

Checking it

You cannot see your own render after the fourth viewing. This is the pass that catches what you've stopped noticing.

The review prompt

Paste this with your render and a reference clip. The phrasing matters more than anything else here.

REVIEW PROMPT

Here is my render and a reference clip that performed well.

List everything amateur in mine compared to the reference. Be specific and name timecodes. Check for:

- transitions that land on the same footage instead of new content
- the same sound firing on more than four cuts
- captions in the same position with the same treatment throughout
- b-roll all in one register, or held as a frozen frame
- cuts spaced evenly, within half a second of each other
- a static shot held past 5 seconds
- the hook not landing by 1.5s, or a first cut after 3s
- a flat CTA slide or a trailing black frame
- hands, garment prints, reflections or posters containing invented content
- framing tighter at the end of the clip than at the start

Return a numbered list of faults. Do not give an overall verdict.

Never ask "is this good?" It returns yes, every time, from anyone including a model, because it invites a verdict and verdicts default to agreement. "List everything amateur compared to this reference" returns an inventory, and an inventory has to name things.

That one swap is most of what makes reviewing your own work possible.

The four checks

Frame every cut. Pull a still at each cut and read them in order. Shows immediately whether each transition landed on new content or just squished.

Contact sheet the whole clip. Highest-value check and the easiest to skip. Scrubbing a timeline hides repetition because you experience it one moment at a time. A grid shows you that eleven of your fourteen captions sit in the same place.

First frame against last frame. Put them side by side. If the subject is larger at the end, your end-state pin is missing and the clip drifted.

Spectrogram, then listen. The picture catches a repeating sound, your ear catches a voice buried under the bed.

What to look for in the frame

The things that read as generated, in rough order of how fast someone clocks them:

TELL

WHERE IT SHOWS UP

Extra faces	garment prints, posters, reflections, photos on a shelf
Duplicated hands or objects	two brushes, three fingers, a second mug
An arm reaching in from nowhere	arm's-length plates, bedrooms especially
Colour-sorted or symmetrical anything	bookshelves, product rows, cushions
Studio-even light	no shadow under the chin, no hot spot, no falloff
A phone visible in shot	you wrote "selfie" or "phone" somewhere
Airbrushed skin	you wrote "photorealistic" or "high quality"
Framing tighter at the end than the start	no end-state pin

Where to stop

An AI-selfie base tops out around 82–85% realism against real footage. That ceiling isn't moving with better prompting.

The last fifteen points are where months go. They're worth less than the ten edit tells in the motion doc, which take an afternoon.

Go compare against the real thing rather than trusting your own eye: GRWM (<https://www.tiktok.com/@leilannyjolie/video/7548934829945982263>) · podcast (<https://www.tiktok.com/@campchaospodcast/video/7641697751721397518>) · street (<https://www.tiktok.com/@chris.stocks/video/7564450806158707982>).

If you build an automated grader

Calibrate it before you trust a single score.

BEFORE TRUSTING ANY GRADER

1. Run it on 3 clips you already know performed well.
2. If they fail, the bar is wrong. Fix the bar, not the render.
3. Re-run those 3 after ANY change to the rubric.
4. Plot the score distribution before quoting a flag rate.

A gate here once rejected 67 out of 67 renders. The obvious reading was that the output was bad. Running it against reference winners took ten minutes and it rejected those too, which meant it had been measuring its own threshold the whole time.

If a gate rejects everything, suspect the ruler before the work.