Report of the Ad Hoc Committee on Artificial Intelligence (AI)

Department of History, University of Washington

December 2025

BACKGROUND AND EXECUTIVE SUMMARY

The Department Chair formed this ad hoc committee at the beginning of AY 2025–26, asking us to consider our department's approach to AI and its potential uses by both students and faculty in research, teaching, and learning. He also invited us to make recommendations to the department about policies, practices, and resources.

"Artificial intelligence" is a field that covers an enormous range of advanced computational approaches, platforms, and applications. Our committee focused its inquiry on the set of tools that have become widely commercially available in the last three years and that now most directly impact our work as scholars and teachers of History: large language models (LLMs) and AI-enabled software currently in use in classroom, research, and professional settings.

We began our work at the start of Autumn Quarter with several open-ended guiding questions:

- How should we as researchers and teachers navigate our relationship to this new technology? What should be the guiding ethics behind AI use in our department?
- What guidance, if any, should we as a department give around best practices or policies?
- How should we specifically approach the potential use of AI in undergraduate studying and learning?
- How should we approach the potential use of AI in graduate student training?
- What does "AI literacy" mean and what is its value to a history education?
- How should faculty and graduate students communicate about the potential use of AI in their coursework, exam preparation, and dissertation work?
- What resources and support would be useful for the department to provide to the History community?

Our committee recognized at the outset that we were unlikely to devise a singular answer to any of these questions, nor should we attempt to do so. We believe that static, "one-size-fits-all' policies will not fit well in a departmental culture that long has allowed a range of pedagogical and research approaches, nor would they be particularly effective given the rapidly changing landscape of AI tools and capacities. We also are keenly aware of the shortcomings of AI-detection software and how difficult it can be to properly and fairly police AI use. Nonetheless, we see a clear need for more policy consistency, including more standardized syllabus statements and clearly articulated departmental principles around AI use by students and by instructors.

We see this "AI challenge" as an opportunity for us not merely to more clearly delineate shared principles of fair use and academic integrity, but for us as historians to teach critical AI literacy and, in so doing, to articulate and defend historical methods, research, writing, and thinking.

To that end, this report recommends that the History Department:

1. Establish a new, AI-informed "Gateway to History" course for all undergraduate History students, taken shortly after declaring the Major or Minor;
2. Provide instructional support and create curricular materials for historically-informed critical AI literacy;
3. Create more standardized AI syllabus statements;
4. Expand departmental limitations on the use of AI in designing and teaching courses;
5. Create a public-facing statement of departmental principles and values, as well as a curated and regularly updated online resource toolkit; and
6. Support faculty in understanding and using AI.

Further details about and rationale for these six recommendations continues after the discussion below. An appendix to this report provides more detailed analysis of the use cases for and limitations of various AI tools.

DISCUSSION

We begin this report with a discussion of what we as a group have identified as the opportunities and limits presented by current AI products and platforms to our work as researchers, teachers, and professionals. The discussion is the result of the collective research by the committee, which surveyed the policies and practices of other units on this campus and elsewhere, reports by professional associations (the AHA and others), available AI literacy toolkits, industry data and technical publications, and interviews with colleagues in this department and beyond.

We detail some use-cases for AI in historical research and teaching in the Appendix to this report, and offer potential policies for each. Broadly speaking, we acknowledge both the promised and substantiated value of AI in facilitating aspects of historical research, and we recognize that in many parts of the digital ecosystem, the reach of AI technologies is now nearly unavoidable. "Opting out" is not an option.

Nevertheless, we recommend caution in proceeding down the path of explicit recourse to AI tools. Moreover, to the extent AI is employed in research and teaching, we believe it should only be done with carefully-considered guardrails in place, and under a clear ethical obligation to acknowledge any such use.

AI and research

Like other tools that have been added to the historian's kit in recent years, current generative AI models have the capacity to augment historical work and to automate some tasks. Generative AI is not a single technology, however, and some uses are more prone than others to historically illegitimate implementation.

To emphasize a fundamental point, AI does not do what historians do. At base, large language models and other generative AI products process data and orient them in high-dimensional space,

sometimes reorienting that space with the addition of new data, with the effect of refining the model. LLMs are complicated, costly, and resource-intensive word guessing machines.

Like historical analysis, creating an LLM is an iterative process. Unlike historical analysis, it is only iterative: generative AI programs have no capacity to "think," and they are not capable of distinguishing between opposing frameworks of understanding, or mounting novel arguments that are not included in or derivable from their training data. These models are continuously refined and expanded, but a thousand generations of refinement will not render a LLM into something that it is not.

