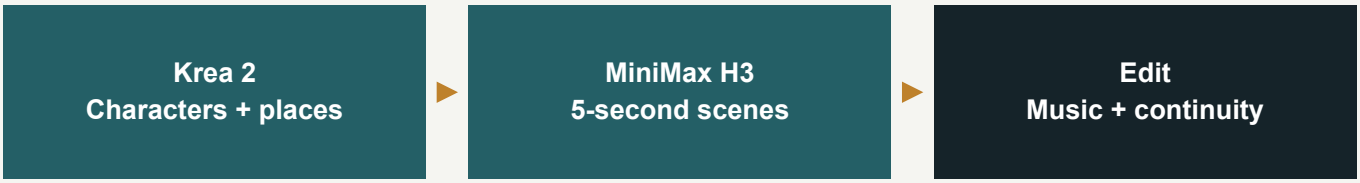


Reference-to-video / Image-to-video / First frame + last frame

Build the scene in layers

Use the project assets to make a small, consistent film world before asking it to move.

The workflow reel pairs finished shots with the ComfyUI graphs behind them. This guide turns that reel into a repeatable production method: design Maya and Mateo, establish the spaces, choose the right conditioning mode, then direct a single readable action.



Start here	What you will make
Characters and places, pages 4-8	Two reusable character sheets, an environment library and scene-ready stills.
Video workflows, pages 9-12	Correct reference roles, opening-frame control and matched endpoint pairs.
Scene recipes, pages 13-26	Lighter, corridor, driving, city, portrait, mirror, map, motel, camera and revolving door.
Finish and resources, pages 27-32	Editing, troubleshooting, location index, workflow downloads and sources.

What to prepare

ComfyUI with your chosen MiniMax H3 workflow installed, access to Krea 2, the supplied JPG assets, and a video editor. Begin at a resolution your installed workflow supports. The reel records a local setup; its performance is not a promise for every machine.

How to use the prompts

Each scene recipe names the files to load and provides a copyable prompt. Asset names in this document are relative to the project root. The companion text file contains the prompts without page breaks. Reconstructed recipes are identified where no standalone source still exists.

Choose the right input

A reference, an opening frame and an ending frame perform different jobs.

Your starting point	Workflow	Use it for
Character sheets + environment	Reference-to-video	Recompose known people in a new place or camera angle.
One finished scene image	Image-to-video	Keep the composition and animate a specific action.
Two matching scene states	First frame + last frame	Connect a dark and lit frame, or a clean and marked mirror.

What stays fixed across the film

Maya: blunt black bob and fringe, silver hoops, worn black leather jacket, oxblood dress and black boots. Mateo: swept dark waves, stubble, charcoal blazer, washed-black top and dark trousers. Use the actual sheets as the final visual authority when a written description disagrees with an approved image.

The world is coastal noir: humid air, petrol-blue shadows, amber practicals, wet streets and restrained red signage. Close-ups isolate a face or a hand; wides show how small the fugitives feel inside the city. Preserve this contrast instead of giving every shot the same shallow focus.



Intimate: warm light and close bodies.



Exposed: cold city scale and deep space.

Keep the story legible

One shot should deliver one visible change: a lighter ignites, a finger touches glass, a door rotates, or two runners hold distance from a retreating camera. Leave complicated story information to the sequence of shots. The project brief avoids weapons, drugs, drinks and cigarettes.

Create the character sheets

Krea 2 is the character-design stage for this tutorial.

1. Open Krea image generation and select Krea 2. Choose a landscape 3:2 canvas for the three-view sheet. If using a local Krea 2 graph, set the latent width and height there.

2. Paste the full character prompt from the next two pages. Generate Maya and Mateo separately. Neutral studio light makes identity and clothing easier to inspect than dramatic scene lighting.

3. Compare the portrait, front and rear views. Reject inconsistent hair length, changing dress hems, mismatched jackets or missing boots. Approve one definitive sheet per character.

4. Save the approved sheets as Maya Chen.jpg and Mateo Cruz.jpg. Keep them unchanged through the first production pass; changing the sheet can change every later shot.

5. Load a sheet into reference-to-video as an identity and wardrobe source. Explicitly say that its three panels depict one person. A sheet is not an opening scene frame.

Use the sheet for identity, then light the scene

Krea documents style references and moodboards as creative controls. For this project, use them for palette and photographic texture; judge character consistency by the approved sheet. If a local workflow exposes only text-to-image, do not assume that attaching a picture adds identity conditioning. Use a compatible reference workflow for that task. [1]

Approval check	Accept only when...
Identity	Face structure, fringe, stubble and distinctive features agree.
Wardrobe	The front and back describe the same garment and accessories.
Framing	Head, hands and both boots are visible; figures are not stretched.
Reuse	The sheet is clean, readable and free of scene props or graphic labels.

Maya Chen

Krea 2 character-sheet prompt | Save as Assets/Maya Chen.jpg



Create a photorealistic character reference sheet for the female lead of the cinematic music video "No Rehearsals."

CHARACTER: Maya Chen, a 25-year-old Asian American woman of Chinese heritage. Beautiful, magnetic, mysterious, with a quietly defiant presence. Oval face, defined cheekbones, dark brown almond-shaped eyes, straight natural brows, softly full lips, warm light-medium skin with visible natural texture. A small beauty mark beneath her left eye. Jet-black chin-length blunt bob with a slightly uneven, eyebrow-length fringe; a few loose strands soften the precise haircut. Subtle smoky eyeliner and muted rose lip gloss. Her expression is guarded and searching, as though she knows the camera has caught something she meant to hide.

BUILD AND WARDROBE: Approximately 168 cm tall, slim build with natural proportions. A deep oxblood satin slip dress with narrow straps, a gently draped neckline, mid-calf hem, and a modest slit on her left side. Worn black leather jacket ending at the waist, worn open, sleeves down. Black leather ankle boots with a narrow square toe and low block heel. Small silver hoop earrings in both ears and one plain silver ring on her right index finger. No other accessories.

COMPOSITION: One landscape 3:2 image divided into three clean vertical panels. Left: a detailed close-up portrait, head and shoulders, facing directly toward the camera, hair and entire chin visible. Center: full-body front view, standing naturally, arms slightly separated from her torso, hands relaxed, both boots entirely visible. Right: full-body rear view, facing directly away, showing the exact back of her haircut, jacket, dress, and boots. Front and rear figures at identical scale and aligned at the feet.

VISUAL TREATMENT: Cinematic photographic realism, restrained urban noir, subtle 35mm film grain, rich fabric texture and believable leather wear. Plain muted gray studio background. Soft neutral light clearly reveals her face, skin tone, and clothing, with faint teal and amber edge accents. Keep shadow detail readable.

Maintain exactly the same woman, facial structure, haircut, body proportions, outfit, and accessories across all three panels. The rear view must be anatomically consistent with the front. Sharp detail throughout each figure. No props, text, labels, logos, watermarks, extra people, cropped feet, exaggerated posing, plastic skin, or illustration.

