What Generative AI Reveals About Art History Pedagogy

Samantha Chang

Samantha Chang is a PhD Candidate in the Department of Art History at the University of Toronto, where she is also the Team Lead, Pedagogy and Educational Technologies at the Centre for Teaching Support & Innovation.
–samantha.chang@utoronto.ca

In the months following the public emergence of generative artificial intelligence, conversations in art history classrooms, faculty meetings, teaching centres, and conference discussions were often framed in defensive terms. How do we stop students from using generative AI to write their papers? How do we detect it? How do we regulate it? These questions are understandable. They respond to real anxieties about authorship, assessment, and the future of disciplinary expertise. But they are also incomplete. They assume that the integrity of art history resides primarily in the production of written artifacts and that the principal challenge posed by generative AI lies in the outsourcing of prose rather than in deeper epistemological questions about how knowledge is made, demonstrated, and evaluated.

What generative AI has made visible is not simply a problem of misuse, but a longer-standing problem of design. Many of our most familiar assessments, particularly when a single genre such as the research essay functions as the primary or sole measure of learning, have been unevenly aligned with how students learn, how art history produces knowledge, and how equity is enacted or denied through assessment practices. This is not an argument against the research essay, which remains central to art historical training and scholarly practice. Rather, it is an argument for examining how assessment ecologies come to rely too heavily on any one form, often without revisiting the assumptions embedded in its use. Generative AI has not created these misalignments; it has exposed them by making visible the gap between what many assessments appear to measure and the forms of disciplinary understanding they are meant to elicit.

This polemic reflects on a course assignment I developed in response to these conditions, the Making Connections Essay. The assignment does not ask students to generate text or images using generative AI. Instead, it asks them to analyze instructor-prompted, AI-generated images using art historical methods. In doing so, it treats generative AI not as a writing aid or a technological novelty, but as an object of critique, one that allows students to demonstrate synthesis, judgment, and historical thinking through a different configuration of visual analysis, textual engagement, and interpretive reasoning.

What I argue here is not that generative AI offers a solution to the challenges facing art history pedagogy. Rather, I suggest that generative AI functions as a disciplinary stress test. By revealing epistemic slippages between visual plausibility and historical meaning, between stylistic patterning and contextual knowledge, generative AI forces us to confront uncomfortable truths about what our assignments actually measure. Assessment practices that rely heavily on standardized prose, linguistic fluency, and surveillance-oriented enforcement have long been fragile as indicators of disciplinary understanding. By contrast, assignments that foreground connection-making, interpretation, and visual reasoning bring different aspects of art historical thinking into view and invite a reconsideration of what, precisely, we value as evidence of learning.

THE MAKING CONNECTIONS ESSAY: DESIGN AND INTENT

The Making Connections Essay was introduced as the final, summative assignment in a second-year undergraduate art history course focused on seventeenth-century Western European art. The course has no prerequisites and enrolls a heterogeneous group of students: majors and non-majors, students with strong writing backgrounds alongside those with more limited confidence in academic prose, and students well acquainted with early modern art as well as learners encountering the period for the first time. Throughout the term, students develop skills in close looking and close reading, engage with primary and secondary sources, and analyze a range of art objects through lectures, in-class discussions and exercises, readings, and audiovisual materials.

For the Making Connections Essay, students select one of two AI-generated images provided by the instructor. These images are generated using

prompts informed by course materials, including lecture content, class discussions and exercises, assigned readings, and artworks analyzed in class, and are intended to resemble seventeenth-century artworks without reproducing any specific historical object. Students are then asked to write an essay that interweaves multiple forms of analysis: formal analysis of the AI-generated image, comparison with course artworks, engagement with assigned readings, and connections to broader course themes. No independent research is permitted, ensuring that all students work from a shared set of materials and that the assessment foregrounds synthesis rather than accumulation. The rubric emphasizes clarity of argument, historical grounding, and the quality of connections made across visual and textual sources.

Several aspects of this design are worth emphasizing. First, the AI-generated images are produced by the instructor rather than by students, ensuring that all students work from the same curated visual materials and shared objects of analysis. Second, writing remains central. Students must articulate a thesis, structure an argument, and write clearly and persuasively. What changes is not the presence of writing, but its epistemic role. Prose is no longer treated as a proxy for disciplinary competence; it is one medium through which students demonstrate historical reasoning, interpretive judgment, and the ability to synthesize ideas across materials.

