STRETCHLEARN · WORKBOOK

AI Video & Generative Content Creation — Workbook

This workbook is your production companion for the AI Video & Generative Content Creation course. Each section maps directly to a course module and contains hands-on exercises, structured worksheets, and action checklists to move you from concept to published content. Complete every section in order — the exercises build on each other, and by the final section you will have a working production system, a tested prompt library, and a live piece of published AI video content.

How AI Video Generation Works

Build your foundational vocabulary and diagnostic eye before spending any credits.

Exercise: Artifact Identification Drill

Find three publicly available AI-generated video clips (search YouTube for 'Runway Gen-3 sample' or 'Sora video demo'). Watch each clip twice — once for enjoyment, once on mute with the 10-second quality audit checklist from Lesson 3 in front of you. Log your findings below.

• For each clip, name every artifact type you observe (temporal drift, limb morphing, physics violation, motion blur incoherence, text illegibility, background instability). Be specific about the timestamp where it occurs.

• Which artifact appears most frequently across all three clips? What does this tell you about what to suppress in your own prompts?

• If you were hired to reprompt each clip to fix its most critical artifact, what single sentence would you add to each prompt?

Worksheet: Tool Stack Planning Sheet

Before you subscribe to anything, map your specific content goals to the tools that serve them. Fill in every row based on the content types you actually plan to produce in the next 30 days.

My primary content format (e.g. Instagram Reels, YouTube Shorts, LinkedIn)

My primary content category (e.g. product demo, educational, lifestyle brand)

Target clip length per video

Generation tool I will start with (and why)

Editing tool I will use

Voiceover tool I will use (or N/A if no voiceover needed)

Music source I will use

Estimated monthly budget (USD)

One tool I will NOT subscribe to yet (and why I am deferring it)

Checklist: Module 1 Readiness Checklist

[ ] I can define temporal drift and name one prompt suppressor for it
[ ] I can name the three layers of a video prompt (subject+action, camera+motion, aesthetic+mood)
[ ] I have compared Runway, Sora, and Kling on at least one capability dimension
[ ] I have created a free Kling account for cost-free prompt testing
[ ] I have identified my primary content format and matched it to a generation tool
[ ] I can explain why text in AI-generated video is always illegible and what to do about it

Prompt Engineering for Video Generation

Build your personal prompt library and validate your first formulas before spending significant credits.

Exercise: 5-Part Prompt Construction Workshop

Choose one specific video idea from your content plan. Write the same scene three times: once as a naive prompt (how you would naturally describe it), once using the 5-part formula from Lesson 4, and once with the artifact suppressors from Lesson 3 added. Generate all three on Kling's free tier. Compare and score each output on the 5-point quality scale.

• Write your naive prompt here (1–2 sentences as you naturally think of it). Then write your 5-part structured prompt using all five components. Show both.

• After generating, score each variant: subject consistency (1–5), motion quality (1–5), aesthetic match (1–5). Which prompt scored highest overall and by how many points?

• What did the structured prompt add that the naive prompt was missing? Be specific — name the exact words or phrases that drove the quality improvement.

• Now add artifact suppression language to your highest-scoring prompt and generate one final version. Did quality improve, stay the same, or decline? Note the result.

Worksheet: Image-to-Video Reference Asset Planner

Plan the reference images you need for your first image-to-video project. For each subject that must appear consistently across multiple clips, document what reference image you will create and how.

Project name or video series title

Subject 1 that needs consistent identity across clips

How I will create the reference image for Subject 1 (photograph, Midjourney, Adobe Firefly, etc.)

Target framing for the reference image (headshot, medium shot, full body)

Background type for reference image (neutral grey, white, specific environment)

Subject 2 (if applicable) that needs consistent identity

How I will create the reference image for Subject 2

Product or location that needs consistent visual treatment (if applicable)

Reference image source for that product/location

Checklist: Prompt Library Launch Checklist

[ ] I have set up a prompt library document (spreadsheet or Notion) with the six required columns: category, prompt, model+settings, output rating, notes, file reference
[ ] I have logged at least three tested prompts with full scores and notes
[ ] I have identified my highest-performing camera/motion formula for my primary content category
[ ] I have identified my highest-performing aesthetic/grade formula for my primary content category
[ ] I have run the 3-variant testing protocol on at least one new prompt concept
[ ] I am using Kling for prompt validation before spending Runway credits
[ ] I have at least one reference image ready for image-to-video generation

Editing, Audio, and Post-Production with AI Tools

Produce your first complete assembled video with voiceover, music, and captions from end to end.

Exercise: End-to-End Assembly Exercise

Using only the tools covered in this module (CapCut AI, ElevenLabs, Suno), produce one complete 30–60 second video from scratch. Start with at least three generated clips, write a voiceover script, generate the audio, generate a music track, and assemble everything in CapCut AI. Export and watch the result before answering these questions.

