Consistency Decision Sheet

Fora Soft Learn · AI for Video Engineering

1 - The core problem (why this matters)

```
A model has NO memory between generations; a text prompt is too loose to pin a face.
Consistency = feed an EXTRA signal (reference / LoRA / shared attention / camera path).
```

2 - The four levers (read 'best for' first)

Lever	Train?	Best for
Reference conditioning	No	Fast default - same face/place from an image
Fine-tuning (LoRA)	Yes	Recurring host / mascot across many shots
Attention sharing	No	A storyboard generated together in one go
Explicit camera control	No	A repeatable dolly / orbit / crane move

3 - 2026 model feature to reach for

Need	Reach for
Same face, a few shots, fast	Runway Gen-4 ref / Veo 3.1 ingredients / Sora 2 cameo
Recurring host, thousands of shots	Open-weights model + a trained character LoRA
Matching storyboard in one go	Kling 3.0 multi-shot / StoryDiffusion (open)
Same location / object	Runway Gen-4 locations & objects, or an object LoRA
A specific repeatable camera move	CameraCtrl (open) or a model camera-path control

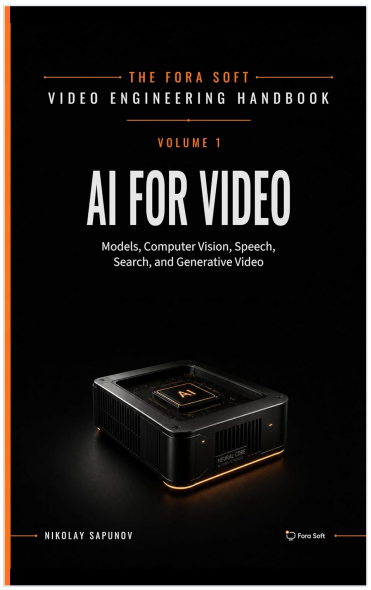
4 - Production order (layer the levers)

- [ ] Lock identity once - a saved reference set (closed model) or a trained LoRA (open model).
- [ ] Fix the world - reference-condition the location and key props so the place won't drift.
- [ ] Change one variable per shot - hold lighting and wardrobe wording constant.
- [ ] Specify the camera per shot with an explicit move, not a typed phrase, where the tool allows.
- [ ] Batch shots that must match so attention sharing works; fall back to the saved identity later.

5 - Pitfalls to check before you build

- [ ] Drift with a locked reference is almost always the PROMPT contradicting the reference - let the reference describe identity, the prompt only action and scene.
- [ ] Lighting changes regenerate faces; change one variable per shot, not five.
- [ ] Reference strength is a dial: too high locks identity but ignores the prompt; too low loosens the face.

Companion to: 'Character, Scene, Camera Consistency in Production'
www.forasoft.com/learn/ai-for-video-engineering/articles-ai/character-scene-camera-consistency-production



GO DEEPER · VOLUME 6 OF 7

AI for Video: A Complete Guide to Intelligence in Video Products

From Models and Inference to Computer Vision, Speech, Search, and Generative Video

By Nikolay Sapunov, CEO at Fora Soft · 1,444 pages · Kindle & paperback on Amazon

Read it on Amazon → amazon.com/dp/B0HH6TKPZ1

This sheet compares reference conditioning, LoRA, attention sharing and camera control for keeping faces and places consistent. Volume 6 covers generative video in depth, along with the AI feature lifecycle, model artifacts such as LoRA adapters and the deployment topologies that decide where generation runs.

WHAT YOU WILL FIND INSIDE

- Generative video models
- Lifecycle of an AI video feature
- Model artifacts and fine-tuning
- Deployment topologies for generation

THE FULL HANDBOOK · 7 VOLUMES · 9,000+ PAGES · ALL ON AMAZON

Vol. 1 · Video Encoding — pixels and color to H.264, HEVC, AV1, quality metrics and AI
Vol. 2 · Video Streaming I — network transport, contribution and ingest, HLS, MPEG-DASH, CMAF
Vol. 3 · Video Streaming II — ABR, CDNs, players, WebRTC, DRM and the economics of delivery
Vol. 4 · Audio for Video I — PCM, codecs, loudness, accessibility, real-time voice processing
Vol. 5 · Audio for Video II — jitter buffers, lip sync, Atmos, quality metrics, neural codecs
Vol. 6 · AI for Video I — models and inference, computer vision, speech, search, generative video
Vol. 7 · AI for Video II — real-time pipelines, agents, serving, cost, compliance, verticals

Need the same character and scene across hundreds of generated shots?

We will set up references or a trained LoRA for your characters and a pipeline that keeps every shot on model. 250+ video products shipped since 2005.

Book a 30-min call → **WhatsApp →** **Email us →**