STRETCHLEARN · WORKBOOK

AI Image Generation for Designers — Workbook

This workbook turns the course into reps you can do in Midjourney, Adobe Firefly, and Stable Diffusion. Each section maps to one module and mixes prompting drills, fill-in planning worksheets, and pre-delivery checklists. Keep a project folder and the prompt-log template open while you work, because the volume of generations gets unmanageable fast without one. Do the exercises in real tools, not on paper.

The Generative Image Landscape for Designers

Build a working mental model of diffusion and learn to choose the right tool for each design task.

Exercise: Same Brief, Three Tools

Write one short brief, for example a hero image for a sustainable coffee brand, and generate it in Midjourney, Firefly, and a Stable Diffusion interface using as close to the same description as each tool allows. Spend about twenty minutes and compare the results.

• Which tool produced the most polished aesthetic with the least effort, and why?

• Which output would you feel safest delivering to a paying client, and on what grounds?

• Where did each tool struggle: text, hands, composition, or color?

• For this specific brief, which tool would you actually use and what is your one-sentence reason?

Worksheet: Tool Selection Decision Sheet

For a real or invented project, fill in the constraints that decide the tool so the choice is deliberate rather than habitual.

Project name and one-line goal

Primary constraint: aesthetics, legal clearance, or precise control

Where the image will be used: internal, client deliverable, or public campaign

Final aspect ratio and minimum pixel dimensions required

Chosen tool and one-sentence justification

Subscription tier needed for commercial rights on that tool

Checklist: Workspace Setup Check

[ ] Midjourney account active on a plan that grants commercial use
[ ] Adobe account ready, with Generative Fill available in Photoshop if needed
[ ] A Stable Diffusion on-ramp chosen, hosted or local, and tested with one render
[ ] Project folder created with raw, selects, edited, and final subfolders
[ ] Prompt-log template open and ready to record prompt, seed, tool, and version

Prompting With Design Intent

Practice the structured prompt formula and learn to direct style and iterate reproducibly.

Exercise: Build a Prompt From Slots

Take a vague prompt like a nice product photo and rebuild it using the seven-slot structure: subject, medium, style, composition, lighting, color and mood, parameters. Generate both the vague and the structured version and compare.

• Write your weak version and your structured version side by side.

• Which single slot changed the result the most when you added it?

• Add the correct parameters for your tool: aspect ratio, version, stylization or CFG.

• Rewrite the structured prompt as a natural sentence the way Midjourney v6 and v7 prefer.

Exercise: Style Without Naming an Artist

Pick a look you admire and recreate it by describing visual qualities only, with no living-artist names. Use movement, era, medium, material, and lighting terms.

• Name the style you want in plain descriptive terms across at least three axes.

• Generate it, then refine the description until the look lands.

• If matching a reference, supply an image prompt or style reference instead of a name.

• Note why your description is both more original and more defensible than an artist name.

Worksheet: Reusable Prompt Recipe Card

Capture one prompt that worked so you can reuse and adapt it. Fill in each slot and the exact parameters.

Subject (concrete)

Medium

Style or era reference

Composition and shot type

Lighting

Color and mood

Tool, version, and exact parameters used

Seed of the best result

Checklist: Iteration Discipline Check

[ ] I picked the closest result, not an unrelated prettier one, before iterating
[ ] I captured and locked the seed of my chosen image
[ ] I changed exactly one clause or parameter per regeneration
[ ] I matched variation strength to my stage: wide early, subtle near the end
[ ] I logged prompt, seed, tool, and a one-line note for each saved image

Refining and Editing in Your Design Tools

Take raw generations the last mile with inpainting, upscaling, and compositing in real design software.

Exercise: The Flaw Pass

Take one of your best generations into Photoshop or a Stable Diffusion inpainting interface. Zoom to 100 percent, find every flaw, and fix each one locally without regenerating the whole image.

• List every flaw you found: warped text, extra fingers, odd objects, artifacts.

• Fix each with Generative Fill or masked inpainting and note the strength you used.

• Use Generative Expand to add negative space for a headline if the frame is tight.

• Compare before and after at full size and confirm nothing important changed.

Worksheet: Resolution and Output Planner

Plan the final resolution and format for a deliverable so you upscale once, correctly, at the end.

Final use: screen, social, or print

Final dimensions in inches or pixels

Required pixels at 300 ppi if print (inches multiplied by 300)

Upscaler chosen and target multiplier (2x, 4x)

Color space for delivery (sRGB or CMYK)

Files to deliver: layered working file plus flattened export format

Exercise: Composite Into a Layout

Place an AI image into a real layout in Figma, Photoshop, or InDesign as a hero with a headline and logo. Treat the image as one element under your normal design rules.

• Did you generate to the final aspect ratio and leave negative space for type?

• What color grade or filter did you apply to unify the AI image with the layout?

• Where did you set type yourself rather than letting the model render text?

• Is the AI placed as a smart object or linked layer so revisions stay easy?

Checklist: Edit and Integration Check

[ ] Flaw pass complete: image scanned at 100 percent and fixed
[ ] Upscaled last, after edits, to the correct final dimensions
[ ] AI image color graded to match the rest of the layout
[ ] Type and precise alignment done in the design tool, not the model
[ ] Editable working file saved with prompt log attached for revisions

Rights, Ethics, and Professional Practice

Clear the legal, ethical, and disclosure hurdles that decide whether an AI asset is safe to deliver.

Exercise: Rights Audit of a Real Deliverable

Take an AI image you intend to use commercially and audit whether it is safe to deliver. Check the tool, your plan, and the use case against the course rights rules.

• Which tool made it, and what is that tool's training-data and indemnification status?

• Does your current subscription tier actually grant commercial use for this use case?

• Is there any identifiable real person or copied living-artist style to remove?

• What is your risk verdict: safe as is, needs human editing, or switch tools?

Worksheet: Client AI Agreement Sheet

Fill this in at project kickoff so AI usage, tools, and disclosure are agreed in writing before you generate anything.

Client and project name

Is AI image generation permitted: yes, no, or limited

Approved tools for this project

Disclosure required: how and where AI use will be noted

Rights and indemnification that apply to the chosen tools

Real-person likeness or cultural references requiring extra permission

Checklist: Pre-Delivery Ethics and Rights Checklist

[ ] Commercial rights cleared for the tool and my subscription tier
[ ] No identifiable real person generated without explicit consent
[ ] No deceptive deepfake or fabricated documentary-style imagery
[ ] Bias and demographic representation reviewed and corrected deliberately
[ ] Cultural references handled with research and respect
[ ] Flaws inpainted and resolution correct for the final use
[ ] Provenance handled and AI use disclosed per the client agreement

Your Action Plan

1. Set up all three tools and create a project folder with raw, selects, edited, and final subfolders plus a prompt log
2. Run the same brief through Midjourney, Firefly, and Stable Diffusion and write a one-line rule for when to use each
3. Adopt the seven-slot prompt formula and build three reusable recipe cards from prompts that worked
4. Practice describing style by movement, medium, and material instead of naming living artists
5. Make seed-locked, one-change-at-a-time iteration your default and log every saved generation
6. Run a flaw pass with Generative Fill or inpainting on every image before calling it done
7. Plan resolution up front and upscale once at the end to the correct final dimensions
8. Composite AI imagery into real layouts with your normal grid, type, and color-grading discipline
9. Audit commercial rights for the tool and plan before any AI image enters paid work
10. Agree AI usage with each client in writing at kickoff and run the pre-delivery ethics and rights checklist on every asset