The assignment emerged not as a technological experiment, but as a disciplinary intervention. It responds to a familiar pedagogical challenge in art history, helping students move beyond isolated observations toward more synthetic thinking. By asking students to analyze an image that is deliberately unmoored from a specific historical referent, the assignment requires them to supply that context themselves. To explain what does not work in an AI-generated image, such as why a composition feels incoherent, why a devotional scene fails to align with seventeenth-century conventions, or why a figure's representation reproduces racialized stereotypes, students must draw on course knowledge in an integrated and historically situated way.

Across multiple iterations of the assignment, recurring pedagogical patterns began to emerge. Some students initially approached the images as exercises in identifying errors or inconsistencies, while others struggled to move beyond description toward historically grounded interpretation. Over time, it became increasingly important to model how formal analysis, historical context, course readings, and broader thematic concerns could be brought into conversation rather than treated as separate tasks. As the generated images became more visually convincing, students were increasingly required to articulate not simply what seemed incorrect, but why apparent coherence did not fully sustain historical meaning. The assignment thus shifted from an exercise in identifying isolated inaccuracies toward a more demanding practice of historically situated judgment.

GENERATIVE AI AS A DISCIPLINARY STRESS TEST

I use the term "disciplinary stress test" to describe moments when a new technology exposes the assumptions, shortcuts, and fragilities embedded in how a discipline teaches and assesses knowledge. In the case of generative AI images, the stress test is epistemological. These images often appear generically plausible at a superficial level, even when they register immediately as implausible to art historians. They draw on recognizable stylistic cues while remaining riddled with epistemic slippages: generalizations that flatten historical specificity, iconographic inconsistencies, and visual motifs that gesture toward meaning without sustaining it within a particular historical framework.

In early iterations of the assignment, these slippages were often obvious, such as distorted anatomy, anachronistic costumes, and inconsistent lighting (fig. 1). Over time, however, the AI-generated images have become more visually

fig. 1 AI-generated image created with Microsoft Copilot on January 1, 2024, for an early iteration of the Making Connections Essay assignment. Produced through prompts referencing seventeenth-century Northern European genre painting, portraiture, landscape, and still-life traditions, the image combines recognizable stylistic conventions with epistemic slippages, including bodily distortions, unstable object relationships, ambiguous animal forms, and the repeated presence of corgis within a domestic interior. Although the image draws on familiar visual conventions associated with seventeenth-century Northern European painting, its visual logic becomes increasingly difficult to situate within historically grounded conventions of representation and material culture.

convincing, even as my prompting practices have remained largely consistent. This improvement has not diminished the pedagogical value of the task. On the contrary, as surface coherence increases, the distinction between stylistic plausibility and historical meaning becomes more difficult, and more necessary, to articulate. What changes is not the presence of epistemic slippage, but its subtlety, which requires students to articulate more precisely why an image that appears coherent still fails to produce historical meaning.

An unfamiliar seventeenth-century artwork could also support interpretive and comparative analysis. However, AI-generated images introduce a different pedagogical condition: they are not simply unfamiliar objects, but historically unstable ones. Their partial coherence and epistemic slippages require students not only to interpret meaning, but also to articulate why apparent stylistic plausibility fails to fully align with historical conventions, contexts, and systems of representation.

Crucially, students are not asked to evaluate the quality of generative AI as a technology, nor to determine whether an image is “good” or “bad” AI output. The Making Connections Essay is not an essay about generative AI. It is an essay about seventeenth-century Western European art and the application of art-historical methods. To demonstrate their learning, students must identify how and why an image fails to align with the visual, cultural, and historical frameworks of the period, and make critical and interpretive connections grounded in the course materials. This requires synthesis between form and content, between image and text, and between individual artworks and broader thematic concerns developed over the course. Generative AI functions here not as an object of fascination, but as a foil that sharpens disciplinary judgment.

Questions of bias and representation emerge naturally in this process. When students analyze AI-generated depictions of historical figures or

devotional scenes, particularly those involving racialized subjects such as Black queens or Black saints, they encounter the limits of training datasets and the consequences of historical erasure. Certain figures prove difficult to render without distortion, refusal, or stereotype, while others are represented through visual conventions that collapse difference or reinscribe familiar hierarchies. These moments are not treated as technical glitches to be bypassed, but as evidence of how visual systems, human and algorithmic alike, are shaped by power, omission, and repetition. Bias is encountered not as an abstract ethical concept, but as a concrete obstacle to historical interpretation.

