

Midjourney for Designers — Workbook

This workbook turns the course into a repeatable Midjourney workflow you can run for any brief. You will set up the tool, write a structured prompt from a real brief, tune it with parameters and references, lock a consistent set, refine and finish the winner, and assemble a client-ready presentation slide with honest licensing notes. Work through the sheets and templates on one real project so you finish with both a skill and a delivered, presented image.

Getting Into Midjourney and Thinking in Briefs

Set up your account and modes, then write a six-line image brief for a real project before you generate anything.

Exercise: Run Your First Single-Variable Test

Generate one simple subject, then rerun it changing exactly one word, so you feel how literally Midjourney reads you. Keep both grids side by side.

- Generate: a single ripe lemon on a pale blue background, studio photograph. What did the grid of four give you?

- Rerun changing only studio photograph to watercolour illustration. How completely did one word change the result?

- Now swap only the lighting word (e.g. soft diffused light to hard direct sunlight). Did the mood shift as expected?

Worksheet: Six-Line Image Brief

Fill this in for your real project BEFORE you write any prompt. Each line will map onto a part of the structured prompt later.

Purpose (where it lives and the job it does)

Subject (the literal thing in the frame)

Mood and tone (3-4 adjectives)

Style and medium (photo / 3D / illustration / era)

Format (exact aspect ratio + minimum resolution)

Constraints and exclusions (brand colours, what to avoid, room for type?)

Checklist: Setup and Readiness Pass

- Paid Midjourney plan active (commercial rights confirmed for client work)
- Logged into the web app and located the imagine bar
- Tried the Discord /imagine flow at least once so both interfaces are familiar
- Understand Fast vs Relax mode and which I am currently using
- Six-line brief written and saved for the real project
- Decided the final aspect ratio from the brief, not by default

Prompt Structure and the Vocabulary of Control

Convert your brief into a structured prompt, then practise the named vocabulary that gives you precise control over the look.

Worksheet: Structured Prompt Builder

Translate each brief line into a prompt block in order. Front-load the subject and your most important quality. Assemble the blocks into one prompt line at the bottom.

Block 1 - Subject (concrete)

Block 2 - Medium (e.g. 35mm film photograph)

Block 3 - Style and references

Block 4 - Composition and lighting

Block 5 - Mood and colour

Block 6 - Parameters (e.g. --ar 16:9)

Final assembled prompt (blocks joined in order)

Exercise: Vocabulary Dial Study

Pick one subject and hold it fixed. Generate it three times, changing only the lighting term each run, to build intuition for the vocabulary as dials.

- Run 1 with golden hour, Run 2 with soft diffused light, Run 3 with low-key chiaroscuro. Which mood did each produce?
-

- Which single lighting term best matched your brief's mood line?
-

- Now swap only the medium term (e.g. macro photograph to flat vector illustration). How much did the feel change?
-

Exercise: Clean Up With --no and Positive Phrasing

Take a prompt that keeps producing an unwanted element and fix it two ways: first by describing the positive scene more fully, then by adding the --no parameter.

- What unwanted element keeps appearing (people, text, clutter, extra objects)?
-

- Did fully describing the positive scene reduce it without using --no?
-

- What did adding --no [element] at the end change, and how clean is the result now?

Checklist: Prompt Quality Pass

- Subject and top quality are at the FRONT of the prompt
- Every block describes an appearance, not an abstract intention
- Used at least three named vocabulary terms (medium, lighting, composition)
- Prompt is focused, not a wall of forty adjectives
- Aspect ratio set with --ar to match the brief
- No request for readable text inside the image (type will be added later)
- Stubborn unwanted elements handled with --no

Parameters, References, and Consistency

Tune your prompt with parameters, direct it with reference images, and lock a consistent multi-image set.

Worksheet: Parameter Tuning Log

Record the parameter values you tried and the effect, so you learn what each dial does for your own work. Change one parameter at a time.

Aspect ratio (--ar) used

Stylize (--s) value tried

Effect of stylize (more literal / more flair)

Chaos (--c) value tried

Effect of chaos (tight variations / wide spread)

Weird (--w) value tried, if any

Best combination for this brief

Worksheet: Reference Direction Sheet

Plan your reference images and tag each to the right channel. Record the strength values so the result is reproducible.