In short, AI does not think like an historian because it cannot. What we do is synthesize and create new frameworks for understanding; what AI does is synthesize by fitting data into existing frameworks. Although AI systems can aid in humanistic endeavors and streamline some tasks, to think that an LLM could perform historical work is to commit a category error.

One reason that we urge caution is that AI tools have the capacity to add efficiency to humanistic inquiry, and yet efficiency is often at odds with the historical task. Historical research is inefficient, and purposefully so; expertise in archives and the humans who created them requires deep knowledge of sources which comes only from prolonged engagement with primary evidence and secondary analysis thereof. We especially caution against ceding "low-level" tasks like transcription (excluding OCR, addressed below) and translation to LLMs, an operation that poses the real danger of eroding skills required to gain expertise in an archive. Such "low-level" skills are not themselves the constitutive elements of expertise, but rather they are the skills that allow one to gain the historical expertise that is and can only be the remit of human intelligence.

Furthermore, the forefront of AI implementation at the moment comprises, in large part, work in "prompt engineering": devising questions that will elicit the correct answer from a model that is largely a black-box. Asking *historical* questions, on the other hand, requires attention to aspects and entireties of archives to which large language models are not and cannot be attuned. Asking historical questions requires expertise of a deeply human nature that is not reproducible by a word guessing machine.

Although each use-case has its own considerations, there are a few general concerns that have structured our reservations regarding the use of AI:

- The historical craft relies fundamentally on attention to particularity of sources (linguistic and otherwise), to the cultural specificity of the context in which they were created, and to the situated perspective of the historian. We are concerned that a generalized LLM trained on undiscriminating (and largely modern) datasets will often elide particularity, specificity, and perspective, leading to flat historical analysis, and at times misleading conclusions.

- Use of automated systems naturally presupposes access to appropriate digitized sources. On the one hand, AI promises both to facilitate widespread digitization, and to allow large corpuses (once digitized) to be more thoroughly and systematically mined than could be done through traditional methods. On the other hand, historical research based on AI tools will inevitably be warped by the nature of the archive of digitized sources. This issue is not novel for historians; we have grappled with the arbitrary nature of

archive collections for centuries. Nevertheless, the ways in which AI source availability will skew historical analysis are likely to be particularly unpredictable and inequitable.

- Any guardrails for the use of AI will of course also need to protect against the well-known issues of hallucinations, errors, and misinformation in AI outputs.

AI and teaching

Students already use AI for a variety of tasks, from generating discussion points to brainstorming ideas when stuck on essays. Students also use AI to write papers and do other coursework for them. Instructors' ability to contain and identify AI-enabled academic misconduct have been confounded by the rapid advance of AI capabilities, shortcomings of policing software, and the lack of clear university-wide policy guidance on this issue.[1]

Furthermore, while many students find AI tools helpful for sparking creativity and outlining work, reliance on AI can substitute for direct engagement with instructors. (This highlights the need to encourage students to attend office hours and/or schedule meetings outside of office hours for guidance.) The documented proclivity of plagiarism detection systems (like Turnitin and GPTZero) to flag original work as generated by an LLM can cause panic and, sometimes, false accusations of dishonesty.

In the meantime, our instructors have taken a range of approaches in how much or how little AI they are allowing in their classrooms, as well as what kinds of tools and tasks are acceptable. The lack of a consistent baseline has created confusion and some anxiety among undergraduate students who may be receiving different messages on permissible AI use in each course they take (even within History).

Our graduate students can receive similarly mixed messages on the acceptable use of AI in their training, instruction, and research. In many domains, it has been proposed that interactive and iterative exchange between AI and student will produce a whole greater than the sum of the parts. As in all such domains, however, the prospect of skill erosion looms large. With regular reliance on AI, it seems nearly unavoidable that un/underutilized professional skills will atrophy. This is a particular concern with regard to graduate education. To the extent AI tools become widely integrated into dissertation research, the risk arises that future historians will fail to develop certain fundamental professional skills in the first place, thus undermining their capacity to complement and cross-check automated systems.

