A controlled workflow for generating and editing images

This practical reading turns Lecture 3 into a repeatable experiment. You will plan a small image grid, compare one setting at a time, and repair a local error while preserving the useful parts of an image.

Assigned technical reading

Hugging Face Diffusers documentation • current pages checked 15 September 2026. Read the sections below; running their Python examples is optional. The exercise uses controls found in compatible diffusion tools and does not require a particular interface.

Allow 20 minutes for selected sections and 20–30 minutes for the exercise. If generation is unavailable, prepare the experiment plan and annotate a grid from class.

What each control changes

ControlNet supplies spatial information such as edges, depth, or pose through an additional network. Its conditioning scale controls the contribution of that input. The control model must be compatible with the base model and expected control type. An edge map constrains boundaries; it does not by itself specify material, identity, or the meaning of a scene. [1]

A seed helps control randomness. Diffusers passes a Generator object whose state changes when used; resetting it matters when comparing repeated calls. A seed alone does not guarantee identical results across hardware, software releases, or pipelines. Record the environment as well as the number. [2]

Image-to-image begins from an existing image. In the documented diffusion pipelines, greater strength adds more noise and allows a larger departure from that starting image. Guidance adjusts text conditioning and is a separate control. The useful range depends on the model and pipeline. [3]

Inpainting uses a mask to select the area to regenerate. In the documented convention, white marks the editable area and black marks the area to keep. Check the particular tool’s convention. Compare the resulting image with the original, including the boundary and supposedly preserved regions. [4]

Plan the image before choosing settings

Classroom exercise: make a poster image for a ceramics exhibition. Show one blue bowl on a table, leave the upper third clear for a title, and use side lighting that makes the rim readable. Add the final title as editable typography after image generation. These are exercise choices, not a tested recipe or a claim about a model's performance.

Draw a simple layout with the bowl's position and the empty title area. Write two requirements that must stay fixed and two appearance choices you may vary. This keeps composition separate from surface detail.

Build a six image comparison

Use a compatible ControlNet pipeline. Set the prompt, input sketch or edge map, image size, model version, sampler, steps, and guidance once. Choose three distinct seeds. For each seed, generate one image with lower spatial conditioning and one with higher spatial conditioning. Select supported values from the tool's documentation and record the actual numbers. Reset the random generator for each matched pair.

Place the results in three rows by two columns. Each row uses one seed; each column uses one control strength. Compare across a row to inspect the control change. Compare down a column to see how much the result varies across seeds. Do not switch the model or rewrite the prompt halfway through the grid.

Record the experiment

Record	What to save
Inputs	Exact prompt, negative prompt if used, sketch or control image
Model	Base model and revision, control model, any LoRA or adapter
Sampling	Seeds, sampler, steps, guidance and control strength
Environment	Tool and version, device if available, image dimensions
Outputs	All six images, run time, failures and later edits

Judge the images with visible criteria

For each image, record whether the title area is clear, the bowl count is correct, and the rim separates from the background. Then describe material and lighting quality. Choose a result using those observations. Keep a failure in the record: a selected image alone cannot show whether a setting works consistently.

Repair one local error

Suppose the best composition has a broken bowl rim. Save the selected image before editing. Mask the damaged area with enough surrounding context to form a plausible repair. Write an instruction about the rim. Start with the tool's documented settings, then compare a small number of edit strengths. Keep the original available for a direct before-and-after check.

Inspect the repaired contour, the light on the rim, and the texture at the mask boundary. Also inspect the table and empty title area. If unchanged pixels are essential, verify them or composite the selected edited region over the original in an image editor. A mask is an instruction to a pipeline; it is worth checking what the output actually preserves. [4]

Explain the mechanism

Answer these questions using your grid and repair. Why can different seeds produce different bowls? How does a control image differ from a text instruction? Why might a stronger edit repair the rim but also change its material? What information would another student need to repeat your comparison?

Connect the workflow to the safety debate

This is a low-stakes experiment, but its record makes a decision inspectable. A reviewer can see the inputs, alternative outputs, changed setting, and reason for selecting a result. A system that publishes images automatically would need an additional decision about permission and responsibility. Image quality tests alone do not answer that question.

For an exhibition poster, use references you are permitted to use and keep a record of their origin. If a generated image could be mistaken for a photograph of an actual exhibited object, decide how viewers will be told what it represents. Treat these as design and communication decisions, rather than assuming an attractive image settles them.

Bring to class

- One labelled six image grid, or a clear experiment plan if you cannot run the tool.
- One before-and-after repair with its mask, if editing is available.
- A settings record and a 150 word explanation of the tradeoff you observed. Report a failed attempt as well as your selected result.

Connection to Lecture 3

Review slides 48–64 for seeds, guidance and sampling; slides 65–72 for adaptation, spatial control and editing; and slides 90–98 for fair comparisons. The workflow above is a classroom synthesis of the four linked technical readings, not a benchmark result.