

GUIDE

How to write prompts that produce usable AI video

A video prompt has four jobs — subject, action, camera and style. Write them in that order, keep one idea per clip and iterate on the weakest element instead of rewriting everything.

Updated November 13, 2026

Contents

1. Subject first ····· 5

2. Action second ····· 5

3. Camera third ····· 5

4. Style last ····· 5

A template that works ····· 5

Iterate on the weakest element ····· 5

Practical habits ····· 6

A good video prompt is a shot description, not a wish. The models that make usable clips respond to four ingredients, in roughly this order of importance: the subject, the action, the camera and the style. Write them deliberately and you will need fewer generations per finished second.

1. Subject first

Name the thing that must be in frame, specifically. "A stainless steel water bottle on a kitchen counter" beats "a product." If the subject is a person, describe only what matters — age range, clothing, action — and leave the rest to the model. Faces are where generations fail most often; keep people small in frame or out of frame when you can.

For product work, the honest path is image-to-video: generate or photograph the real product as a still, then animate it. A text prompt cannot reliably invent your actual product, label and all.

2. Action second

Say what moves and how. "The camera slowly pushes in" is an action. "It looks cinematic" is not. Useful action words: rotates, slides, pours, unfolds, rises, tracks alongside, settles into place. One clear motion per clip beats three competing motions. If the shot needs a hand pouring liquid and a label coming into focus, split it into two clips and cut them together.

3. Camera third

Camera language carries more weight than most beginners expect. Shot size (wide, medium, close-up), angle (eye level, overhead, low) and movement (static, pan, push-in, orbit) each change the output. Luma's Dream Machine page highlights camera-motion controls; Runway and Kling expose similar motion options on their generation pages. Learn each tool's vocabulary — the words that move the camera are the tool's own terms, not universal ones.

4. Style last

Style is the finishing layer: lighting, palette, film look, mood. "Soft morning light, neutral tones, shallow depth of field" tells the model more than "professional" or "viral." Keep style consistent across a series by reusing the same style sentence in every prompt. That consistency is what makes five generated clips cut together like one shoot.

A template that works

> *[Shot size] of [subject, specific] [action, one motion]. [Camera movement]. [Lighting and palette]. [Style references: film look, aspect ratio hints].*

Example: "Medium shot of a matte black coffee mug on a wooden table, steam rising slowly. Camera pushes in gently. Warm morning window light, neutral wood tones, shallow depth of field. Clean product-ad look."

Iterate on the weakest element

When a generation misses, do not rewrite the whole prompt. Find the element that failed and change only that: the subject got the wrong shape, sharpen the subject description; the motion was wrong, rewrite the action; the mood was off, adjust the style line. Each retry costs credits — Pika's pricing page puts a 5-second 720p clip at about 122 credits on Starter — so targeted fixes beat wholesale rewrites.

Practical habits

- Generate three takes of the hero shot and pick the best; budget for it.
- Keep every prompt that produced a keeper in a document; they are your reusable shot library.
- Match prompt length to clip length: a 5-second clip needs one clear beat, not a paragraph of action.
- For anything with text in frame — prices, labels, headlines — plan to add the text in your editor afterward. Generative video garbles lettering reliably.

Frequently asked questions

How long should a video prompt be?

Long enough to pin down subject, action, camera and style — usually two to four sentences. Longer prompts dilute the important parts.

Why do my generations look different each time?

Most models add variation on purpose. Fix the elements that matter with more specific wording, and expect to generate several takes.

Should I write the prompt before or after I know the tool?

Write the shot list first. Each tool's prompt field takes the same four-part structure; only the syntax for camera moves and style tags differs.

Sources

Runway AI video generator (fetched 2026-11-13) — https://runwayml.com/ai-video-generator
Luma Dream Machine (fetched 2026-11-13) — https://lumalabs.ai/dream-machine
Pika pricing (fetched 2026-11-13) — https://pika.art/pricing