[1] The UW Student Conduct Code's current guidelines on AI are as follows: "...cheating includes the unauthorized use of assistance, including technology, in completing assignments or exams. While some instructors may encourage you to utilize technology to enhance your learning experience, other instructors may prefer that you do your own work without seeking outside help. It is your responsibility to read the syllabus for each course you take so that you understand the particular expectations of each of your instructors. If you are unsure of expectations, you are encouraged to ask for clarification before you use specific resources in completing assignments." Community Standards and Student Conduct, "Academic Misconduct," https://www.washington.edu/cssc/for-students/academic-misconduct/ (accessed 18 November 2025).

In summation, we acknowledge that AI poses problems and solutions to our teaching mission. On one hand, students can use it as a complementary tool to generate ideas, proofread work, and aid in preliminary research tasks. On the other hand, students can use AI to undermine (and or avoid) the intellectual work expected of students to complete the major or minor in History, as well as skills foundational to doctoral training in our discipline. Equipping students with critical skills for responsible AI use are essential steps toward maintaining academic integrity and fostering deeper learning.

RECOMMENDATIONS

1. ***Create a required "Gateway to History" course for all undergraduate History students, taken shortly after declaring the major or minor.***

AI poses significant challenges with respect to academic integrity, and it also can mislead students who lack discernment skills to distinguish credible sources from unreliable ones. To address this, our proposed solution is a required "Gateway to History" course for all History majors and minors. This idea has long been considered by our department; we believe this "AI moment" is a ripe opportunity to implement it. AI literacy and expectation-setting would be a component of this course, but it would not be the sole emphasis. Instead, it would emphasize the fundamentals of historical research and writing, introduce students to sources, methods, and resources at the UW, and help them understand the uses and limitations of AI tools and platforms. In doing so, it would show students how using one's original and authentic voice in crafting and communicating history is a form of exercising agency.

We believe that not making the course solely AI-focused will make it a more durable and flexible vehicle to adapt over time. It also harmonizes well with longstanding and continuing interest in the department for a gateway/prerequisite course for the major/minor, including providing education in sources and methods earlier in the curriculum than the 388/494 seminars. The course would be offered regularly, ideally quarterly (depending on volume of entering students), and it would be a prerequisite for the major. A potential model would be the first-year interest group classes (FIGs) or introductory classes offered by other units.

The committee discussed several different approaches to the credit-hour load for this course. A one-credit course would likely not be enough to deliver the content we propose here. Another possibility would be to fold this curriculum into our existing five-credit 288 seminars (conceived as introductory sources-and-methods courses), if this did not impinge on existing courses or require a new designation or catalog description. The optimal approach may be to make this a 2-3 credit-hour course–sufficient for some breadth and rigor, but not an onerous load on the existing curriculum. We ultimately leave open the question of credit load, to be decided by the department if this recommendation is implemented.

Ideally, this would be an in-person class. However, in the interest of promoting accessibility and enrollment, it also could be offered as a synchronous online course. Commuting students and those who work would benefit from having this option. We do not recommend offering the

gateway to history as an asynchronous course, as students are required to do many of these, and part of the intent of the course is to emphasize the importance of person-to-person instruction, human-driven research, and interactive pedagogy.

This course could incorporate some of the same content conveyed through in-class research presentations by library staff, but it would not replace them entirely (see recommendation 2 below). Incorporating History research presentations in a regularly offered course would provide a predictable structure and workload for the instructional staff and librarians who deliver these presentations. It would further institutionalize these positions in an age of budget cuts and AI-driven efficiencies, and strengthen the case for these non-instructional staff FTEs being maintained permanently. Another advantage of this stand-alone course would be that it would exist apart from regular class time and allow more time for instructors to teach content rather than address issues of AI literacy and other core concerns of every history major that are better addressed in a coherent, consistent fashion.

2. ***Provide instructional support and create curricular materials for historically-informed Critical AI Literacy***

Historians have a responsibility as teachers to ensure that student learning is preserved in an environment in which AI is ubiquitous. At the same time, we can help our students understand where and how AI can be used critically and constructively to enhance the learning environment.

Critical AI literacy encourages scholars and students to use their own values, needs, and experiences to make mindful choices about when or if to use AI tools. The "Gateway to History" course will teach these skills to our majors and minors, but other curricular resources will be needed to reach non-History students in our courses, especially lower-division ones.

We suggest that the department collaborate with library staff to develop a History-specific critical AI literacy teaching module that can be regularly offered to our undergraduate classes. This would be similar to the in-class library research presentations that the former and current History librarians long have delivered. It would not be a tutorial or how-to, however. Instead, the presentation would help students to think critically and historically about where these tools come from, the social biases and power structures AI can reflect and reinforce, and the ethical implications of AI use. The module could be designed for in-class presentation or, for when there is not time or staff resources for live demonstrations, as an online tutorial.