Mateo Cruz

Krea 2 character-sheet prompt | Save as Assets/Mateo Cruz.jpg



Create a photorealistic character reference sheet for the male lead of the cinematic music video "No Rehearsals."

CHARACTER: Mateo Cruz, a 29-year-old Latin American man from Mexico. Handsome, magnetic, mysterious, with restrained intensity and a trace of exhaustion. Medium olive-brown skin with natural pores, an angular face, pronounced cheekbones, strong jaw, slightly crooked nose bridge, deep-set dark brown eyes, thick eyebrows with a small natural notch in his right eyebrow. Short, even dark stubble. Thick dark brown wavy hair, medium length on top, loosely brushed back with one strand falling toward his forehead; shorter sides, natural neckline, no skin fade. His mouth is relaxed, his gaze direct and difficult to read, as though he is weighing a confession.

BUILD AND WARDROBE: Approximately 183 cm tall, lean athletic build, broad shoulders with natural proportions. A slightly worn charcoal single-breasted blazer, hip length, worn open over a fitted washed-black ribbed tank top. Straight-leg black tailored trousers with a medium rise, plain black leather belt, and worn black leather Cuban-heel ankle boots. A thin silver chain resting at his collarbone and one small silver hoop in his left ear. No other jewelry. The clothes feel carefully chosen but lived in, with subtle creasing.

COMPOSITION: One landscape 3:2 image divided into three clean vertical panels. Left: a detailed close-up portrait, head and shoulders, facing directly toward the camera, entire hairstyle and chin visible. Center: full-body front view, standing naturally with shoulders relaxed, arms slightly separated from his torso, hands open and visible, both boots entirely visible. Right: full-body rear view, facing directly away, clearly showing his hair at the nape, blazer construction, trouser fit, and boots. Front and rear figures at identical scale and aligned at the feet.

VISUAL TREATMENT: Cinematic photographic realism, restrained urban noir, subtle 35mm film grain, tactile cloth and leather, believable skin texture. Plain muted gray studio background. Soft neutral light clearly reveals facial features, true skin tone, and garment details, with faint teal and amber edge accents. Keep shadow detail readable.

Maintain exactly the same man, facial structure, hairstyle, stubble, body proportions, outfit, and accessories across all three panels. Preserve the eyebrow notch and left-ear hoop on their anatomically correct sides. Sharp detail throughout each figure. No props, text, labels, logos, watermarks, extra people, cropped feet, exaggerated muscles, plastic skin, or illustration.

Build a reusable location library

Create empty spaces first. Add the characters when you compose the scene.



Motel room: map, mirror and planning scenes.



Service corridor: flight and silhouettes.



Broadcast complex: surveillance and arrival.



Revolving-door lobby: glass and reflections.

Generate each place as one continuous image

Use a landscape 16:9 canvas and a self-contained description. Describe architecture, materials, practical lighting, time of day and camera position. Keep open floor space for the characters. Do not generate a stacked location sheet or a collage.

The project includes 18 location JPGs, plus a vehicle reference and two city backgrounds. Use one approved image per location for the first pass. A later angle may change the viewpoint; it should retain the same doors, windows, furniture and lighting sources.

A useful folder rule

Store the character sheets in Assets, the environment images in Assets/Locations, approved scene stills beside 01.jpg through 06.jpg, and generated clips in a separate shot folder. Record which asset versions were connected to every accepted take.

The complete 18-location index appears on pages 29-30. Recovered location prompts are included in the companion prompt file.

Generate the motel room

Krea 2 location example | Start from the project prompt, then inspect the actual canvas.



Generate one photorealistic cinematic location image in landscape 16:9. One continuous composition, no panels or collage.

A small, worn coastal motel room at the end of evening. A double bed against the left wall, a round wooden table and two mismatched chairs beside a large front window on the right. Venetian blinds cast thin shadows. At the rear, a narrow vanity holds a rectangular mirror, an empty shallow stainless-steel tray, and an amber lamp. A bathroom doorway stands beside the vanity. Muted olive walls, faded beige bedding, dark patterned carpet, an old air conditioner beneath the window.

Photograph from the entrance with a moderately wide lens, clearly establishing the table, bed, window, and vanity. Leave clear space around the table for two characters. Cool residual daylight, red exterior neon reflected in the glass, subdued amber practical light.

Dark intimate cinematography, anamorphic optical character, faint haze, fine film grain, gentle highlight bloom and edge distortion. No people, weapons, drugs, drinks, cigarettes, captions, or watermarks.

Check dimensions outside the prompt

One inspected Krea PNG embeds an 8-step graph with CFG 1.0, er_sde sampling and a simple scheduler. Its prompt requests 16:9, but the latent canvas is 1920 x 1920. Text alone did not set its shape. For a landscape result, change the actual canvas dimensions before generating. Those observed settings belong to that saved graph, not to all Krea 2 variants.

For the character sheets, choose 3:2. For locations, choose 16:9. For final movie frames, use the common aspect ratio selected for the edit.

Reference-to-video

Compose a new shot from known people, objects and spaces.

