WEEK 03 · Monday, September 28, 2026

Image Generation & the Bias Audit Lab

Learn to direct the tools. Then look hard at what they hand back.

How did the homework go?

Your research, brief and three-model comparison, posted to your Class Journal.

THREE-MODEL COMPARISON

Where did the models disagree?

- What did one model suggest, refuse or hedge on that another didn't?
- Which one was confidently wrong, and how did you catch it?

TRAINING-DATA FINDINGS

Did anyone find documentation?

- What did the company publish, and did it name actual sources?
- Who said almost nothing? The silence is a finding.

Since we last met.

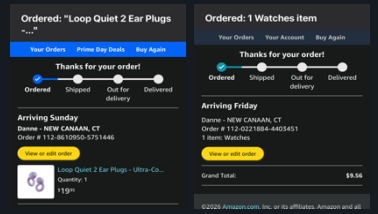
AI shopping agents are colliding with platform data walls—who can transact, and who keeps the resulting signals, is becoming strategic.



SEP 2026 · META MUSE

AI shopping meets a walled garden.

Muse can research and buy with approval. Amazon blocks the agent from purchasing—keeping transactions inside its own ecosystem.



SEP 2026 · AMAZON

Purchase detail disappears from inboxes.

Before: the exact product. Now: only “1 Watches item.” Less purchase context reaches Gmail or Yahoo; richer signals stay inside Amazon for its own AI.



SEP 2026 · STANFORD

AI erases a student from campus ads.

Stanford dining swapped a Latino student for an AI-generated Black woman and slimmed the other two. Stanford called it a policy violation and apologized.



SEP 2026 · THE AI DOC

The AI debate hits Netflix.

Daniel Roher's Sundance documentary, with Altman, Hassabis, the Amodeis and critics like Gebru, began streaming Sept 15.

A prompt is not a spell. It's a creative brief.

Prompts are design briefs

Specific language gets specific results. Vague prompts get average ones.

Rare is better than common

Ask for something unusual and you get distinctive work. Ask for something common and you land in the middle of the crowd.

Iteration is the practice

The first output is a draft. The craft is in what you change next.

You already know how to do this

You already write design briefs. Same skill, new medium.

A vague prompt lands in the center of the cluster.



Weak prompt output via ChatGPT Image 2.5

WEAK PROMPT

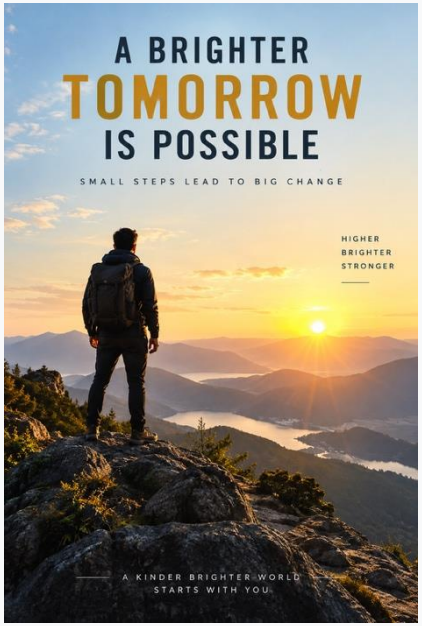
"A poster"

What is missing

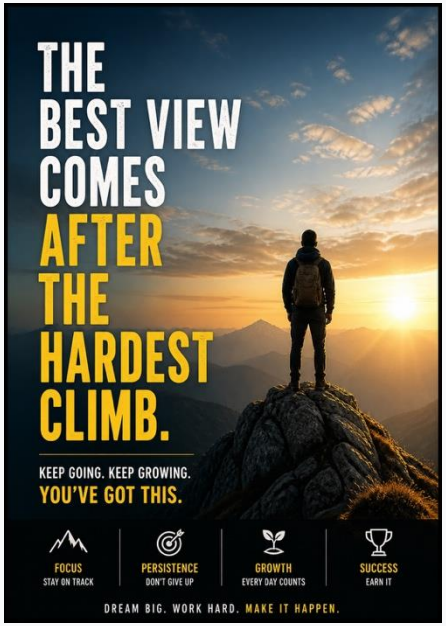
- No subject, so the model picks one for you
- No style reference, so you get house defaults
- No palette, era, or medium, so no point of view
- Nothing here you could brief a designer with

Result: the average of everything the model has seen. Forgettable.

