

Guiding Principles for Artificial Intelligence in History Education

Approved by AHA Council, July 29, 2025

In 2023, the American Historical Association Council charged an Ad Hoc Committee on Artificial Intelligence in History Education with exploring the implications of generative artificial intelligence (AI) for history teaching and learning (see the glossary in Appendix 1). A separate committee is focused on research and publication, which falls outside our purview.

This committee recognizes that generative AI tools offer significant opportunities to improve teaching and student learning. At the same time, we respect the concerns expressed by history educators, many of whom feel overwhelmed, distracted, or frustrated by these technologies.

While generative AI is undeniably powerful, it cannot replace human teachers. The most extreme proposals to automate education betray a fundamental misunderstanding of teaching and learning, the core competencies we aim to cultivate in students, and the deeply human-centered work of education. Indeed, the rapid adoption of AI tools suggests that it has never been more important to appreciate the complexity of our shared past and what it means to be human.

History educators have been seeking guidance on how to responsibly and effectively incorporate generative AI into their teaching practice. Some have voiced concerns about the challenges of maintaining academic integrity; others have raised important ethical, environmental, and economic objections to these technologies and their application. We minimize none of these concerns. Given the speed at which technologies are changing, and the many local considerations to be taken into account, the AHA will not attempt to provide comprehensive or concrete directives for all instances of AI use in the classroom. Instead, we offer a set of guiding principles that have emerged from ongoing conversations within the committee, and input from AHA members via a survey and conference sessions.

I. Historical Thinking Matters

Historical thinking remains essential in an age of AI.

The rapid emergence and continuing evolution of generative AI is transforming our relationship with technology. We approach this moment with confidence in the value of human-authored interpretations of history and humility about our predictive capacities. Peter N. Stearns's classic essays—[“Why Study History?”](#) and [“Why Study History? Revisited”](#)—are made no less relevant by the advent of large language models (LLMs) and generative AI.

Many disciplines and professions are changing; the historical discipline will too.

While we cannot predict the future, generative AI is already reshaping many disciplines and professions. As has been the case with other technologies, historians will find distinctive ways to work with generative AI. The need for history and history education will not disappear.

Generative AI can mimic some of the work done by historians and history educators. This should not be mistaken for teaching or for learning. Far from rendering the discipline obsolete, generative AI may increase the demand for historians' specific skills as societies and workplaces navigate an increasingly complex information landscape. The ability to act as subject matter experts, undertake extensive research, synthesize complex secondary literature, and look for biases, inaccuracies, and limitations are invaluable in an age of generative AI. Critical, too, is our disciplinary commitment to accuracy, complexity, and nuance, which remains at the center of historical training.

II. Generative AI and Its Limitations

AI produces texts, images, audio, and video, not truths.

Generative AI is a remarkable technological achievement, but it has undeniable limitations. An awareness of these limitations is important for instructors and students alike. LLMs produce text using an algorithm to select each word from existing books, articles, images, and other media, including AI-created sources. AI texts do not reflect truth; rather, they echo and synthesize, sometimes poorly, sources on which the model has been trained. Generative AI reproduces the limitations of its own training material. By contrast, historians learn to identify and dissect author biases, experiences, social environment, and hidden motivations. Students need to learn to interpret AI-generated content with a critical lens, using their historical training to assess material rather than passively accept it as true or complete.

For all its capacities, generative AI regularly hallucinates content, references, sources, and quotations.

AI models are trained to identify and reproduce patterns, not to comprehend the world in all its complexity and contradictions. If a pattern leads to a false, biased, or imagined output, AI has no way to self-correct. Commercially available generative AI algorithms prioritize speed over accuracy. Given a large task, an AI tool will eagerly invent fictional answers that complete its prompt more quickly, a process often referred to as hallucination. It is essential for students to understand that generative AI can hallucinate data and that historians work to counter these hallucinations when they appear. AI introduces new possibilities for fabricated sources; students must be trained to critically assess all outputs and to recognize that any information provided by a generative AI tool could be false unless properly verified. Evaluating the reliability of sources and assessing the validity of claims are core components of historical thinking and remain especially relevant today.

AI introduces a false sense of certainty where uncertainty exists.