FROM POLICING TO DESIGN

Much of the early discourse surrounding generative AI in higher education has emphasized detection, prohibition, and policy responses, framing the problem primarily in terms of regulation rather than pedagogy. Analyses of institutional guidance show a rapid proliferation of policies that foreground risk management, academic misconduct, and technological control, often in the absence of sustained attention to assignment design or learning goals.[1] While institutional policies are necessary, an overreliance on policing approaches is pedagogically regressive. Detection tools are widely acknowledged as unreliable, enforcement disproportionately affects already marginalized students, and attempts at prohibition rely on a belief in technological containment that is difficult to sustain.[2] More importantly, these approaches sidestep a more difficult question: why are so many of our assignments so easily vulnerable to automation?

As John Warner argues, the problem is not that generative AI can write, but that many writing assignments have been designed in ways that do not meaningfully require thinking.[3] When students are asked to produce generic summaries or formulaic essays, generative AI exposes the hollowness of those tasks. The solution is not to intensify surveillance, but to redesign assignments so that writing and thinking are inseparable and disciplinary judgment is required.

The Making Connections Essay is one such redesign. Its specificity makes outsourcing more difficult—not impossible, but less attractive—because the task is grounded in shared visual material and course-specific content. More importantly, the assignment shifts the locus of integrity from authorship alone to disciplinary judgment. What matters is not simply who typed a sentence, but whether a student can develop historically grounded interpretations by integrating evidence drawn from course artworks, assigned readings, class discussions and activities, and the instructor-generated AI image within the deliberately curated evidentiary framework of the course. This is not an argument against writing in art history. It is an argument against treating standardized academic prose as the sole evidence of disciplinary understanding.

This shift aligns with guidance emerging from the broader historical profession. The American Historical Association emphasizes that generative AI produces texts and images, not historical knowledge, and that core historical practices such as synthesis, contextualization, and critical evaluation remain essential in an AI-rich environment.[4] In this framing, generative AI does not replace expertise; it makes disciplinary judgment more visible and more necessary.

LEARNER VARIABILITY AND EVIDENCE OF LEARNING

Although the Making Connections Essay was initially motivated by concerns that increasing reliance on generative AI might diminish students' opportunities to develop disciplinary knowledge, interpretive judgment, and habits of critical engagement, it quickly became apparent that it also addressed a longer-standing issue in art history pedagogy: uneven access to academic writing conventions. Students who struggled with grammar or prose fluency often demonstrated strong visual reasoning and historical insight during discussions, yet those strengths were not always legible in traditional essays. Conversely, students who

were adept at producing polished prose sometimes relied on surface-level analysis that obscured gaps in historical understanding.

Recent revisions to the Universal Design for Learning (UDL) framework make clear that such variability is not incidental; it is predictable. The UDL Guidelines 3.0 emphasize that barriers to learning reside in design rather than in learners themselves and call on educators to assume diversity in prior knowledge, language, attention, and modes of engagement as a baseline condition.[5] Importantly, UDL 3.0 moves beyond accessibility understood solely in terms of accommodation to address questions of bias, belonging, and learner agency as integral to learning design.[6]

Viewed through this lens, generative AI does not simply introduce new inequities into art history classrooms; it intensifies and renders more visible inequities already embedded in assessment practices. Writing-only assessments, rigid genre expectations, and high-stakes surveillance disproportionately disadvantage students who are multilingual, neurodivergent, or unfamiliar with disciplinary conventions. The Making Connections Essay responds by requiring students to apply disciplinary skills across course artworks, assigned readings, class discussions and activities in order to interpret and evaluate a shared AI-generated image whose historical instability requires disciplinary judgment. Students still write, but writing is no longer treated as the sole or primary measure of disciplinary competence. What is assessed is the ability to connect ideas meaningfully within a historically grounded framework.

Crucially, this is not a matter of lowering standards. Evaluating synthesis and interpretation demands a high level of disciplinary engagement and historical reasoning. What changes is not rigour, but access to demonstrating rigour.

