

Generative & AI Art Direction — Workbook

This workbook turns the course into reps. You will write briefs, run single-variable prompt experiments, build a reference and style library, lock a brand across a set, and finish a production-ready image with clean documentation. Work one section per module and finish with a small, on-brand image set and a prompt system you could reuse on real client work.

From Prompting to Art Direction

Shift from typing wishes to writing briefs, and learn to judge outputs against intent rather than delight.

Exercise: Pretty vs Right: A Judgement Drill

Pick one simple brief (for example, a hero image for a coffee roaster with space for a headline). Generate a batch of four to nine images in any tool. Then judge each strictly against the brief, not against how much you like it.

- For each image, write yes or no on whether it actually fits the stated purpose, subject, mood, and format.
- Identify your favourite image, then state honestly whether it is the right one for the brief or just the prettiest.
- Name one image that is technically beautiful but wrong for the brief, and say exactly why it fails.
- Pick the single best brief-fit image and list what still needs fixing before it could ship.

Worksheet: The Six-Line AI Image Brief

Fill this brief completely before generating anything for the project you will carry through this workbook. Each line should map directly to a block of your future prompt.

Purpose (where it lives and the job it does)

Subject (the literal thing in the frame)

Mood and tone (3 to 4 adjectives)

Style and medium (e.g. 35mm photo, flat illustration, 3D render)

Format (exact aspect ratio + minimum resolution for final use)

Constraints and exclusions (brand colours to hold, things to avoid, room for type)

Checklist: Before You Generate

- I have written a complete six-line brief for this image
- I know the exact aspect ratio the final use requires
- I know the minimum resolution the final medium needs (screen vs print)

- [] I have decided the three or four mood adjectives
- [] I have collected at least 6 reference images that point at the target look
- [] I can state what the image must NOT contain

Prompt Anatomy and Style Control

Build a structured prompt, learn the vocabulary dials by isolating them, and steer with parameters.

Exercise: Single-Variable Vocabulary Study

Keep one subject fixed and change only one block at a time so you learn what each term actually does. Use the same generator for all runs so the comparison is fair.

- Generate your subject three times, changing ONLY the lighting term (e.g. golden hour, then soft diffused light, then low-key chiaroscuro), and note how the mood shifts.

- Generate the same subject three times changing ONLY the medium (e.g. 35mm photograph, watercolour, 3D clay render).

- Generate the same subject twice changing ONLY the composition (e.g. overhead flat lay vs shallow depth of field close-up).

- Write one sentence per term describing the effect it had, so you build your own dial reference.

Worksheet: Structured Prompt Builder

Convert your six-line brief into a structured prompt by filling each block in order. Put the most important words first.

Subject (concrete, leads the prompt)

Medium (photo / illustration / render / painting)

Style and references (movement, technique, aesthetic)

Composition and lighting (framing + light direction/quality)

Mood and colour (feeling + named palette)

Parameters (--ar, --stylize, --chaos, --no, etc.)

Full assembled prompt (all blocks combined, in order)

Exercise: Parameter Sweep

Take your assembled prompt and run it across parameter settings to feel their effect. Record what each change does so you know which dial to reach for next time.

- Run the prompt at --stylize 50, then --stylize 250, then --stylize 750, and note how closely each follows your words vs adds its own flair.

 - Run the prompt at --chaos 0 and again at --chaos 50, and note the difference in how varied the four options are.

 - Set --ar to match your brief exactly and confirm the output is the right shape for its destination.
-

- Add a --no exclusion for something the model keeps inserting, and confirm it is reduced.

Checklist: Prompt Quality Pass

-] The subject and most important quality appear at the START of the prompt
-] Every block describes what is SEEN, not the abstract idea behind it
-] I used real photography or art vocabulary, not vague words like nice or cool
-] The aspect ratio parameter matches the brief
-] I can change one phrase to change one aspect of the image (the prompt is editable)
-] I avoided asking for accurate in-image text or precise object counts where the tool is weak

Style, Reference, and Brand Consistency

Choose the right tool, direct with references and seeds, and lock a brand across a whole set.