Historians understand that there are things we know about the past and much that eludes us. Generative AI tools risk promoting an illusion that the past is fully knowable. Multimodal models, capable of processing input in one medium to generate content in another, can fabricate strikingly clear visual representations of historical moments that never existed, while chatbots simulate conversations with historical figures as if they were speaking with us directly. These outputs do not represent authentic reconstructions of the past—they are fabrications based on statistical patterns in existing, often flawed datasets. A good history class teaches students to work within the gaps and silences of the historical record, stressing that uncertainty is not a failure but a fundamental feature of historical inquiry. Helping students recognize this fact is essential in an age of AI-generated content.

III. AI Literacy

Banning generative AI is not a long-term solution; cultivating AI literacy is.

Students of all kinds already rely on generative AI tools and will continue to do so. Some committed educators have chosen to reject generative AI for its ethical, environmental, and economic consequences, but ignoring this technology will neither halt its spread nor shield our discipline and students from its reach. We have a responsibility to help students understand these issues in historical context and make informed decisions about their future application. Even if history instructors emphasize in-class writing assignments and exams, the influence of generative AI will be felt both in and outside the classroom. Students will want to use generative AI's formidable tools and will need to understand its limitations. This committee believes that blanket bans are neither practical nor enforceable. Even those who choose to advance student learning in an AI-free environment will have to engage with these technologies. We must determine how to do so responsibly and effectively. Our task is to help students build the critical skills to navigate these tools. Students are already seeking guidance. One of the most meaningful contributions we can make is to support the development of intentional and conscientious AI literacy.

Generative AI can be a valuable partner in the classroom.

Generative AI can be a valuable collaborator for users who know what to ask and how to correct errors. It can enhance teaching and provide a resource for classrooms. It can speed up preparation and suggest alternative or enhanced learning assessments. Generative AI allows seemingly limitless possibilities for assignments that cultivate crucial literacies. For example, a student could be asked to compare an AI-generated summary of an academic article with the original text, assessing what the AI engine gets right, what it gets wrong, and whether the article's most important contributions have been recognized. Such tasks help students cultivate analytical skills while fostering a more nuanced understanding of the strengths and weaknesses of generative AI. It can also prompt students to engage with the original article more deeply, building skills of historical thinking while fostering AI literacy.

Creativity is even more essential in an age of generative AI.

Some forms of assessment, even those hallowed by time, may disappear. Assignments such as short summaries can be easily duplicated by generative AI. On the other hand, creative assignments such as the unessay or in-class role-playing exercises along the lines of the *Reacting to the Past* series will likely become even more valuable.

Training future history educators requires clear and transparent engagement with generative AI.

The teachers of tomorrow, whether K–12 instructors or higher education faculty, are students today. Generative AI will likely be one of the most significant professional issues they encounter. Current history educators have a responsibility to model appropriate engagement with generative AI and to equip future teachers with the ethical frameworks and practical skills needed for their careers. It is essential to prepare future teachers not by abandoning traditional historical training but by combining it with new AI literacies. The core skills of the historical discipline—extensive research, careful source evaluation, critical reading—remain foundational. If anything, the rise of generative AI makes these skills even more essential. At the same time, future educators must be AI literate, which means learning how generative AI systems are trained, how to recognize bias and hallucinations in generative AI outputs, how to use AI tools to support (rather than replace) critical work, and how to teach our

students to do the same. Training the next generation of history educators requires that we hold fast to the disciplinary core of history while expanding the professional toolkit available to future teachers.

IV. Concrete and Transparent Policies

History educators must develop concrete and transparent policies for AI usage and communicate these to students.

Students are navigating a vast and rapidly changing technological landscape with few settled rules. It is essential for history educators to provide clear, consistent guidance for students at all levels and to talk openly about these technologies and their limitations. All syllabi should include explicit generative AI policies that specify when these tools may be used and when they are prohibited (for examples, see the additional resources below). Syllabi should also affirm core scholarly principles—most importantly, the obligation to cite all sources, including AI-generated material. The rise of generative AI does not alter the expectations that underpin historical scholarship. The specific citation format is less important than the act of acknowledging the use of generative AI. Students should be taught how to do so routinely and accurately. Setting expectations openly in a supportive atmosphere—perhaps including an AI-use section in each assignment—encourages students to develop responsible habits without fear of penalty for honest disclosure. Vague or inconsistent expectations risk serious consequences, including unintentional academic misconduct or professional harm. As a possible starting point, we include a model table in Appendix 2.