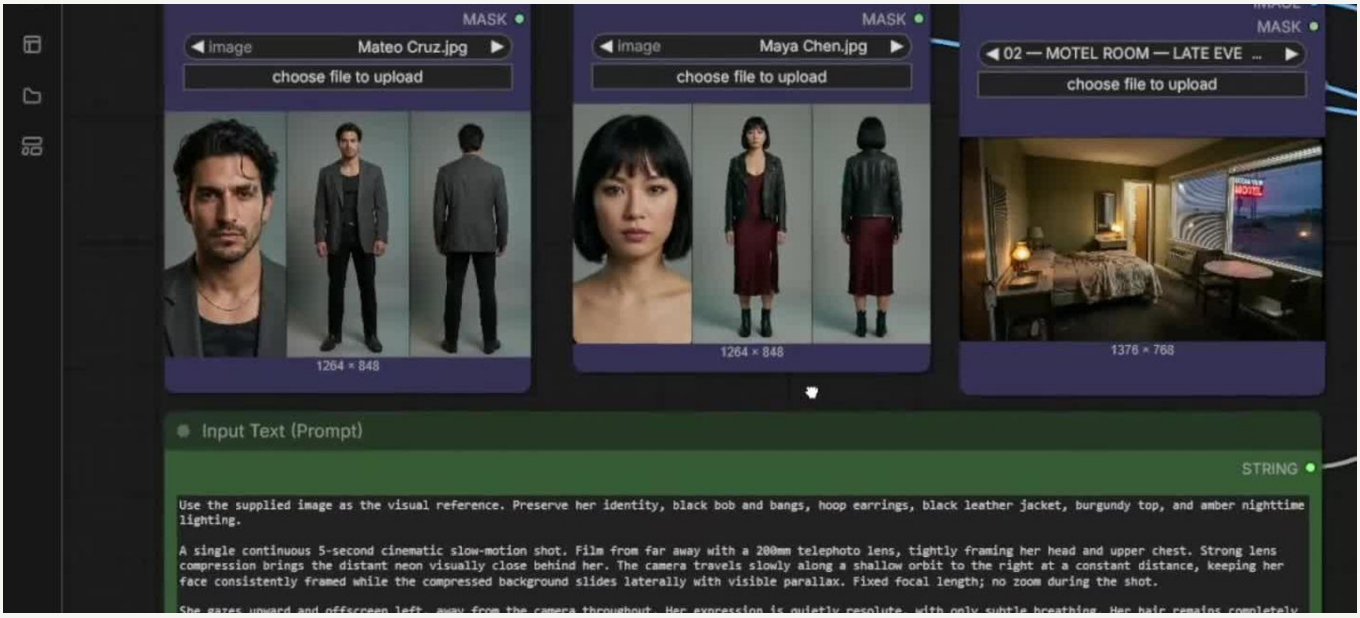
Load images
Identity + location

▶

H3 reference
Prompt + controls

▶

Save video
Review the take



Recorded asset loaders: Mateo, Maya and the motel room. The recipe prompt must match the connected image order.

1. Open the MiniMax H3 reference-to-video graph from the workflow pack. If nodes are missing, install the graph's documented dependencies before substituting anything.

2. Reload every image from your own computer. Connect character sheets and the chosen environment to the reference inputs in a deliberate order.

3. Write a role for every input. Example: Image 1 is Maya, Image 2 is Mateo, Image 3 is the motel room. These are reference roles, not three frames of a sequence.

4. Set a 5-second test duration and supported output dimensions. Generate one take, inspect identity and spatial layout, then alter one variable for the next pass.

H3 separates keyframe-conditioned and reference-conditioned model paths. Load a workflow built for the mode you need; similar-looking image sockets are not proof that two modes are interchangeable. [2]

Image-to-video

Use an approved scene still when the composition already works.

Load scene still
Exact opening

▶

H3 keyframe path
Action prompt

▶

Save video
Inspect start + end



Assets/01.jpg already contains the driver, car, lighting and composition.

1. Load one complete scene still into the first-frame input. Leave the last-frame input unused unless the workflow requires an explicit option to disable it.

2. Match the output aspect ratio to that still. Choose one camera behavior and one physical action, then state what must remain unchanged.

3. Describe the motion after the opening moment. The image already supplies texture, wardrobe and set design; the prompt should direct what changes.

4. Inspect the first frame, the main action and the last frame. Check the whole clip at normal speed before approving it. A good still at the end does not excuse a warped middle.

The driving example

The car and driver stay fixed relative to a camera mounted on the car, while the outside city moves. A locked camera can still show a moving world. The prompt must identify what the camera is locked to, or the model may freeze the entire image.

Use the 01.jpg and 02.jpg recipes on pages 17-18. The city recipe on page 19 keeps the camera stationary in world space instead.

First frame + last frame

Use matching endpoint images to control a specific state change.



First: 05b.jpg / lighter unlit.



Last: 05.jpg / lighter burning.

Load first + last
Same crop + size

▶

H3 keyframe path
Describe the transition

▶

Save video
Check the landing

1. Load 05b.jpg into first_frame and 05.jpg into last_frame. Both supplied files are 1264 x 848, so their aspect and dimensions already match.

2. Keep the same camera axis, person scale and room geometry in both images. When preparing new pairs, align them before generation; resizing unrelated images cannot solve incompatible compositions.

3. Explain why the change occurs. Here, a thumb strikes the lighter, a spark catches, and warm light reveals the faces. Do not ask for a generic fade between two stills.

4. For the reverse action, swap the two images and rewrite the action as a controlled breath extinguishing the flame. Do not assume that reversing a completed clip produces natural human action.

Resolve the format before the final render

The lighter pair is about 3:2, while several other project stills are wider. For a uniform film, prepare both endpoints in the same final crop or composition. Test that neither the flame nor either face is lost. Never change the crop of only one endpoint.

Write a prompt that directs

Fix the anchors. Name the action. Specify the camera. Describe the result.

Prompt element	Example from this project
Anchors	Keep Maya on frame left, Mateo on frame right, and preserve the rain-speckled window.
Action	He strikes the lighter once. The flame catches and settles into a steady burn.
Camera	Locked camera, identical framing and focal length throughout.
Timing	Hold for 1.5 seconds, ignite over 1 second, then settle to the final frame.
Physical cause	The flame changes the light; ambient exposure stays constant.
Finish	Reach the lit reference and hold it. No extra action after the landing.
Sound	One lighter strike and quiet room ambience; no dialogue or generated score.

Plain-language prompt or structured model prompt?

The recipes here are production-facing instructions for a workflow prompt box. Some H3 nodes perform prompt processing for you. A direct base-model path may need the structured prompt format described in the official guides: visual timeline, soundscape and music fields. Reference-mode preprocessing also tracks subject definitions and retention. Follow the input contract of your installed graph. [3, 4]

Change one thing per iteration

If the camera moves correctly but the faces drift, simplify the action and strengthen identity references. If the people look right but the framing changes, tighten the camera instruction. Keep the seed fixed when the graph exposes it, so the next comparison answers a specific question.

All scene durations in the recipes are creative targets. Use the controls supported by your specific H3 workflow. Requested endpoint matching and timestamps still require visual review.

Make the lighter ignite

First frame + last frame | 5 seconds | Dark to lit



First input: Assets/05b.jpg



Last input: Assets/05.jpg

Load the dark frame first and the illuminated frame last. This is the direction shown at the beginning of the workflow reel. Keep both images in the same crop.

Use 05b.jpg, the dark frame, as the exact first frame and 05.jpg, the illuminated frame, as the exact last frame. Create one continuous 5-second cinematic shot. Lock the camera completely: identical framing, perspective, and focal length throughout. Preserve both characters' identities, clothing, positions, the rain-speckled window, and corridor geometry.

[0-1.5 seconds] They face each other in silence, their faces in deep silhouette. He holds the unlit lighter raised between them. Maintain the dim teal ambience and subdued orange wall tones of the opening reference.

[1.5-2.5 seconds] He makes one small, natural thumb movement to strike the lighter. A brief spark immediately catches into a small flame. Warm light appears in direct response to ignition, revealing their faces, his hand, and nearby clothing.

[2.5-5 seconds] The flame settles into a steady burn with gentle natural flicker. He maintains his grip and holds the lighter between them. She remains still, lips relaxed, looking at him. Match the warm illumination, expressions, and poses of the final reference precisely.

Keep ambient lighting and camera exposure constant; the lighter provides the change in illumination. Preserve organic film grain and natural anatomy. No blowing, pursed lips, extinguishing, repeated ignition attempts, oversized flame, cuts, dissolves, or camera movement. Audio: one crisp lighter strike, a soft ignition, and faint corridor ambience; no dialogue or music.