Image prompt used? (what content/composition it brings)

Style reference image (--sref) source and what look it borrows

Style weight (--sw) value

Character reference image (--cref) source, if any

Character weight (--cw) value

What each reference is responsible for (content / look / character)

References and strengths saved with the prompt? (Y/N)

Exercise: Lock a Consistent Three-Image Set

Build a master style string (medium, lighting, palette, mood, plus a fixed --sref) and append a different subject for each of three images, changing nothing else. Aim for a set that reads as one campaign.

- What is your fixed master style string (everything except the subject)?
 - Across the three images, does only the subject change, or did the look drift?
 - Where did consistency still break (face, exact colour, rendering style), and would a seed or finishing fix it?
-

Checklist: Consistency and Control Pass

- Aspect ratio matched the brief on every image in the set
- Stylize tuned deliberately (lowered for accuracy, raised for flair)
- Chaos raised while exploring, lowered while refining
- Same --sref image fed to every image in the set
- Master style string reused with only the subject swapped
- Brand palette named in words in every prompt
- Seed captured and reused where tight variation was needed
- Prompt, parameters, seed, and references saved for each kept image

Refining, Finishing, and Presenting to Clients

Refine the winner with Midjourney's tools, finish it in a design tool, and assemble a curated, honestly-licensed client slide.

Exercise: Refine the Winner, Do Not Reroll

Take your best image and improve it using Midjourney's in-app tools instead of rerunning the prompt. Use Vary Region for a flaw, and Zoom Out or Pan if you need room for type.

- What flaw did you fix with Vary Region, and did it solve in one pass?
 - Did you need Zoom Out or Pan to create negative space for a headline? What changed?
 - Which upscaler (subtle vs creative) gave the cleaner result for your medium?
-

Worksheet: Finishing Handoff Sheet

Record what still must be corrected in a design tool, because Midjourney cannot guarantee these. Plan the handoff before delivery.

Final medium and required resolution (e.g. A2 print at 300 ppi)

Exact brand colour to correct to (hex)

Real type / headline to be added (font and content)

Logo to be composited (file and placement)

Artefacts to retouch out (where)

Design tool used for finishing (Photoshop / Illustrator / Figma)

Layered master file saved? (Y/N)

Worksheet: Client Presentation and Licensing Sheet

Prepare what you will actually show and say. Curate to two or three resolved options in context, and write the licensing line in plain English.

Number of options shown (aim for 2-3, in context)

Mockup context for each (website hero / poster / social)

Framing line (these are directions to be refined and finished)

Plan: subscription plan confirming commercial rights

Disclosure line (imagery is AI-generated, in plain words)

Ownership note (copyright is unsettled for purely AI images)

Iteration promise (what you can turn around fast vs not at all)

Checklist: Delivery and Presentation Pass

- Curated to 2-3 resolved options, not a wall of raw grids
- Each option shown in its real layout and format
- Winner refined with Vary Region / Pan / Zoom Out as needed
- Upscaled to the final medium and inspected for invented texture
- Colour corrected to exact brand hex in a design tool
- Real type and logo composited (not rendered by Midjourney)
- On a paid plan with commercial rights confirmed
- Client told, in plain words, that the imagery is AI-generated
- No mimicry of a living artist or trademarked character for commercial use
- Prompt, seed, and references archived as reproducible source

Your Action Plan

1. Activate a paid Midjourney plan, log into the web app, and confirm commercial usage rights for client work.
2. Write the six-line image brief for your real project: purpose, subject, mood, style, format, constraints.
3. Convert the brief into a structured prompt: subject, medium, style, composition and lighting, mood, parameters, front-loading the subject.
4. Run a single-variable test, changing one word at a time, to confirm Midjourney reads you literally.
5. Tune the result with parameters: set --ar to the brief, adjust --stylize for obedience, raise --chaos to explore then lower it to refine.
6. Direct content and look with references: an image prompt for composition, --sref for the

aesthetic, --cref for any recurring character.

7. Lock a consistent set by reusing one master style string and the same --sref, swapping only the subject, capturing seeds where needed.

8. Refine the winner in-app with Vary Region for flaws and Zoom Out or Pan for headline space, then upscale to the final medium.

9. Finish in a design tool: correct colour to the exact brand hex, composite real type and logos, retouch artefacts, save a layered master.

10. Build the client slide: curate 2-3 options in mockup context, write the AI-disclosure and licensing lines, and archive prompts, seeds, and references.