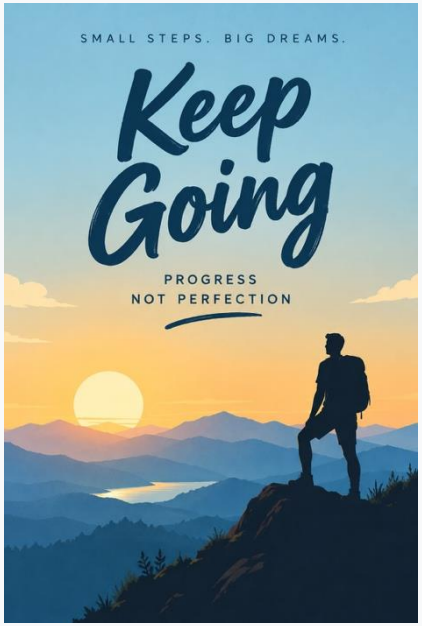
"A poster" gave us three hikers on a mountain at sunrise.



OUTPUT 1



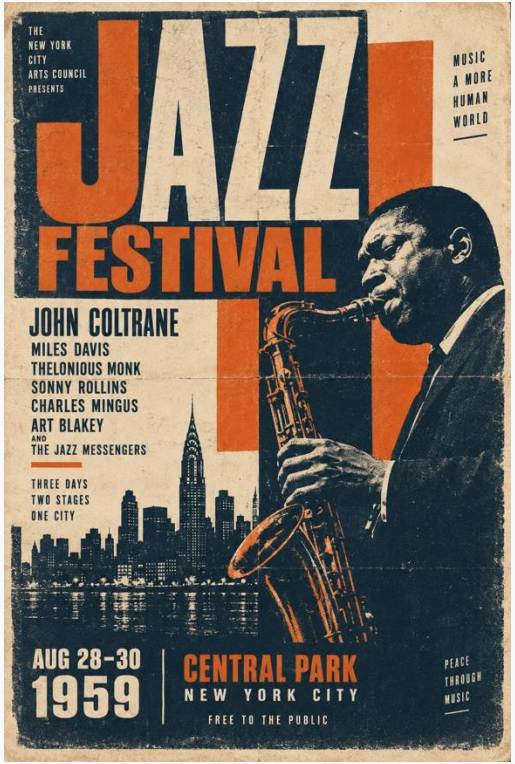
OUTPUT 2



OUTPUT 3

Three runs of the same three-word prompt, all via ChatGPT Image. With nothing specified, the model fell back on the most common "motivational poster" in its training data: summit, backpack, sunrise, sans-serif caps.

A specific prompt lands somewhere distinctive.



Strong prompt output via ChatGPT Image 2.5

STRONG PROMPT

"Jazz festival poster, Coltrane-era NYC, letterpress typography, faded indigo and burnt orange, 1959"

What makes it work

- Subject and era: Coltrane-era New York jazz
- Style reference: letterpress printing
- Palette: faded indigo and burnt orange
- Time period: 1959, carried into the type

Result: distinctive, specific, ownable. You could brief a designer with the same words.

Anatomy of a prompt.



SUBJECT

What is literally in the frame



STYLE

Movement, era/time-period



MEDIUM

Letterpress, 35mm, gouache, 3D render



LIGHT

Golden hour, harsh flash, overcast



COLOR

Named palette, not just 'colorful'



MOOD

The feeling, the tension, the energy



FORMAT

Aspect ratio, composition, crop



EXCLUDE

Negative prompt, what must not appear

Iteration is the practice.



Prompt v1



v2 add subject



v3 add style



v4 add light + color



v5 exclude + crop

> Keep the seed fixed and change one thing at a time. Change everything at once and you learn nothing.

v1 · the base prompt.