Approve when the flame starts once, lights the faces for a physical reason, and stays lit. Reject an exposure ramp that brightens the entire room or a flame that appears before the thumb strike.

Make the flame disappear

First frame + last frame | 5 seconds | Lit to dark



First input: Assets/05.jpg



Last input: Assets/05b.jpg

This alternate uses the same two source images in reverse order. It is a new performance: Maya exhales, the flame bends and goes out, then the scene settles into darkness.

Use 05.jpg as the exact first frame and 05b.jpg as the exact last frame. Create one continuous 5-second cinematic shot. Lock the camera completely: identical framing, perspective, and focal length throughout. Preserve both characters' identities, clothing, positions, the rain-speckled window, and corridor geometry.

[0-1.5 seconds] They face each other in silence. He holds the burning lighter steadily between them. The flame flickers gently, casting warm light across their faces.

[1.5-2.5 seconds] She subtly leans toward the lighter, purses her lips, and gives one short, controlled breath. The flame bends away from her toward screen right, flutters briefly, then extinguishes completely.

[2.5-5 seconds] The flame's warm illumination disappears immediately, allowing their faces and clothing to fall into deep silhouette. The scene settles into the dim teal ambience and subdued orange wall tones of 05b.jpg, without an exposure adjustment brightening their faces. She relaxes her lips and returns to the final reference pose. He keeps the extinguished lighter raised as they hold each other's gaze. Finish precisely on 05b.jpg.

Maintain natural movement, restrained expressions, organic film grain, and consistent anatomy. No cuts, dissolves, camera movement, relighting flashes, reignition, exaggerated blowing, or visible breath cloud. Audio: one soft exhale and faint corridor ambience; no dialogue or music.

Approve when the flame bends away from the breath and extinguishes once. The camera and ambient lighting remain unchanged. Trim any dead hold in the edit only after the action itself works.

Run down the corridor

Image-to-video | One continuous 5-second shot



Load: Assets/04.jpg. For reference-based composition, use both character sheets + Location 05.

Use the supplied image as the exact opening frame. Preserve both characters' identities, dark clothing, the woman on frame left, the man on frame right, and the narrow industrial corridor with cyan overhead lighting and warm orange accents.

One continuous 5-second handheld tracking shot. Both characters run urgently forward toward the camera with natural strides, coordinated arm swings, and focused expressions. The camera retreats down the corridor at exactly their running speed, maintaining a constant distance and consistent full-body framing. They remain approximately the same size throughout while the corridor walls, overhead fixtures, and floor move past with strong depth parallax.

Use believable handheld movement: rhythmic vertical bounce from the camera operator's backward steps, slight lateral sway, and small framing corrections. Keep their faces readable and both bodies within frame. Their clothing and the woman's bob respond naturally to running. Their shadows travel across the right wall as they pass beneath the lights.

Maintain the reference's dark cinematic exposure, cyan haze, warm practical lights, organic film grain, and natural motion blur. Real-time running, no speed ramps. No cuts, zoom, orbit, excessive shake, sliding feet, duplicated limbs, or changing faces. Audio: rapid footsteps, breathing, and clothing movement reverberating through the corridor; no dialogue or music.

Why this motion reads as running

The camera retreats at the runners' speed, so their size stays stable while the walls rush past. Natural bounce and arm swings make the movement physical. If the background stays fixed, the result can look like running on a treadmill. Use real-time motion here; the latest corridor direction in the project chat overrides the earlier slow-motion brief.

Build the driving angle

Reference-to-video | One continuous 5-second shot



Load: Image 1: Assets/Mateo Cruz.jpg. Image 2: Assets/the car.jpg.

Create one continuous 5-second cinematic side-window shot of a man driving through a wet city at night.

Image 1 defines Mateo: preserve his face, swept dark wavy hair, stubble, charcoal blazer and washed-black top. Its panels are different views of one person. Image 2 defines the dark vintage coupe and the cool blue wet-street setting. Preserve its body shape, window design, trim and interior. Render one Mateo in one car; do not reproduce either reference as a collage.

Compose a new tight side view through the driver's window. Mateo sits behind the steering wheel, looking ahead, one hand resting naturally near its upper rim. Lock the filming camera relative to the moving car. Keep the man, window frame and cabin the same size and screen position throughout.

The city moves smoothly from screen right to screen left with layered depth and restrained motion blur. Mateo breathes subtly and makes only a small steering correction. Between 2 and 3 seconds, headlights from one passing car sweep pale light across his face, hand and cabin exactly once, then disappear.

Match the vehicle reference's petrol-blue night ambience, wet glass, soft reflections, dark exposure and organic grain. No cuts, zoom, orbit, changing car design, eye contact, extra passengers or dialogue. Quiet road and engine ambience only; no generated music.

Reconstructed reference-mode recipe based on the approved assets and the project's driving instructions. For tighter control of the exact shot, use the existing 01.jpg opening frame on the next page.

Drive past the city

Image-to-video | One continuous 5-second shot



Load: Assets/01.jpg. Camera locked to the car; city moves behind it.

Use the supplied image as the exact opening frame. Preserve the man's identity, profile, clothing, hand position on the steering wheel, vintage car interior, and cool blue nighttime palette.

One continuous 5-second cinematic shot. The camera is rigidly mounted outside the driver's side window, locked relative to the moving car. Maintain exactly the same framing, angle, focus, and subject scale throughout. He drives steadily forward, eyes fixed on the road, expression calm and serious. His posture remains consistent, with subtle breathing and minimal natural steering adjustments.

Outside, city buildings and streetlights move smoothly from screen right to screen left, with believable depth parallax and restrained motion blur. The driver, window frame, and car remain stable in the composition.

From 2-3 seconds, headlights from one passing vehicle create a single soft sweep of pale light across his face, hand, and cabin, briefly revealing detail before fading back into blue shadow. This light sweep occurs exactly once.

Preserve rain-streaked glass, soft reflections, deep shadows, and organic film grain. No pan, tilt, zoom, orbit, camera shake, cuts, repeated light flashes, eye contact, or dialogue. Audio: low engine hum and subdued road noise; no music.

Three checks before approving

The driver keeps his eye line on the road. The cabin remains stable as the background travels. The passing headlights produce one sweep, not a repeating strobe. If the car appears stationary, strengthen the background parallax instruction before adding more body movement.

Cut to the steering hand

Image-to-video | One continuous 5-second shot



Load: Assets/02.jpg. Keep the grip and wheel geometry stable.

Use the supplied image as the exact opening frame. Preserve the driver's hand, finger placement, silver bracelet, dark sleeve, worn steering wheel, rain-marked windshield, and cool blue nighttime palette.

