

Teaching AI Literacy Through Cultural Critique


Krystian K. Ramlogan

Pages 92–96 of the Georgia Tech Academic AI Playbook, 2026 Edition


About the Playbook

The Georgia Tech Academic AI Playbook collects stories of how artificial intelligence is being used, taught, and responded to in Georgia Tech classrooms. The 2026 Edition gathers 67 articles from faculty, staff, and researchers across every college at Georgia Tech, along with the Library, the Center for Teaching and Learning, the Georgia Tech Research Institute, the Office of Information Technology, and the Provost’s Office.

Each article is a “play”: a specific initiative carried out in a real course, documenting the prior context, the actions the instructor took, the results, and the lessons learned. The volume closes with 20 lessons synthesized across the full collection.



Front cover



Back cover

Copyright and License

Copyright for the content of this volume is retained by each section’s author or authors. Copyright for sections not attributed to a specific author is retained by the editors.

ISBN 979-8-9978407-0-9. Edited by David A. Joyner and Nimisha Roy. <https://hdl.handle.net/1853/81982>

This work is licensed under a Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International License (CC BY-NC-SA 4.0). <https://creativecommons.org/licenses/by-nc-sa/4.0/>

Suggested Citation

Krystian K. Ramlogan (2026). Teaching AI Literacy Through Cultural Critique. In D. A. Joyner & N. Roy (Eds.), *Georgia Tech Academic AI Playbook* (2026 ed., pp. 92–96). Georgia Institute of Technology. <https://hdl.handle.net/1853/81982>

Title: Teaching AI Literacy Through Cultural Critique

Authors:

- Krystian K. Ramlogan; School of Literature, Media, and Communication;
kramlogan3@gatech.edu

Categorization: Teaching AI

Sub-Categorization: AI Ethics, Policy & Critical Perspectives

Formation

Introducing AI tools into a writing course is not a neutral act, because the decision carries a real risk: choosing generic efficiency over authenticity. Students often meet ChatGPT, Claude, Gemini, and similar platforms as neutral assistants, but these systems reproduce a recognizable set of defaults: linear structure, compressed dialogue, individual protagonists, resolution-driven plots, and a version of clarity that mistakes rhetorical density for excess. For students working in culturally specific traditions, whether Caribbean oral storytelling, non-Western dramatic structures, collective and cyclical story logics, or linguistic registers that carry cultural memory, those defaults are not merely stylistic suggestions. They are structural pressure against the very qualities that make a student's voice worth cultivating.

This concern shaped my approach to screenwriting pedagogy before my current student research team was assembled. The theoretical groundwork, drawing on Haile Gerima's *Triangular Cinema*, Tony Hall's *Jouvay* Popular Theatre Process, and decolonial scholarship, was already shaping my course design when a practical teaching question surfaced: before recommending any AI tool to students for creative work, what do these tools actually do to culturally specific voices once they are given editorial authority? The experiment described here was built to answer that question with student researchers, not for them. The initiative emerged from my work in screenwriting and narrative-writing pedagogy, then extended into broader writing-intensive and media-studies course design.

The stakes reach past screenwriting. Any course with a writing, storytelling, design, or interpretive component faces the same underlying challenge: students may take AI feedback as craft guidance without seeing that the guidance is shaped by what the system has learned to value. The transferable question is plain: what disciplinary or cultural standard should students develop before AI enters the workflow, and what happens when an AI tool is tested against that standard? This initiative starts there, with critical inquiry before tool adoption, not alongside it and not after.

Action

I designed this initiative in three connected parts: a structured AI-comparison experiment, a writing-governance standard, and a bounded course-support deployment. The first was a three-phase critical inquiry experiment run with an interdisciplinary student research team spanning Cinematic Arts, Computer Science, Computer Engineering, Computational Media, and Digital Media. I wanted that mix in the room. Students with technical, systems, and UI/UX expertise sat alongside students grounded in narrative and cultural theory, and that combination shaped how the team read the platforms' behavior: they could weigh both the outputs and the design logic that produced them. The test case was the *Midnight Robber*, a masquerade figure from Trinidad Carnival's *Jouvay* tradition whose elaborate, performative speech combines poetic repetition, archaic diction, historical reference, moral authority, and rhetorical excess. I chose the Robber because his oratory performs survival and ancestral memory at once,

while also carrying the very qualities conventional Western screenwriting often treats as problems: length, density, repetition, and a rhetorical logic that refuses to become efficient at the cost of meaning.