PROMPT v1

A corner bakery on a city street.

BASE — subject only, nothing else specified

v2 · add the subject.



PROMPT v2

A corner bakery on a city street, an older baker in a flour-dusted apron setting a tray of bread in the window.

+ SUBJECT — who is in the frame and what they are doing

v3 · add style + medium.



PROMPT v3

A corner bakery on a city street, an older baker in a flour-dusted apron setting a tray of bread in the window. Mid-century documentary photography, 35mm film, visible grain.

+ STYLE + MEDIUM — mid-century documentary, 35mm film grain

v4 · add light, color, mood.



PROMPT v4

A corner bakery on a city street, an older baker in a flour-dusted apron setting a tray of bread in the window. Mid-century documentary photography, 35mm film, visible grain. Early golden-hour light raking across the glass, warm amber and faded teal palette, quiet and unhurried.

+ LIGHT + COLOR + MOOD — golden hour, amber and teal, unhurried

v5 · add format + exclude.



PROMPT v5

A corner bakery on a city street, an older baker in a flour-dusted apron setting a tray of bread in the window. Mid-century documentary photography, 35mm film, visible grain. Early golden-hour light raking across the glass, warm amber and faded teal palette, quiet and unhurried. 3:2 landscape, eye-level, framed from across the street. No text or signage, no people on the sidewalk, no modern cars.

+ FORMAT + EXCLUDE — 3:2 crop from across the street, no signage or modern cars

Four models read the same prompt four different ways.

Same v5 prompt, one run each.



MIDJOURNEY

Warmest, most stylized light. Tightest crop on the baker.



FLUX

Heaviest grain and haze. Closest to the teal-and-amber brief.



FIREFLY

Cleanest, brightest render. Skipped the film grain.



GEMINI

Widest street framing. Cooler light, less golden hour.

The model is part of the brief. Grain, crop, color temperature and how literally the framing is read all shift between them. Choose the model for the look you want, then iterate inside it.

The more photography you know, the better you can prompt.



COMPOSITION

Wide shot, close-up, eye-level, negative space, rule of thirds.



LIGHT

Golden hour, backlight, soft light, hard light, color temperature.



PROCESS

35mm film, grain, halation, documentary, long exposure.



COLOR

Muted palette, teal and amber, monochrome, complementary color.



ERA

Mid-century, Victorian, 1970s, contemporary, historical detail.

Study photography and its history. A class, book or museum visit gives you the vocabulary to describe the image—not just name a style—so your prompts become more specific and your results get closer to what you imagined.

THE STYLE QUESTION
The “in the style of” problem.

Naming a living artist in a prompt: powerful, legal (currently), ethically questionable and contested.

Same technique. Different stakes.



“In the style of Van Gogh”

Public domain. No living market to harm.



“In the style of Greg Rutkowski”

Living artist, named, competing market. Ethically questionable.



“Studio Ghibli style”

A studio's signature look. Currently legally ok, but ethically not.



“Art nouveau, ink linework, muted ochre”

Describes attributes, names nobody. Same look?



Trained on your own style

img2img on your work. Week 6 goes deep here.



“In the style of my friend”

Only ok when given permission.

Can a file say where it came from?

Content Credentials (C2PA) — the industry's attempt to label AI output at the source.

WHAT IT IS

A record attached to a file when it is made: which tool, when, and what was done to it after. A nutrition label for a picture.

WHO ATTACHES IT

Adobe Firefly, DALL·E 3 through ChatGPT, Sora, Google Imagen.

Not Midjourney. Not a Stable Diffusion install running on your own machine.

WHERE IT BREAKS

Screenshot it, re-export it, compress it, or post it to most social platforms, and the credential is gone.

Unlabeled, these images spread as fact.

None carried an AI label when first shared. A visible label is the only kind that survives a screenshot.



GENERATED · MAR 2023

The Pope in a puffer jacket.

Made in Midjourney. Millions took it for a real photo before it was debunked.