One continuous 5-second cinematic insert shot. The camera is rigidly mounted inside the moving car, locked on the hand gripping the top of the steering wheel. Maintain the same framing, angle, subject scale, and shallow focus throughout. Keep the hand and wheel sharp while the city beyond the windshield remains softly defocused.

He drives steadily forward, maintaining his grip with only a tiny, gradual steering correction. Outside, the approaching street expands naturally from the distant vanishing point, with nearby buildings and streetlights drifting outward faster than distant details. Smooth depth parallax and restrained motion blur convey forward travel.

From 2-3 seconds, headlights from one passing car produce a single soft sweep of pale light across his knuckles, bracelet, sleeve, and steering wheel, then fade back into blue shadow. The sweep occurs exactly once.

Preserve natural skin texture, soft reflections, deep shadows, and organic film grain. No camera movement, focus pulls, cuts, large wheel rotations, finger repositioning, hand distortion, or repeated light flashes. Audio: low engine hum and subdued road noise; no dialogue or music.

Use it as a continuity insert

Place this between wider driving shots. Match the travel direction and blue ambience. Watch each finger during the light sweep: extra joints, changing grip and a deforming bracelet are easier to spot in motion than in a thumbnail. Keep steering corrections small.

Let the city breathe

Image-to-video | One continuous 5-second shot



Load: Assets/03.jpg. A world-locked camera; only traffic and water move.

Use the supplied image as the exact opening frame. Preserve the skyline, bridges, river, elevated highways, building geometry, and cool blue nighttime palette.

One continuous 5-second cinematic establishing shot from a completely locked elevated camera. Maintain identical framing, perspective, focus, and exposure throughout. The city remains still while small streams of headlights and taillights move naturally along the highways and across the bridges, following the existing lanes and road curves. Traffic moves steadily at realistic speed, without time-lapse streaks.

Gentle ripples travel across the river, subtly breaking up the reflected bridge and city lights. Distant atmospheric haze softens the skyline. Building windows and bridge lights remain steady.

Retain the subdued cyan highlights, deep blue shadows, fine architectural detail, and organic film grain. No pan, tilt, zoom, orbit, camera drift, cuts, sweeping foreground headlights, changing architecture, or added text. Audio: distant traffic wash and faint river ambience; no dialogue or music.

Use small motion to protect the skyline

The skyline supplies scale while moving traffic supplies life. A locked shot gives the model fewer architectural changes to invent. Check bridge cables, road curves and tower edges from start to finish. Cars should follow existing lanes rather than materialize across the river.

03.jpg is 1584 x 672. Keep that composition for a test, or prepare an approved common-ratio frame for the final edit. Avoid stretching the image to make it fit.

Compose Maya in neon

Reference-to-video | One continuous 5-second shot



Load: Image 1: Assets/Maya Chen.jpg. Image 2: Assets/out of focus bg.jpg.

Create one continuous 5-second cinematic close-up of Maya in an amber-and-teal city at night.

Image 1 supplies the identity, blunt black bob and fringe, hoop earrings, leather jacket and oxblood clothing of one adult woman. Do not reproduce its three-panel sheet. Image 2 supplies only the diagonal neon bands, warm halation, cool color accents, shallow focus and background texture. Do not copy a recognizable face or person from the background reference.

Frame Maya's head and upper chest from a distant camera position with the appearance of a 200mm telephoto lens. Keep her slightly right of center. Her eyes look toward a fixed point offscreen left, away from the camera. Preserve her face and natural skin texture. Her head and hair remain still, with only subtle breathing and no wind.

The camera travels slowly along a shallow arc to the right at a constant distance. Keep her face at a stable size and screen position. Use a fixed focal length; there is no zoom. The compressed diagonal orange neon slides gently behind her with smooth parallax, remaining soft and defocused.

Retain deep shadows, amber practical light, gentle highlight halation and organic film grain. No cuts, head tracking, eye contact, extra people, drifting facial features or added text. Faint city ambience only; no speech or generated music.

Reconstructed reference-mode recipe, adapting the latest telephoto-orbit direction from the shot-prompt chat. For a darker street environment, replace Image 2 with out of focus bg 2.jpg and update its role in the prompt.

Move around a held gaze

Image-to-video | One continuous 5-second shot



Load: One approved single-frame portrait of Maya. Do not load the entire character sheet as the first frame.

Use the approved single-scene portrait as the exact opening frame. Preserve her identity, black bob and bangs, hoop earrings, black leather jacket, burgundy top, and amber nighttime lighting.

A single continuous 5-second cinematic slow-motion shot. Film from far away with a 200mm telephoto lens, tightly framing her head and upper chest. Strong lens compression brings the distant neon visually close behind her. The camera travels slowly along a shallow orbit to the right at a constant distance, keeping her face consistently framed while the compressed background slides laterally with visible parallax. Fixed focal length; no zoom during the shot.

She gazes upward and offscreen left, away from the camera throughout. Her expression is quietly resolute, with only subtle breathing. Her hair remains completely motionless; still air, no wind or loose strands moving. Preserve warm neon halation, deep shadows, shallow depth of field, natural skin texture, and organic film grain. Focus stays on her eyes. No cuts, eye contact, head tracking, or facial distortion. Audio: faint distant city ambience, no dialogue or music.

An orbit is different from a zoom

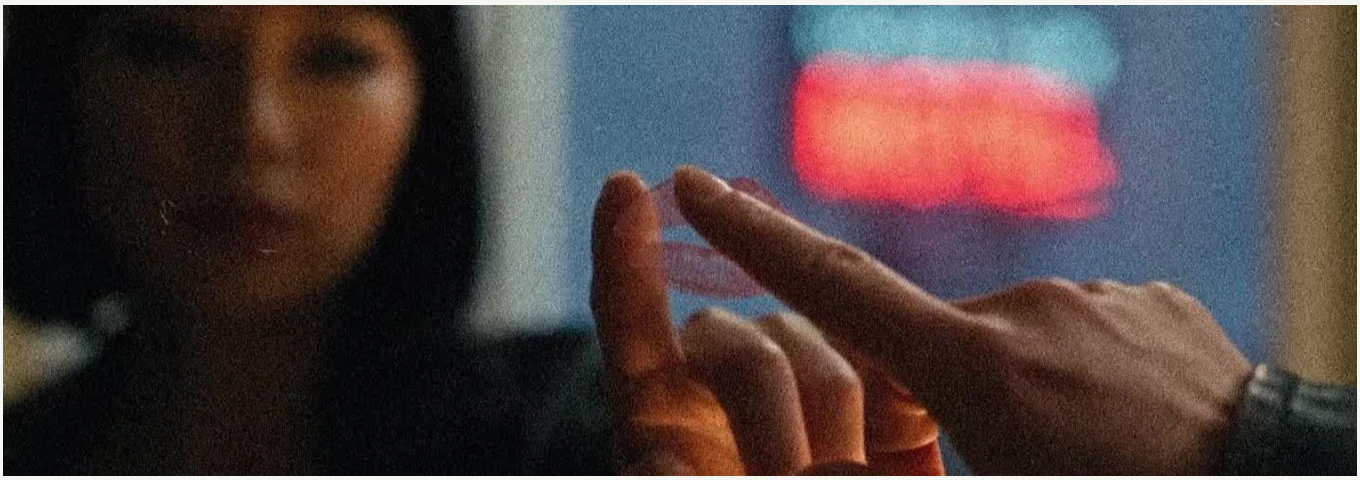
The lens stays fixed while the camera changes position. Background parallax is the visible cue. The 200mm description is a look target, not physical lens metadata. If the scene warps during the arc, reduce the travel rather than adding a zoom to disguise it.