Experiment, reflect, revise.

No single generative AI policy will be perfect. A landscape in which technologies are evolving rapidly calls for a flexible, experimental, and iterative approach: try new tools and policies, observe their effectiveness, gather student feedback, and revise as necessary. What works this semester may require significant adjustment next year. Teaching AI literacies will require engaging with these technologies and modeling ethical AI engagement. We cannot ask our students to cite content produced or adjusted by generative AI if we do not adhere to this rule ourselves.

V. The Value of Historical Expertise

Generative AI cannot replace historical methodology.

Historical inquiry involves gathering information, making connections, and interpreting evidence in ways that reflect both an individual mind and established disciplinary standards. As a deeply human endeavor, writing history is both science and art. Great works of history are transformative because they are neither predictable nor obvious; therefore, they cannot be replaced by a technology that simply reproduces existing patterns. Generative AI systems are powerful pattern-recognition tools that are also fundamentally limited. They do not think historically; they predict based on past data rather than questioning or reinterpreting it. AI cannot surprise us with new historical arguments, creative reframings, unpublished materials, or original narratives that challenge established understandings. The vast wealth of human history contained in gated archives and nondigitized material is inaccessible to AI engines. At the same time, these tools cannot substitute for rigorous historical methods: finding new sources, posing generative questions, weighing evidence, assessing context, grappling with uncertainty, and constructing original arguments.

There are no shortcuts to expertise.

Evaluating AI-generated content requires expertise that can be built only through sustained engagement with the subject matter. LLMs present a crucial paradox: they can produce material that appears polished and credible, but assessing their outputs demands critical skills that the models themselves can neither teach nor foster. If students rely on generative AI without developing their own skills, they risk entering an unproductive loop: minimal engagement leads to an inability to properly assess outputs, which leads to an uncritical acceptance of flawed material. Our goal is to foster a different trajectory, whereby generative AI is seen as a tool that supports the pursuit of knowledge, not a shortcut that replaces meaningful work. Through active engagement and skill-building, students can use AI thoughtfully by integrating outputs that genuinely improve their work and rejecting those that do not. For example, a student who drafts an essay and uses an LLM to refine its language or sharpen phrasing will need strong writing and analytical skills to evaluate whether AI has served its purpose. In short, expertise must precede AI reliance.

History education must continue to cultivate habits of mind that current and future students will rely on to thrive in a world shaped by generative AI.

Generative AI can produce college-level essays, complete complex research tasks, mimic historical figures, and create realistic-looking historical images and media. New capabilities are emerging every day. We cannot predict the long-term trajectory of these technologies, but we can recognize that we are in the middle of a period of profound change and take steps to meet the moment and prepare our students for the future.

Historical thinking will continue to matter. The study of history prepares students to account for change over time, to recognize the complexity of human existence, and to wrestle with the contingencies that define life in an uncertain world. Everything has a history, and to think historically is to contemplate what it means to be human.

This document was researched and written by the AHA Ad Hoc Committee on AI in History Education in 2024–25:

Jennifer Baniewicz, Amos Alonzo Stagg High School and AHA Council, Teaching Division

Doron Ben Atar, Fordham University, Lincoln Center

Theresa Case, University of Houston–Downtown

Adam Clulow, University of Texas at Austin

Sonia Hernández, Texas A&M University

Kelly Schrum, George Mason University

Tamara Shreiner, Grand Valley State University

Lee Vinsel, Virginia Tech

Brendan Gillis, AHA director of teaching and learning and committee liaison

Appendix 1: Glossary

Artificial intelligence (AI)

A field of computer science that develops systems capable of performing tasks that typically require human intelligence, such as reasoning, learning, perception, language understanding, and decision-making.

AI chatbot

A program or interface that applies AI to simulate conversation and provide information. The most sophisticated AI chatbots use large language models to generate text in real time.