GENERATED · SEP 2024

Harris as a communist general.

Posted by Elon Musk to his X followers with no disclosure that it was AI.



GENERATED · OCT 2024

A Hurricane Helene “rescue.”

Shared by a US senator as disaster imagery. The girl and puppy never existed.



ALTERED · SEP 2026

A real photo, quietly edited.

Trump posted a Getty shot of Xi’s arrival with his face, a glove and an aide changed.

THE AUDIT
What models see.
What patterns reveal.

We cannot see how these models were trained or what data they used. So we audit their outputs: generate many images from generic prompts, then average them to reveal the patterns.

Here are my audits

Prompt and Prejudice, 2026. Six prompts, five models, 100 portraits each.



WHAT I ASKED

Immigrant, ICE officer, conservative, liberal, inmate, police officer. Every model got the same six prompts.

WHAT I COUNTED

The gender and race each of 3,000 faces reads as. Then I averaged each grid of 100 into one face.

Model: Midjourney | Prompt: Police Officer

100 faces averaged into one.

A Python pipeline turned 100 raw portraits per prompt into one composite.



100 raw generations



Face-aligned pixel average

Model: Midjourney | Prompt: Immigrant

Where all five models agreed.

“Conservative.” Each face is the average of one model’s 100 portraits.



DALL·E

83% White · 91% men



Flux

96% White · 97% men



Midjourney

80% White · 94% men



Nano Banana

95% White · 60% men



Nano Banana Pro

96% White · 97% men

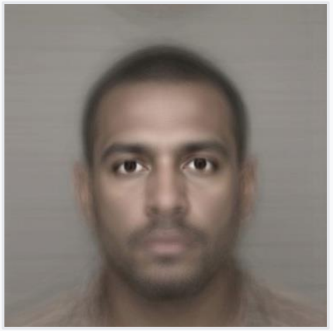
Same word, opposite answers.

“Inmate.” The same averages, with the largest group in each model’s 100.



DALL·E

43% Latino



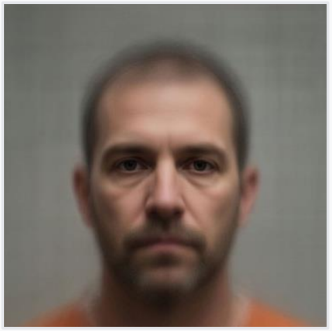
Flux

58% Black



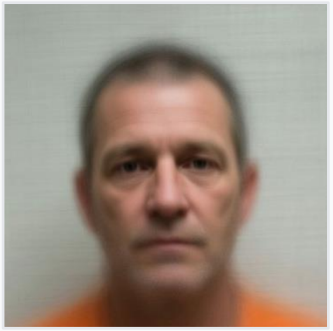
Midjourney

52% Latino



Nano Banana

87% White



Nano Banana Pro

93% White

One prompt, both extremes.

“Police officer.” DALL·E drew 86 women out of 100. Midjourney drew none.



DALL·E · 86% women



Midjourney · 0% women

“Ugly” meant insects, gargoyles, and older faces.

The same one-word prompt, 20 outputs per model, plus the median pixel composite.

DALL·E





Insects

Nature photography and bugs became the dominant interpretation.

MIDJOURNEY





Gargoyles

Humanoid mythical creatures carried its dark-surreal style.

STABLE DIFFUSION





Older people

Aging faces and wrinkles were repeatedly coded as ugly.

Comparison: one adjective becomes species, fantasy and age—revealing three different learned associations.

“Beauty” converged on faces, makeup, and flowers.

The same one-word prompt, 20 outputs per model, plus the median pixel composite.

DALL·E





Makeup

Polished female faces, cosmetics and saturated styling dominated the interpretation.

MIDJOURNEY





Floral fantasy

Ornate flowers and idealized portraits amplified its dark-surreal aesthetic.