To prepare the opening still, compose it with Maya's approved identity and the neon background reference. The local folder has no separately named final portrait still; the figure here is a crop of the completed workflow reel. Use a clean full-resolution source still for a new generation.

The two alternate portrait prompts from the project chat are included in the companion text file. They preserve still hair, an offscreen gaze and fixed-focal-length movement.

Leave a trace on the mirror

First frame + last frame | One continuous 5-second shot



Load: Prepare two aligned mirror frames: hand absent, then trace left behind.

Use the clean mirror image as the exact first frame and the matching marked-mirror image as the exact last frame. Make one continuous 5-second cinematic shot. Preserve the glass plane, Maya's indistinct reflection, the lip-gloss mark, the red motel-sign glow and the warm practical lighting.

From 0 to 1 second, hold on the lip-gloss trace and imperfect glass. Maya's face remains out of focus behind it. From 1 to 3 seconds, her hand with a black leather cuff enters from the lower right. One fingertip touches the glass beside the gloss and drags downward a short distance, leaving a faint fingerprint streak. The finger physically contacts the mirror rather than floating in front of it.

From 3 to 5 seconds, the fingertip lifts and the hand withdraws. Hold on the abandoned trace and converge on the final reference. Keep the camera locked, the glass surface sharp and the face softly defocused. Maintain the same mirror geometry and exposure across the shot.

Photorealistic coastal noir, amber practical light, subdued teal fill, red neon bokeh, soft halation and fine grain. Natural hand anatomy and reflection. No face reveal, extra hands, cuts, dissolves, new writing or moving camera. A faint fingertip-on-glass sound and quiet room ambience; no dialogue or music.

Reconstructed from the mirror action in the project chat and the paired-frame setup shown in the reel. The endpoint files are not separately supplied in Assets. Create them from Maya + Location 02, keeping the same view; do not use a character sheet as either endpoint.

Turn the photograph

Reference-to-video | One continuous 5-second shot



Load: Image 1: Maya. Image 2: Mateo. Image 3: Location 02, motel room.

Generate one continuous 5-second photorealistic cinematic insert across a worn wooden motel table.

Image 1 defines Maya's identity, black leather jacket and silver ring. Image 2 defines Mateo's identity, charcoal blazer and dark clothing. Each sheet shows one adult person in several views; do not reproduce the panels. Image 3 defines the motel's room, round table, olive walls, window and amber practical lamps. Create a new close camera angle inside this same room.

A small architectural photograph lies partly over a folded city map. Mateo's thumb, emerging from a charcoal sleeve, pins one corner. Maya's hand enters from the opposite side, its black leather cuff visible. She puts two fingertips on the opposite corner. Mateo releases his thumb, allowing her to rotate the photograph a quarter-turn toward herself. Their hands do not touch. Finish with the photograph still between their hands.

The camera makes a slow, short lateral slide. Keep fingertips, the worn photograph edge and the paper folds in a narrow plane of focus. The image on the photograph is architecture only. Map lettering remains indistinct. Amber lamplight grazes the paper while a soft red neon reflection lies across the tabletop.

Preserve natural hand anatomy, paper thickness, jacket textures and consistent light. Dark intimate coastal noir, shallow focus, restrained anamorphic bokeh and fine film grain. No faces, extra hands, duplicate photographs, melting paper, cuts, exaggerated camera movement, dialogue or music. Quiet paper movement and room ambience only.

Adapted as a standalone shot from the project's motel-planning scene. The map and photograph are described props; no separate prop files were found. Keep the hand action simple if the papers distort.

Hold the motel wide

Reference-to-video | One continuous 5-second shot



Load: Image 1: Maya. Image 2: Mateo. Image 3: Location 02, motel room.

Create one continuous 5-second cinematic wide shot inside the supplied motel room after sunset.

Image 1 supplies Maya's identity, black bob, leather jacket, oxblood dress and boots. Image 2 supplies Mateo's identity, dark wavy hair, stubble, charcoal blazer and dark outfit. Render one of each character; do not copy the reference sheets' panels or backgrounds. Image 3 supplies the actual motel layout: olive walls, bed on the left, vanity and warm lamp, round table beside the large window, mismatched chairs and Venetian blinds. Preserve these spatial relationships.

Shoot from the entrance through a dark, softly blurred doorway. Maya and Mateo sit opposite each other at the table, small inside the room. Maya is on the right, seen from behind. Mateo is near the window in a shadowed three-quarter view with his head lowered. A folded map and small architectural photograph lie between them.

Lock the camera. Maya slowly draws her hands away from the papers and rests them in her lap. Mateo remains leaning forward, studying the photograph. Neither looks toward the other or at the camera. Their separation is the emotional center of the composition.

Keep enough depth of field to read the room. Amber practical lamps and cool dusk through the blinds provide the light; the exterior red sign is soft in the distance. Fine organic grain, gentle halation, restrained haze and rich shadow detail. No extra people, moving furniture, camera motion, cuts, readable overlay text, drinks, cigarettes, dialogue or generated music. Quiet room ambience only.

Adapted from the project's motel-planning prompt. Check the bed, round table, window and chairs first: if their arrangement changes, strengthen the location role before adding further gestures.

Watch through the camera

Image-to-video | One continuous 5-second shot



Load: A clean CCTV-dome opening still. Build it from Location 03 if needed.

Use the supplied security-camera close-up as the exact opening frame. Create one continuous 5-second live-action cinematic shot beneath the entrance canopy of a waterfront broadcast building at blue hour.

Preserve the black smoked-glass dome, its mounting, moisture beads, small red status light, reflected amber entrance lights and two tiny reflected approaching silhouettes. The outer housing is fixed to the building. Only the camera lens inside the transparent dome slowly pivots toward the entrance. The red indicator glows steadily.

The filming camera makes a very short, slow lateral slide. Begin with focus on the fine moisture beads, then gently shift attention toward the internal lens. Keep the reflected people indistinct; never resolve their faces or add more figures. The reflected forecourt should change naturally with the filming camera's movement, without turning into a second scene.

Retain cool blue ambient light, an amber edge flare, humid atmosphere, soft highlight halation and organic grain. End when the internal lens faces the approaching silhouettes. No moving outer housing, bending dome, multiplying reflections, cuts, zoom, added text, dialogue or music. Quiet exterior ambience and a faint mechanism sound only.