The experiment ran in three phases. In Phase 1, students prompted ChatGPT, Claude, and Gemini to generate a screenplay structure from a *Midnight Robber* premise, explicitly invoking Hollywood conventions. In Phase 2, students supplied original Robber dialogue and asked each platform to make it more “cinematic,” testing how each system defines cinematic value. In Phase 3, students attempted to correct the platforms’ outputs by providing explicit cultural context and requesting restoration of the Robber’s authentic register. Student researchers documented their observations across all three phases; those reflections became primary evidence in the analysis. The core method is adaptable: select a meaningful disciplinary standard, ask AI to act on material governed by that standard, compare what changes across platforms, then ask students to name what the tool valued and what it failed to see.

Alongside the experiment, I developed the Writers’ Room AI Use Standard: a governance framework, analogous to citation practice in research writing, that treats AI as a tool requiring critical evaluation rather than a source of editorial authority. The Standard does not prohibit AI use. It sets the terms under which AI may be consulted without transferring authorship or editorial authority away from the student. Its six operative principles are:

- **Authorial Intentionality Precedes AI Engagement.** A writer must articulate, in writing, the specific purpose of any AI consultation before that consultation occurs. AI tools may not be used to generate initial ideas, first responses to a prompt, or opening structural decisions.
- **AI Output Has No Authorial Status.** Text, structure, or content produced by an AI system is source material, not authorship. It must be disclosed, assessed for fit, and either discarded or substantially transformed by the writer before entering submitted work.
- **Cultural Voice Is Non-Negotiable.** When AI output conflicts with a writer’s cultural, linguistic, or rhetorical identity, the writer’s voice governs. No AI suggestion that compresses or homogenizes culturally specific language may be accepted without explicit written justification.
- **Every AI Interaction Requires Interrogation.** Writers must record what was asked, what was returned, what was kept or rejected, and why. This interrogation record is a required element of the writing process, not an optional supplement.
- **Coaching Functions Are Distinct from Editorial Functions.** AI may pose questions, surface contradictions, or identify underdeveloped arguments. AI may not make decisions. A question the writer must answer is coaching; a revision the writer is expected to accept is editing. Only coaching functions are permissible.
- **Editorial Authority Remains with the Writer.** All final decisions regarding structure, argument, voice, evidence, and form belong to the writer. Submission of work constitutes the writer’s attestation that every substantive decision in the document is their own.

Students working under this model keep an AI Use Log for each major assignment: a running record of what was asked, what was returned, what was kept or discarded, and why. The log is not a policing mechanism; it is where the critical thinking becomes visible and assessable.

The same distinction between support infrastructure and editorial authority shaped how I used AI in LMC 3253: Animation Studies during the Maymester 2026 intensive. The compressed format, fifteen fully online synchronous sessions across three weeks, created a level of student need that the ordinary rhythm of instructor-student response could not easily meet. The question was the one the *Midnight Robber* experiment had already surfaced, now at a different layer: how do you put AI in service of student learning without displacing the intellectual work that is the learning? I built two support tools with

explicit boundaries. Ani, a Custom GPT course assistant drawn from the syllabus and lecture concepts, handled course navigation and concept clarification, but it was not positioned as a tool for writing reflections, essays, or student arguments. I organized Google NotebookLM notebooks by function and used NotebookLM to support my thematic analysis of student project proposals. In both cases, the boundary mattered more than the tool itself. AI could support access, organization, and pattern recognition, but it could not become the place where student interpretation, judgment, or voice was produced.

Result

Across all platforms and all five student datasets, the experiment produced three convergent findings that describe one pattern. The research team named it contextual flattening: the systematic reduction of culturally specific forms to generalizable, recognizable structures. The learning shift was that students could move from sensing that an AI output felt wrong to naming the pattern that produced the loss. That reduction often stays invisible unless students have been prepared to see it, which is what makes it a teaching problem and not only a technical one. The immediate research impact involved five student researchers; the teaching deployment extended the same framework into LMC 3253 through bounded course-support tools. Three findings establish it.

The first was default architecture bias. Every platform imposed linear, resolution-driven structure on a premise that was deliberately non-linear and non-resolving. Each system invented conventional character roles, audience surrogates, discrete antagonists, and clear win conditions that the source material neither contained nor needed. The platforms were not interpreting the story. They were replacing it with a version they recognized.

The second was cultural compression. Asked to make the Robber's speech more "cinematic," every platform treated compression as improvement, shortening and smoothing and stripping the dialogue of its rhetorical density. What the systems marked as excess was exactly what carried the cultural and epistemological work of the original. Audrey Chung, the team's graduate lead researcher, documented the pattern across her Phase 2 data: ChatGPT shortened every sentence; Claude added stage directions for physical movement in place of oratorical weight; Gemini deleted cultural markers, broke the speech into simpler sentences, called the original "soapbox oratory," and then expressed pride at having "modernized" it. Her own reading was that each platform's sense of the cinematic stayed bound to a modern American definition, offering speeches short and to the point on the assumption that less is more, until what came back read, in her words, "like a speech out of a contemporary superhero film." None of the platforms flagged anything as lost. Each framed the revision as an improvement.