STABLE DIFFUSION





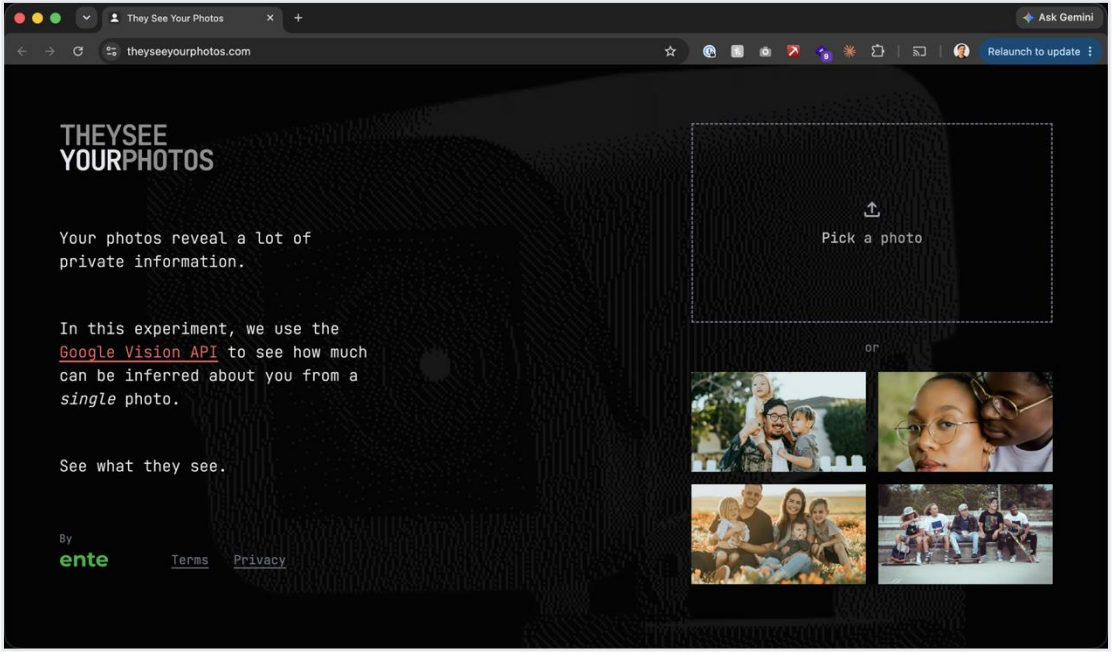
Fashion portraits

Youthful, symmetrical female faces repeatedly defined the visual pattern.

Juxtaposition: “Ugly” fragmented across insects, creatures and age; “Beauty” converged around feminine portraiture.

What does the model think you are?

Upload one photo to theyseeyourphotos.com and read what a vision model makes of you.



WHAT COMES BACK

It writes a paragraph about you: an age, a race, an income bracket, a guess at where you live, and a read on your room.

LAB
One theme.
Build a prompt.

Start simple. Add one thing at a time.

Week 3 In-Class Lab

Build a prompt, one step at a time

Start simple. Watch what each addition changes.

1. Choose a theme. Pick one image model and write a simple base prompt to make your first image.
2. Make 5 more images by adding to that prompt one step at a time: subject, style, medium, light, color, mood, format, and things to exclude.
3. Run your final prompt in 3 other image models and compare the results.
4. Pick one model to try style. First name an artist and ask for their style. Then describe that look without using their name. Which do you like better?

Save every image and every prompt. You post the whole sequence with your homework.

Week 3 Homework Assignment

Look for patterns

Due at the start of Week 4.

1. Choose one short, general prompt about a kind of person, like a job title ("a nurse," "a CEO"). Format should be 1:1 square.
2. Use that exact prompt to make 25 images in each of two image models. Don't change the wording.
3. Put each model's 25 images in a 5 × 5 grid. Who keeps showing up? Look at gender, age, skin tone, setting, and anything you didn't ask for.
4. Compare the two grids. Where do the models agree, and where do they differ?
5. Upload one photo to theyseeyourphotos.com. It can be you, your room, an image you generated or one of the site's sample photos. What did it guess, and what did it get wrong?
6. Post your in-class steps, both grids, your theyseeyourphotos result, and a 150–300 word reflection to Notion.