Adapted from the surveillance shot in the project's broadcast-complex scene. No standalone CCTV still was found. First compose a dome close-up using Location 03, optionally referencing both character sheets for the distant silhouettes; approve it before animation. The figure is taken from the workflow reel.

Pass through the glass

Reference-to-video | One continuous 5-second shot



Load: Image 1: Maya. Image 2: Mateo. Image 3: Location 04, revolving-door lobby.

Generate one continuous 5-second cinematic shot from inside the supplied waterfront lobby, looking obliquely through its revolving door at blue hour.

Image 1 defines one Maya in her black leather jacket, oxblood dress and boots. Image 2 defines one Mateo in his charcoal blazer and dark clothing. Preserve their identities and wardrobe; the sheets' multiple views do not represent multiple people. Image 3 defines the lobby: glazing, bronze door frames, central revolving-door axis, reception area, materials and amber interior light. Keep the architecture rigid.

Maya enters the first compartment and Mateo follows in the next compartment. They move at an unhurried pace. The door rotates naturally around its fixed central axis. Lock the filming camera while the wings and mullions pass across their bodies. Blue exterior reflections and amber interior reflections overlap the glass without creating extra people.

Maya clears the door and moves toward the lobby interior, pausing near the edge of frame. Mateo emerges behind her. Finish with both oriented deeper into the building. Keep their faces obscured by hair, mullions and reflected glare; do not ask them to follow the camera.

Selective focus on clothing and the nearest door frame, soft anamorphic bokeh, light haze, gentle halation and organic grain. No cuts, camera orbit, bending glass, duplicated limbs, sliding feet, extra characters, dialogue or generated music. Quiet footsteps and a subdued revolving-door mechanism only.

Adapted as one shot from the project's broadcast-complex scene. If two people cannot clear the door cleanly in five seconds, generate one character's passage per take or use a supported longer duration. Preserve door mechanics before adding a more elaborate performance.

Turn the takes into a sequence

Build the rhythm in the edit after each individual shot works.



1. Choose the clearest version of each action. The workflow reel is a demonstration montage, so its short cuts are not a requirement to keep every generated take at its full duration.

2. Put the original song on a dedicated audio track. Mute or remove generated music from the clips. Add selected ambience only where it helps the cut and does not compete with the track.

3. Build an initial order: lighter, corridor, driver, portrait, city, mirror, map, motel, surveillance, revolving door, final portrait. Reorder for the music rather than treating this as a fixed narrative chronology.

4. Match direction across cuts. Keep the car traveling consistently, preserve Maya and Mateo's screen positions in connected actions, and check that hands and props match between inserts and wides.

5. Normalize the grade gently. Protect the distinction between amber intimacy and blue exposure. Avoid lifting dark faces until the noir mood disappears. Add final grain sparingly if the source is already textured.

6. Export a test and watch it from start to finish. Confirm the correct song, natural movement, consistent framing and the absence of unwanted generated speech.

Choose one film canvas

The project contains several source ratios. For a widescreen film, approve a common frame such as 21:9 and prepare crops consistently. For a 9:16 workflow reel, use a vertical presentation canvas: a landscape film strip above a readable workflow capture, as in the supplied reel. Reframe scenes deliberately for a full-screen vertical version; do not squeeze a wide image into a tall shape.

The supplied workflow reel is 1440 x 2560 at 24 fps. The separate teaser is 1920 x 1080 at 24 fps with cinematic framing inside it. These are observed exports, not mandatory H3 generation settings.

Fix the visible problem

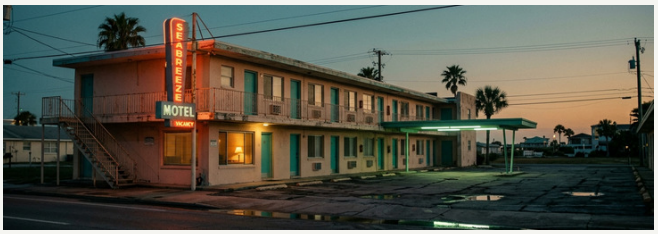
Diagnose the failed relationship before rewriting the entire prompt.

What you see	What to change next
Several copies of Maya	Say that the sheet shows one person in several views. Use it only as identity reference. For I2V, load a single composed scene frame.
The set becomes the gray studio	Remove sheet-background cues from the scene description. Assign the environment image its own location role.
The car freezes with its background	Lock the camera to the car and explicitly animate the outside city with depth parallax.
Repeated headlight flashes	Request one passing vehicle and one sweep in a short time window. Remove other flicker or flashing-light language.
Lighter scene fades between stills	Describe thumb movement, spark, flame and the resulting illumination. Keep exposure and ambient lights constant.
Faces brighten after the flame dies	Keep automatic exposure compensation out of the desired action. Reassert the final dark state and its reference.
Hair moves in the portrait	Use the latest direction: still air, motionless hair, fixed offscreen gaze. Restrict motion to the camera and subtle breathing.
The orbit looks like a zoom	Specify constant distance, constant subject scale, fixed focal length and a shallow camera arc.
Door glass bends or doubles people	Reduce the action to one clear passage. Lock the camera and preserve the axis, wings and rigid mullions.
Square output despite a wide prompt	Set width, height or aspect ratio in the generator. Verify actual file dimensions before the next stage.
Good first and last frames, broken middle	Reduce the change between endpoints or simplify the gesture. Review the entire take, not just its thumbnails.

Log the failure and the one change you made. Keep a successful identity or camera setup while solving the next problem.

Locations: motel to hideout

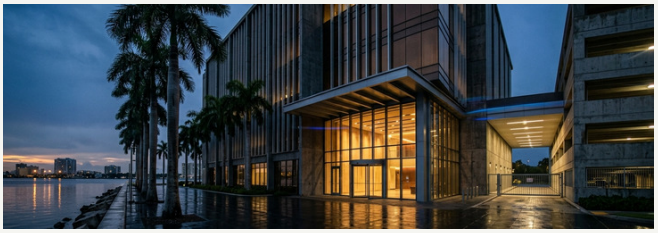
Each JPG is one environment image, ready to combine with the character references.



