



INO AI LAB

[AI EDUCATION • COURSE MATERIAL]

AI Video Generation with Sora & Veo

Sora 2, Veo 3, Runway Gen-4

LEVEL	DURATION	LESSONS
Intermediate	4 hours	8



[00]

Table of Contents

LESSON 01	The State of Video Generation in 2026
LESSON 02	Prompting for Video: Direction, Not Description
LESSON 03	Image-to-Video and Keyframe Workflows
LESSON 04	Sound, Music & Final Polish



[LESSON 01]

The State of Video Generation in 2026

Three tools dominate: Sora 2 (OpenAI) for cinematic narrative clips, Veo 3 (Google) for physical realism and synced audio, and Runway Gen-4 for editor-friendly workflows with frame-level control. All three produce minute-long, coherent video from text or image prompts. None is yet ready for unedited final delivery in high-budget productions, but all three are production-grade for social, ads, music videos, and explainer content.

Pricing has dropped 10x since 2024. A minute of 1080p generation costs around \$1–3 depending on tool and quality tier. The constraint is no longer cost — it's iteration time. Each generation takes 30–120 seconds. Plan shoots like a director: storyboard, shot list, then batch generations overnight to wake up to a full reel ready for selection.

// KEY TAKEAWAYS

- › Sora for narrative, Veo for realism, Runway for control.
- › Cost is no longer the constraint — iteration time is.
- › Plan and batch like a real shoot.

[LESSON 02]

Prompting for Video: Direction, Not Description

A weak video prompt describes a scene. A strong one directs the camera. Write like a cinematographer: 'Medium close-up, 35mm lens, shallow depth of field. Subject turns toward camera, golden-hour light from frame left. Slow push-in over 4 seconds. Warm color grade, slight film grain.' Models reward this level of specificity with dramatically more coherent, usable shots.

Specify what you don't want as well as what you do. Negative prompts cut common artifacts: 'no morphing limbs, no text overlays, no zoom, no cuts'. Lock your aspect ratio and length from the first generation — switching mid-project breaks visual consistency. Generate variations at fixed seeds for the same composition with subtle changes; this is how you find the keeper without burning credits on random rerolls.

// KEY TAKEAWAYS

- › Direct like a cinematographer — lens, light, movement.
- › Negative prompts cut common artifacts.
- › Lock aspect ratio and length from the start.

[LESSON 03]

Image-to-Video and Keyframe Workflows

The most reliable workflow in 2026 isn't text-to-video — it's image-to-video. Generate the perfect still in Midjourney or Flux, then animate it in Runway or Sora with motion-only prompts. You get full control over composition, character look, and brand consistency, while the video model only handles motion. This hybrid pipeline ships 5x more usable footage per hour than pure text-to-video.

For longer sequences, use keyframes: generate frame 1, frame 4, and frame 8 as stills, then animate the transitions. This gives you control over story beats while leveraging the model's strength at interpolation. Runway Gen-4's keyframe mode is the most mature; Sora 2 is catching up. Always export master files at the highest resolution available — you'll downscale for delivery but want headroom for grading.

// KEY TAKEAWAYS

- › Image-to-video > pure text-to-video for quality.
- › Keyframes give you story-beat control.
- › Export masters at max resolution; downscale for delivery.

[LESSON 04]

Sound, Music & Final Polish

Veo 3 generates synced dialogue and ambient audio natively — a genuine leap. For Sora and Runway clips, layer audio in post using Suno v5 (music), ElevenLabs v3 (voice), and a small library of Foley samples. AI-generated music is now indistinguishable from production music libraries for background use; reserve human composers for hero tracks where the music is the story.

Always color-grade. Even the best AI clip looks 'AI' until you take it through DaVinci Resolve or CapCut for a consistent grade across shots. Add subtle film grain, a slight contrast curve, and a unified color palette. This is the single highest-leverage post step — it can turn a passable AI reel into something a client signs off on without realizing it wasn't shot on a camera.

// KEY TAKEAWAYS

- › Veo 3 = native sync audio; layer audio for others.
- › AI music ≈ stock for backgrounds.
- › Color grading is the AI tell-killer.



[NEXT]

Keep Going

You've completed this course material. The real learning starts when you apply what you've read. Pick one idea from this PDF and run a small experiment this week. Document what worked and what didn't. Share your findings with the community.

Explore more free courses, daily AI tips, and curated tools at:

innovationailab.com

Have feedback or want to suggest a topic? We read every message.

hello@innovationailab.com

— Innovation AI Lab Team —

// part of LumiLife Tech