AMCC 5160 · Lecture 10 readings

All three connect directly to tonight's stories: one on reading an AI image as data, one on how an aesthetic score was built, and one on a real museum dispute. For Quiz 10, reflect on any one of the three. Optional readings are listed last.

Reading 1 / Art and technique: How to Read an AI Image

Eryk Salvaggio · Cybernetic Forests | essay — https://www.cyberneticforests.com/news/how-to-read-an-ai-image

What to do: Read the essay and follow its method on one AI image of your own, from Assignment 2 or anywhere else. Skip no steps: start literally.

Summary: Salvaggio argues that every AI-generated image is an infographic about its dataset. He generates a set of images of people kissing, describes what repeats and what is missing, then looks at the training data to explain it: the data contains stock photographs where actors were asked to kiss. His course, Critical Topics: AI Images, starts from this idea, and Hamilton College has built an art history course around it.

Questions to keep in mind:

- What did you see too often in your image? Could the dataset explain it?
- What is missing from your image, and who would notice?
- Where is the line between what the data did and what the artist did?

Connection: Story 1: "Every AI image is an infographic about the dataset" and "How to read an AI image".

Reading 2 / Technique: LAION-Aesthetics

LAION · LAION blog — https://laion.ai/blog/laion-aesthetics/

What to do: Read the post (short) and look closely at the example images grouped by score. You do not need the code.

Summary: LAION trained a small model on CLIP image embeddings to predict how much people like an image, using human ratings from sources including AVA (photos from the DPChallenge contest site). They scored LAION-5B and released subsets by score. The Stable Diffusion v1 team trained on the subset scoring 5 or more, about 600 million images. The team chose a linear model partly because its results looked best to them.

Questions to keep in mind:

- Whose taste do these ratings capture? Who is not in the room?
- Look at the images scored highest. Do you agree? What is missing?
- Tonight Vera Molnár scored below the cut-off and a recoloured Van Gogh beat the original. What would you change, if anything?

Connection: Story 2: “How taste got into Stable Diffusion” and “We ran that predictor on art history”.

Reading 3 / Culture: Madison art museum debuts new exhibition using generative AI amidst pushback

Wisconsin Public Radio · 25 September 2026 — https://www.wpr.org/news/madison-art-museum-new-exhibition-generative-ai-pushback

What to do: Read the report (about 5 minutes) and look at the photographs. Note every reason given on each side.

Summary: The Madison Museum of Contemporary Art opened Beeple’s Diffuse Control, 12 synchronized screens that transform works from the museum’s collection through a diffusion model. A local artist’s open letter, signed by more than 900 people, objected that artists had not given explicit consent and that diffusion models are trained on scraped art. The director replied that one can listen and hear but still not agree, and warned against blanket policies around artists.

Questions to keep in mind:

- The letter makes two different complaints. Which did the museum answer?
- Which of our four questions (AI? good? fair? allowed?) does each side answer?
- You run the museum. What would you do differently, and what would you keep?

Connection: Story 3 and tonight’s debate.

Optional / Art: The Power of Patience

Jennifer L. Roberts, Harvard Magazine · November–December 2013 — https://www.harvardmagazine.com/2013/11/the-power-of-patience

Three hours with one painting. If you read only one optional piece, read this.

Optional / Culture: An open letter to Pope Leo XIV in defence of AI-generated art

Sarp Kerem Yavuz, The Art Newspaper · 27 October 2026 — https://www.theartnewspaper.com/2026/10/27/comment-open-letter-pope-leo-xiv-artificial-intelligence

An artist’s reply to the Pope’s 2 October address to artists.

Optional / Art: In Search of the Future, Artists Give New Life to Dead AI Models

Shanti Escalante-De Mattei, ARTnews · 12 October 2026 — https://www.artnews.com/art-news/reviews/1234801925-1234801925/

Post Generation: three artists respond to an archive from the now-defunct Veo 1.

Optional / Technique: Evaluating Generative Models (lecture slides)

Jun-Yan Zhu, CMU 16-726 · Lecture 20, Spring 2025 — https://learning-image-synthesis.github.io/sp25/static_files/lectures/Lecture20_evaluating_generative_models.pdf

The full version of tonight's FID, FVD and Elo slides.