01 / Motel Exterior
Late Evening



02 / Motel Room
Late Evening



03 / Waterfront Broadcast Complex
Blue Hour



04 / Revolving-Door Lobby
Blue Hour



05 / Service Corridor
Night



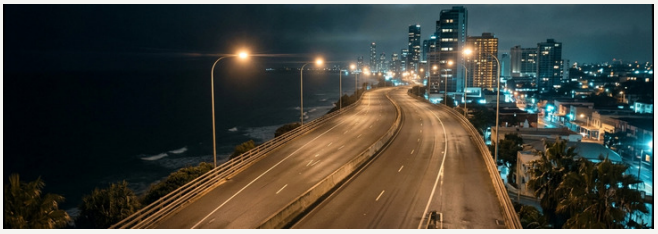
06 / Private Media Archive
Night



07 / Empty Recording Studio
Night



08 / Waterfront Parking Garage
Night



09 / Elevated Coastal Highway
Late Night

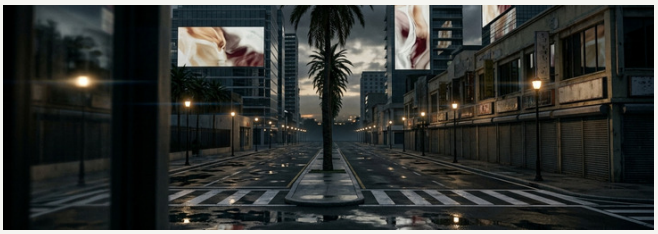


10 / Industrial Overpass Hideout
Late Night

Find the original JPGs in Assets/Locations. Preserve their descriptive filenames when relinking a workflow. This library supports the larger story as well as the scenes shown in the reel.

Locations: city to first light

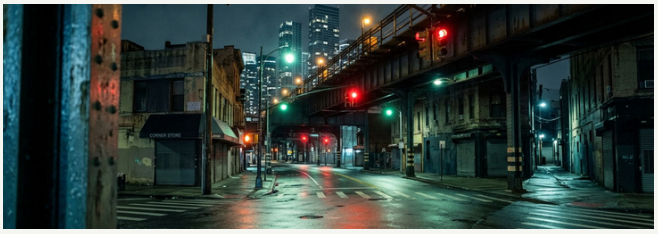
Each JPG is one environment image, ready to combine with the character references.



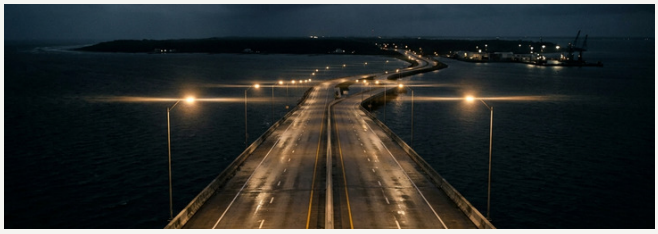
11 / Palm-Lined City Boulevard
Late Night



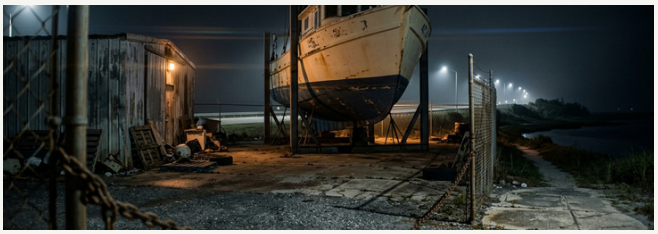
12 / Glass-Fronted Newsroom
Late Night



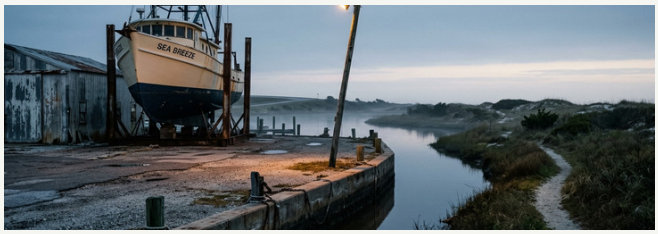
13 / Elevated-Rail Intersection
Late Night



14 / Open-Water Causeway
Deepest Night



15 / Abandoned Boatyard
Night



16 / Boatyard Waterfront
Predawn



17 / Dune Path
First Light



18 / Remote Beach
Early Morning

Find the original JPGs in Assets/Locations. Preserve their descriptive filenames when relinking a workflow. This library supports the larger story as well as the scenes shown in the reel.

Additional reusable assets

Assets/the car.jpg is the vehicle reference. Assets/out of focus bg.jpg supplies the soft orange-and-teal neon. Assets/out of focus bg 2.jpg supplies a darker night street. Assets/01.jpg through 06.jpg are scene images, not empty environments; 05b.jpg is the lighter's alternate state.

Package a repeatable workflow

Keep the creative decision and the files that produced it together.

Before every generation

1. Confirm the mode: reference-to-video for a new composition, I2V for a finished opening still, or paired frames for a controlled state change.

2. Check the input order against the prompt. Name each file and its role. Remove unused references that might compete with the scene.

3. Confirm the actual output duration and dimensions. Make matched endpoint pairs identical in size and crop.

4. Write one camera behavior and one primary action. Keep the final state clear.

Before accepting a take

1. Watch the whole clip at normal speed. Scrub faces, hands, props, reflections and architecture around the main action.

2. Check the opening and ending against their reference images. Reject obvious drift, duplication or a transition that hides rather than performs the action.

3. Listen for unwanted speech or music. Save the successful prompt, connected asset names, workflow version, seed if exposed, and output file together.

Suggested handoff contents

Folder	Keep inside it
Characters	Approved Maya and Mateo sheets, plus the Krea 2 text prompts.
Locations and frames	The environment JPGs, approved opening stills and matched endpoint pairs.
Workflows	The working H3 graphs and their dependency notes. Relink local images after import.
Renders and edit	Selected takes, review notes, the original song and final exports.

Example take name: lighter-ignite_fl_v03.mp4. Keep your naming consistent, and do not overwrite the accepted take during experiments.

Workflows and source notes

Download the companion graphs from the shared project folder.

Open the workflow download folder

Scan to open the same Drive folder.

The workflow files are distributed separately. This edition documents the local reel and project assets; it does not bundle or claim to validate unpublished workflow JSON.

Project material used

Reviewed: Export/Workflow-reel/no-rehearsals-video-workflow-teaser.mp4; the separate teaser; the two character sheets; numbered scene JPGs; 18 location JPGs; the vehicle and background references; and Krea PNG metadata. Read the related chats "Analyze No Rehearsals reference" and "Write Minimax H3 shot prompts" for the original character, story and motion instructions.

The character prompts and several I2V prompts are reproduced from project chats. The reference-mode scene recipes and mirror/CCTV preparation steps are adapted or reconstructed for this tutorial. They are creative instructions, not claims that a new render was generated and tested during writing.

Technical references

[1] Krea 2 documentation - Model selection, ratios and reference/style controls.

[2] MiniMax H3 repository - Keyframe and reference mode families; supported system outputs.

[3] H3 base/keyframe prompt guide - Structured prompt format for direct keyframe paths.

[4] H3 reference prompt guide - Reference roles and structured reference processing.

[5] fal H3 prompting examples - The guide linked in the original shot-prompt chat.

[6] Taste Skill design guidance - Adapted for print: strong hierarchy, actual project imagery and restrained surfaces.

Prepared 17 October 2026. Model controls and node packages can change; use the dependency notes supplied with the downloaded graph.