The third finding mattered most for teaching: intervention failure. When students supplied explicit cultural context and asked the platforms to repair their own outputs, none restored what had been lost. Surface features could be mimicked: archaic vocabulary, extended length, elevated diction. The underlying logic of the form stayed absent. Harsha Vegiraju's results complicated the picture in a productive way. His background in computer science oriented him toward highly specific prompt engineering, and in Phase 3 he achieved more successful correction than the others working the same protocol. This did not disprove the finding. It exposed its deeper implication. Harsha read the bias as plain on its face: the systems did not see past Western ideologies on their own, and even when prompted for detail they kept deviating back toward Western norms, so that he expected he would have to "continuously prompt the AI systems and force them into accepting other cultures." What his data shows is that the correction did not come from the AI's cultural knowledge. It had none. It came from prompting pressure so specific and sustained that it amounted to supplying by hand what the system lacked. That is

not a prompting skill. It is a labor burden distributed asymmetrically: a burden no student working inside Western narrative conventions would ever carry, and one that stays invisible until you are the student working outside them.

The teaching implication is direct. The limitation is not in how students prompt these systems. It is architectural, built into what the systems were trained to value. Saying this plainly, and showing it through structured comparison, is the most honest and useful thing an instructor can offer. It sets the right scope, it keeps students from over-trusting AI feedback in high-stakes creative work, and it returns authority over the work to the student. What a student does with that authority, once they can see what the tool cannot, is where the real writing begins.

The LMC 3253 deployment showed the same principle in course infrastructure. The tools that worked best were the ones with bounded purposes and explicit exclusions. Ani widened access to course guidance without standing in for student reflection. NotebookLM helped me see patterns across student project proposals without grading or revising any of them. The model that emerges has two parts: AI is first studied as an object with values built into it, and only then deployed as infrastructure with defined limits. For instructors outside narrative fields, the transferable part is not the Midnight Robber itself. It is the structure of inquiry: set a disciplinary standard, test AI against it, compare what changes across platforms and phases, and ask students to account for what the tool changed and why. In that sense, the example remains culturally specific, but the method can travel wherever students are being asked to use AI in work that depends on judgment, voice, interpretation, or meaning.

Lessons

- **Lesson 1: Interrogate the tool before you assign it.** When students meet AI for the first time inside a course, they often meet it as a neutral assistant. It is not neutral. That assistant carries training priorities, design assumptions, and defaults about what good work looks like, and most students have no reason yet to suspect any of it. So the most useful thing I have found is to build an encounter with those assumptions before the tool enters the workflow, and to treat that encounter not as a warning but as the first analytical task of the course. The three-phase experiment here is one way to do it. Its core is simple: put meaningful material in front of the tool, watch what it changes, attempt a correction, then study the pattern. Comparison matters here. One prompt to one platform gives you one output; it does not show you a pattern. Testing the same material across platforms and phases is what lets systemic behavior read as systemic. That move, from intuition to analysis, is the AI literacy outcome I think is worth designing for.
- **Lesson 2: Establish the standard before introducing the tool.** Without a standard already in place, students have no way to recognize what AI has removed. They will read the output as reasonable, because they have nothing to measure the loss against. In this project, that meant grounding the research team in the Midnight Robber's rhetorical and cultural logic before any platform was allowed to touch it. The standard was set in the room before the tool arrived. The same holds in any field, and the questions are yours to pose: what does good evidence look like here, what does culturally specific language carry that compression strips away, what does field-specific argument do that a generalist default cannot do? Working through those questions is not preliminary to the course. It is the course. Standard first, tool second.
- **Lesson 3: Treat AI governance as pedagogy, not paperwork.** The coaching/editing distinction has to be named early and held structurally. A tool that helps a student find a course policy, clarify an assignment, or troubleshoot a technical pipeline is doing support work: it extends access without displacing thought. A system that drafts an argument, rewrites sentences, or decides whether an interpretation is convincing is doing something else entirely. The AI Use Log,

the Writers' Room AI Use Standard, and the bounded design of Ani all exist to keep that distinction visible. Students are not just disclosing AI use. They are articulating an intention before they consult the tool, weighing what comes back, rejecting what does not serve the work, and keeping authority over meaning. That sequence is not compliance documentation. It is close reading turned on the student's own process. A student who can explain why they turned down an AI suggestion on their own writing has not satisfied a governance protocol. They have built the critical literacy this moment asks for, and they will carry it well past the course.

Acknowledgements

The three-phase Midnight Robber experiment described here was conducted with a student research team whose careful observation and systematic documentation made the analysis possible. I am grateful to Audrey Chung, Grace Layman, Harsha Vegiraju, Shoroqu Othman, and Camtien Nguyen for their intellectual engagement and their commitment to this work.