KNOWING STUDENTS, KNOWING WORK: INTEGRITY BEYOND DETECTION

If learner variability requires us to rethink what counts as evidence of learning, it also requires us to rethink how integrity is recognized and enacted in everyday teaching, particularly in AI-saturated environments. In discussions of generative AI, academic integrity is often framed as a technical problem to be solved through tools, policies, and detection mechanisms. In classroom practice, however, integrity cannot be reduced to technical measures alone. It is also a matter of relationship, context, and judgment, particularly in disciplines such as art history where interpretation, justification, and situated reasoning are central to knowledge-making.

Throughout the course, I build relationships with students through one-on-one meetings, start-of-term surveys, and in-class exercises that allow me to become familiar with their oral and written voices. These practices do not eliminate the possibility of inappropriate or misaligned use of generative AI, but they fundamentally change how concerns are addressed. When a piece of work raises questions, the response is conversation rather than accusation. Students are asked to talk through their thinking, to explain connections, and to articulate their process. In a course structured around close looking, interpretation, and synthesis, this approach is not an add-on or a special accommodation. It extends the discipline's ordinary methods. Art history has long relied on reasons that can be given, claims that can be defended, and interpretations that can be situated in relation to evidence. A student who can speak to those moves demonstrates more than compliance. They demonstrate historical reasoning.

This relational approach also clarifies what detection cannot do. Detection tools promise certainty, but in practice they often substitute disciplinary judgment with suspicion by focusing narrowly on linguistic patterns or surface features of text as proxies for authorship. When integrity is made to depend too heavily on automated inference about textual authorship, the result is not clarity but distortion. Practices we claim to value, including careful justification, interpretive

transparency, and accountability to evidence, risk being overshadowed by an overemphasis on algorithmic provenance claims about authorship in student writing rather than on the quality of historical reasoning itself. Integrity becomes a policing problem rather than a learning problem.

This pedagogical stance aligns closely with Sarah Elaine Eaton's concept of postplagiarism, which she defines as an era in which advanced technologies, including generative artificial intelligence, are a normalized part of how people learn, teach, and produce knowledge.[7] Eaton argues that postplagiarism requires a shift away from punitive, compliance-driven models of academic integrity toward ethical, educative, and relational frameworks that acknowledge technological change without abandoning responsibility.[8] In this view, integrity is not primarily enforced through surveillance. It is cultivated through design, transparency, and trust, and it is sustained through relationships that make learning legible. Knowing students as thinkers makes it easier to recognize when work is misaligned and to respond in ways that support learning rather than criminalize uncertainty. The goal is not to eliminate risk through control, but to build conditions in which students have both the support and the responsibility to do the work that the discipline actually demands.

WHAT GENERATIVE AI MAKES VISIBLE

If generative AI has become a flashpoint in higher education, it is partly because it makes visible what many instructors have sensed for years but could more easily set aside. It reveals the extent to which some assessment practices rely on proxies. Prose stands in for thinking. Polish stands in for understanding. Compliance stands in for learning. When generative AI can produce fluent text on demand, these substitutions become harder to defend, and the resulting anxiety often gets routed into detection, prohibition, and policy. Yet what is being exposed is not simply student behaviour. It is the fragility of the measures we have used to decide what counts as knowledge.

In art history, this visibility has particular force because our discipline already knows that plausibility is not the same as meaning. A surface resemblance to a style does not guarantee historical specificity. A convincing image does not necessarily sustain an argument. The epistemic slippages that appear in AI-generated images concentrate, in visual form, a problem that has long existed in written assessment. A student can produce stylistically competent prose without making historically grounded claims, just as an algorithm can produce stylistically competent imagery without historical understanding. This is why the distinction between style and meaning, between appearance and interpretation, becomes newly urgent. The technology intensifies the need for judgment, and it clarifies that judgment, developed through disciplinary training and practice, is not an optional enrichment but the core of art historical practice.

Generative AI also makes longstanding inequities more visible, particularly when institutions respond through surveillance-oriented approaches. When assessment depends primarily on standardized academic prose and when integrity is framed as a detection problem, students who are multilingual, neurodivergent, or unfamiliar with disciplinary conventions bear disproportionate risk. The question is not whether art history should abandon writing or abandon research but whether we can design assessment practices that make disciplinary thinking visible for a wider range of learners without reducing rigour to a single communicative form.