• What was the most time-consuming step in the assembly process? How could the workflow be made more efficient on the next video?

• Review your captions in the final export. How many errors did Auto Caption produce? What types of words were most commonly mis-captioned (proper nouns, technical terms, fast speech)?

• Describe the audio mix in your final video. Is the voiceover clearly intelligible over the music? If not, what specific level adjustment would fix it?

Worksheet: Script-to-Audio Delivery Tracker

For each voiceover segment you produce, fill in this tracker to document the ElevenLabs settings that produced the best result for your voice and content type. This becomes your audio style guide.

Script segment title or ID

Word count

Target duration (seconds)

ElevenLabs voice used

Stability setting (0.0–1.0)

Similarity setting (0.0–1.0)

Delivery rating (1–5)

Notes on what to adjust next time

Final WAV filename

Checklist: Post-Production Quality Checklist

[ ] Every AI clip has been trimmed 0.5 seconds from head and tail to remove generation warm-up and drift
[ ] Color Match has been applied across all clips to unify color temperature
[ ] Voiceover peaks at -6 dB; background music sits at -20 to -22 dB under narration
[ ] Auto Ducking is active on the music track during voiceover segments
[ ] Captions have been reviewed and corrected — proper nouns and brand names are accurate
[ ] No AI-generated text is legible in any clip (all text overlays are added in post)
[ ] Final export is at 1080p or higher with correct aspect ratio for target platform
[ ] SRT caption file has been exported alongside the video for platform upload

Publishing, Distribution, and Content Strategy

Launch your publishing workflow, set up your production calendar, and build the compliance habits that protect your content operation.

Exercise: Platform Optimization Audit

Take the video you assembled in Module 3 and prepare it for publication on two platforms. Do not publish yet — instead, complete every preparation step for both platforms and document what changed between the two versions.

• List every specification difference between your two chosen platforms (aspect ratio, maximum length, hashtag format, caption position, thumbnail requirements). How many differences are there?

• Write the title and first 100 characters of the description for each platform separately. How does the copy change between a YouTube Shorts audience and a LinkedIn audience for the same video?

• What disclosure language did you add (or decide not to add) to each platform, and why? Cite the specific policy or guideline that informed your decision.

Worksheet: Weekly Production Calendar Template

Map out your first four weeks of AI video production before you begin. Assign specific topics to specific pillars, estimate generation credit costs, and identify which slots are for testing new prompts versus executing proven ones.

Week number

Content Pillar 1 title

Content Pillar 2 title

Content Pillar 3 title

Content Pillar 4 title (or leave blank if using 3 pillars)

Video topic for this week - Pillar 1

Video topic for this week - Pillar 2

Video topic for this week - Pillar 3

Estimated Runway credits needed this week

Primary platform for this week's content

New prompt formula to test this week (or 'none')

Publishing days and times for each video

Checklist: Legal and Compliance Readiness Checklist

[ ] I have read and can summarize the commercial use terms for every tool in my stack (Runway, Sora/OpenAI, ElevenLabs, Suno)
[ ] I have a generation log template set up to document timestamps, prompts, and model versions for all commercial projects
[ ] I understand YouTube's synthetic media disclosure requirement and have added the required disclosure to my channel settings
[ ] I will not generate AI video featuring real, identifiable people without their explicit written consent
[ ] I have decided on my disclosure language for AI-generated content and it appears in my publishing checklist
[ ] I understand that AI-generated text visible in video will be garbled and I have a post-production step to add all text overlays manually
[ ] I have set my ElevenLabs, Suno, and Runway accounts to paid tiers before using outputs commercially

Your Action Plan

1. Create accounts on Kling (free), ElevenLabs Starter ($5/month), and Runway Standard ($15/month) — do not subscribe to more tools until you have completed your first 5 videos
2. Set up your prompt library document with the six columns and log your first three test prompts using Kling's free tier before spending a single Runway credit
3. Define your 3–4 content pillars and write 12 video topics (3 per pillar) before your first production week — front-load the creative decisions
4. Generate your reference images for any character or product that needs visual consistency across multiple clips, and store them in a dedicated project folder
5. Complete the end-to-end assembly exercise in Module 3 to produce your first full video — this is the most important practice step in the course
6. Export your SRT caption file from CapCut and verify caption accuracy before publishing — fix all proper noun and brand name errors
7. Prepare your first video for two platforms simultaneously to train yourself on the specification differences before they become habitual oversights
8. Add disclosure language to your publishing checklist as a non-optional step — treat it as technically required even on platforms where it is currently optional
9. After your first week of production, review your generation log and identify which prompt

formulas had the highest keeper rate — promote those to your core library
10. Set a 30-day review date to assess your cost per published video and identify which tool in your stack is delivering the least value relative to its cost