Worksheet: Tool Routing Decision

For your project, decide which generator fits and record why, especially the licensing reason. Confirm the current terms of the exact plan you are on.

Brief in one line

Chosen tool (Midjourney / DALL-E 3 / Firefly / other)

Reason for choice (aesthetic control / fast ideation / commercial safety)

Commercial use allowed on my plan? (Yes / No / Need to check)

Training-data / licensing risk note for this client

Fallback tool if the first does not deliver

Exercise: Reference and Seed Control

Practise directing by showing rather than describing, and lock variables so a set holds together. Use a tool that supports image and style references (e.g. Midjourney --sref and --seed).

- Build a small moodboard of 4 to 6 reference images and use one as a style reference (--sref) while your text supplies a different subject; note how the aesthetic transfers.
- Generate one image you nearly love, fix its --seed, then make ONE small prompt edit and confirm the result is a controlled variation, not a different picture.
- Tune style strength (e.g. --sw) up and down and note how strongly the reference look is applied.
- If using a recurring character or element, test a character/omni reference and note how consistent it stays across three images.

Worksheet: Brand Consistency Lock Sheet

Define the fixed elements that every image in your set will share, so only the subject changes. This becomes your reusable house-style template.

Master style string (medium + lighting + mood, reused every time)

Locked palette in words (exact brand colours named)

Style reference image(s) used for the whole set (--sref source)

Medium and light that must NOT vary across the set

Recurring element to keep stable (mascot / face / product) and how

What is allowed to change between images (usually only the subject)

Checklist: Set Coherence Gate

- All images in the set share the same medium (no mixing photo and 3D unintentionally)
- All images share the same lighting style and mood
- The same style reference or reference set fed every generation
- Brand colours are named identically in every prompt
- I have noted where colour will need exact correction in a design tool afterward
- A stranger could tell these images belong to one brand

Finishing and Integrating into a Real Workflow

Repair and upscale to a real standard, slot AI into a pipeline, and handle licensing, ethics, and client communication.

Exercise: Generation to Deliverable

Take your best brief-fit image and finish it to a shippable standard. Use inpainting, outpainting, retouching, and upscaling as needed.

- Use inpainting (Vary Region / Generative Fill / edit) to fix at least one flaw such as a bad hand, artefact, or melted detail.
- Use outpainting or Generative Fill to extend the frame so there is clean negative space for a headline.
- Correct the colour to the exact brand value in a design tool, since the prompt could not guarantee it.
- Upscale to meet the final medium (note target ppi and pixel dimensions) and check for invented texture you must retouch.

Worksheet: Asset Documentation Record

Document the kept asset so it is reproducible and the licensing is traceable. Fill one record per final image.
Project / asset name

Generator and plan used

Full prompt (exact text)

Parameters and seed

Reference / --sref images used

Final size (pixels) and ppi

Commercial licence status confirmed (Yes / No)

Checklist: Ship-Ready and Ethics Gate

- Visible flaws (hands, text, artefacts) are repaired
- Resolution and ppi meet the final medium (print vs screen)
- Brand colours, logos, and type are pixel-accurate, composited in a design tool
- Commercial use is permitted on my plan for every tool used
- No living artist's name, trademarked character, or real brand logo was generated to pass off as original
- The client has been told the imagery is AI-generated
- Prompt, seed, and references are saved so the asset is reproducible

Your Action Plan

1. Write a complete six-line brief for one real image before opening any tool.
2. Collect a moodboard of 6 or more references that point at the target look.
3. Convert the brief into a structured prompt, most important words first.
4. Run a single-variable study on lighting, medium, and composition to learn the dials.
5. Sweep --stylize and --chaos and set --ar to match the brief.
6. Route the job to the right generator and confirm its commercial licence on your plan.
7. Lock a master style string plus a style reference and build a small set where only the subject changes.
8. Fix the seed to refine your best image into controlled variations.
9. Finish the winner: inpaint flaws, outpaint space for type, correct brand colour, and upscale to the final medium.
10. Document each kept asset with its prompt, seed, references, size, and licence, and disclose AI use to the client.