The Making Connections Essay suggests one path through this problem. It recentres analysis, interpretation, and debate by positioning instructor-generated AI images as shared objects of inquiry through which students demonstrate disciplinary judgment. It asks students to connect course materials in a historically accountable way. It brings bias and representation into view not as abstract ethical concerns but as interpretive problems with consequences. Most

importantly, it insists that art history's strengths are not located in an ability to produce text alone, but in an ability to make and justify claims about visual and historical meaning. In an environment saturated with automated outputs, that insistence becomes both more difficult and more necessary.

CONCLUSION

The most pressing question posed by generative AI is not whether students will use it, nor whether institutions can contain it. Rather, it is whether we are willing to reconsider the assessment habits that the technology has exposed as brittle. In art history, where disciplinary knowledge depends on interpretation, contextualization, and judgment, the response cannot be a retreat into surveillance. It must be a renewed commitment to designing assessments that require the work we actually value.

The Making Connections Essay is not a universal solution, and it does not claim to replace research or writing in the discipline. It offers a specific intervention that treats generative AI as an object of critique and a catalyst for clearer thinking about evidence of learning. By shifting attention from policing to design, from authorship to judgment, and from uniformity to learner variability, we can respond to generative AI in a way that is intellectually rigorous, ethically grounded, and pedagogically honest. The task is not to make students write without machines. It is to ensure that what we ask students to do remains accountable to historical meaning and to the human work of making connections.

ENDNOTES

1 Nora McDonald, Aditya Johri, Areej Ali, and Aayushi Hingle Collier, "Generative Artificial Intelligence in Higher Education: Evidence From an Analysis of Institutional Policies and Guidelines," *Computers in Human Behavior: Artificial Humans* 3 (March 2025): Article 100121, https://doi.org/10.1016/j.chbah.2025.100121; Jennie Miron and Laura Facciolo, "Policy Analysis of Higher Education Institutes in Ontario, Canada: A Focus on Artificial Intelligence," *Canadian Perspectives on Academic Integrity* 8, no. 5 (2025): 1–12, https://doi.org/10.55016/ojs/cpai.v8i5.81779; Eli Rasky, *Generative AI Policy in Higher Education: A Preliminary Survey* (Centre for International Governance Innovation, 2024), https://www.cigionline.org/publications/generative-ai-policy-in-higher-education-a-preliminary-survey/.

2 Doraid Dalalah and Osama M.A. Dalalah, "The False Positives and False Negatives of Generative AI Detection Tools in Education and Academic Research: The Case of ChatGPT," *The International Journal of Management Education* 21, no. 2 (2023): Article 100822, https://doi.org/10.1016/j.ijme.2023.100822; Loleen Berdahl, "Focusing on GenAI Detection is a No-Win Approach for Instructors," *University Affairs*, December 11, 2024, https://universityaffairs.ca/career-advice/focusing-on-genai-detection-is-a-no-win-approach-for-instructors/; Maddy Dwyer and Elizabeth Laird, *Up in the Air: Educators Juggling the Potential of Generative AI with Detection, Discipline, and Distrust* (Center for Democracy & Technology, 2024), https://cdt.org/insights/report-up-in-the-air-educators-juggling-the-potential-of-generative-ai-with-detection-discipline-and-distrust/.

3 John Warner, *More Than Words: How to Think About Writing in the Age of AI* (New York: Basic Books, 2025), 5–6.

4 American Historical Association, "Guiding Principles for Artificial Intelligence in History Education," last modified August 5, 2025, https://www.historians.org/resource/guiding-principles-for-artificial-intelligence-in-history-education/.

5 CAST, *Universal Design for Learning Guidelines*, version 3.0, last modified on July 30, 2024, https://udlguidelines.cast.org/.

6 Anne Meyer, David H. Rose, and David Gordon, *Universal Design for Learning: Principles, Framework, and Practice* (Lynnfield, MA: CAST Professional Publishing, 2026), 61–74.

7 Sarah Elaine Eaton, "Postplagiarism: Transdisciplinary Ethics and Integrity in the Age of Artificial Intelligence and Neurotechnology," *International Journal for Educational Integrity* 19 (2023), Article 23, https://doi.org/10.1007/s40979-023-00144-1.

8 Eaton, "Postplagiarism."