Generative AI

A type of artificial intelligence that can create new content, such as text, images, audio, or video, by learning patterns from existing data.

Hallucination (in AI)

When generative AI produces content that sounds plausible but is factually incorrect or entirely made up.

Large language model (LLM)

A type of generative AI specifically trained on large-scale textual data to understand and generate human-like language through statistical pattern recognition.

Multimodal generative AI

A type of generative AI capable of operating across multiple media and modalities, such as text, image, video, and audio. A multimodal model might, for instance, process textual input and an audio recording to produce a video.

Prompt

The input or question a user gives to a generative AI model, which the model then uses as a basis to produce a response.

Token

A piece of text, often a word or sub-word unit, that language models employ to process and generate language.

Training data

The data used to “teach” an AI model how to recognize patterns or perform tasks.

Appendix 2: Example of an AI Policy Table for Use in History Education

Students and colleagues have many questions about if and when it is appropriate to use generative AI in research, analysis, and writing within the classroom context. Providing clear expectations is essential. As a preliminary model, the AHA Ad Hoc Committee on AI in History Education has composed the following table for individual faculty, departments, and institutions to consider as they develop and refine their own policies. In characterizing many of the uses of AI for particular tasks as acceptable or not acceptable, our goal is to advocate for transparency and offer a format for communicating expectations. We urge educators to confirm that their course and departmental policies align with institutional requirements. We also expect that departments may decide to adopt different policies for courses at different levels (primary, secondary, general education, advanced undergraduate, and graduate). We appreciate that our colleagues will have a wide range of responses to this table, and we hope that this document will serve as a starting point for discussions.

Task	Could this be acceptable use?	Under what conditions?
Ask generative AI to identify or summarize key points in an article before you read it	Yes	Acceptable without explicit citation
Use an AI chatbot as a writing partner to help generate and develop ideas	Yes	Acceptable, may require explicit citation depending on circumstances
Ask generative AI to produce a starter bibliography	Yes	Acceptable without explicit citation only if each reference is checked and additional databases and sources are mined
Ask generative AI to produce a historical image for a paper or presentation	Yes	Image should be clearly marked as AI generated and with explicit discussion as to how the image was created. Images should not be shared beyond the classroom
Ask generative AI to fix the structure or formatting of your footnotes	Yes	Acceptable without explicit citation
Ask generative AI to write an essay or chapter. Submit that essay or chapter as your own work	No	Never acceptable
Write an essay/chapter. Ask AI to sharpen the language but not modify, add to, or replace the main points	Yes	Acceptable use without explicit citation only if changes suggested by AI are limited to grammar and syntax

Write an essay/chapter. Ask AI to add additional points	Yes	Acceptable with explicit citation only if fact-checked and adapted in your own words
Ask AI to summarize a book or article in your field. Use this as a starting point for critical engagement	Yes	Acceptable without explicit citation
Ask AI to summarize a book or article in your field. Reproduce that summary in your literature review without reading the book or article	No	Never acceptable, as there has been no engagement with the book or source itself
Use an AI generated summary of scholarship to critique another scholar's approach	No	Never acceptable, as AI frequently makes basic errors and there has been no engagement with the scholarship
Include a reference generated by AI in a footnote without checking the original	No	Never acceptable

American Historical Association

AI Policy Template for Use in History Education

Task	Could this be acceptable use?	Under what conditions?
Ask generative AI to identify or summarize key points in an article before you read it		
Use an AI chatbot as a writing partner to help generate and develop ideas		
Ask generative AI to produce a starter bibliography		
Ask generative AI to produce a historical image for a paper or presentation		
Ask generative AI to fix the structure or formatting of your footnotes		
Ask generative AI to write an essay or chapter. Submit that essay or chapter as your own work		
Write an essay/chapter. Ask AI to sharpen the language but not modify, add to, or replace the main points		
Write an essay/chapter. Ask AI to add additional points		
Ask AI to summarize a book or article in your field. Use this as a starting point for critical engagement		
Ask AI to summarize a book or article in your field. Reproduce that summary in your literature review without reading the book or article		
Use an AI generated summary of scholarship to critique another scholar's approach		
Include a reference generated by AI in a footnote without checking the original		