For those faculty who are not able to devote class time to these presentations, we recommend that the department make available a standard informational slide (or slides) on historically-informed critical AI literacy, developed by the Digital History Librarian, that faculty can incorporate in their own class lecture presentations on the first day of class.

3. ***Create more standardized AI syllabus statements.***

Although we recognize that instructors should have freedom to determine how AI is used, or not used, in their classrooms, our students will benefit not only from better AI education but also clearer standards and predictability. For both our undergraduate and graduate programs, any

cases of potential academic misconduct involving AI will be easier to adjudicate if expectations and policies are set out clearly in advance.

Rather than attempt to develop a one-size-fits-all AI policy, the department could develop different sets of "high-medium-low" standards, i.e. options for standardized language that an instructor can choose from and then convey to their students on their class syllabus. We believe the Undergraduate Studies Committee would be an appropriate body to develop these different sets of standards. This approach would allow instructors to choose their own approaches but also give students more clarity and consistency.

4. Expand departmental limitations on the use of AI in designing and teaching courses.

The committee strongly supports the Department's recently adopted policy to limit faculty from using AI to grade student work unless a waiver is requested and granted. We recommend that the Department also consider extending that policy to include authoring of course content (slides, handouts, discussion questions and activities, quizzes, exams, paper prompts, etc.) and providing feedback to student assignments (papers, exams, etc.).

At a most basic level, students already have the ability to ask an LLM to generate content or provide feedback related to any historical topic. It is our belief that such AI responses are likely to be of limited utility, and clearly far inferior to the value provided by a member of the history faculty. Any time savings that might result from guided use of AI by a faculty member in conducting routine teaching tasks is, in our view, likely to be far outweighed both by the substandard and generic results thus generated, and by the associated degradation of our teaching vocation and the values we seek to instill in students.

If faculty feel that generative AI use is merited in day-to-day tasks of instruction, they could petition and receive a waiver from the department to do so. We recommend that faculty state their AI use (or non-use) explicitly on their syllabus.

5. Create a public-facing statement of departmental principles and values, as well as a curated and regularly updated online resource toolkit.

We recommend that a short, public-facing statement of principles be published on the departmental website that emphasizes the importance of historical thinking, methodologies, and writing in the age of AI. We believe this will be much more effective, enduring, and reflective of our departmental values and culture than specific policies around AI use, policing, or surveillance.

Such principles would communicate the value of historical thinking and original writing as well as the limitations of AI (as discussed above). They would emphasize that the choice to use them should be made with discretion and critical and historical awareness. This includes understanding who is behind the creation of AI systems, the ideologies and biases shaping these systems, the aims of these tools, and their societal impacts.

The principles would be accompanied by a curated and regularly updated online library of resources on AI for students, faculty, and others interested in the practice of history. This living

document, regularly updated, would include links to resources on teaching, research, critical AI literacy, and recommended reading by historians and others. As many fine resources already exist (on this campus and elsewhere), this toolkit would aggregate key policies, best practices, and critical readings on AI.

The toolkit could include both public-facing and instructor-only sections. The latter might also include scholarship on alternative modes for knowledge assessment and assignments, or even a space for instructors to share successful assignment practices. (To ensure appropriate confidentiality, this section could be user-restricted by UWNetID.) The analysis of AI tools in this report's Appendix could be part of this resource as well.

Regular maintenance and updates of this toolkit could fall under the purview of the RA who manages the department's website, newsletter, et al.

6. Support faculty in understanding and using AI.

We recommend that the department give ongoing support to our faculty and instructors by offering one-on-one tutorials on AI tools 'explaining their strengths, limitations, how students may be using them ethically and otherwise, and how such tools can be both helpful and detrimental to research. This task might be performed by a graduate student SA/RA.

Respectfully submitted,

Eric Johnson
Mark Letteney
Margaret O'Mara (chair)
Ben Schmidt
Chris Tounsel (ex officio)
Aubrey Williams
Charity Urbanski (ex officio)
Jake Beckert (graduate student representative)
Willa Hoey (undergraduate student representative)

APPENDIX - USE CASES

In order to foster more concrete thinking around these issues, we decided to assess the problem through potential use-cases, which are detailed below. In each case we have tried to explore in fair and open-minded terms both the potential benefits and pitfalls of employing AI as part of our professional toolkit. In so doing, we have consciously excluded potential uses of AI that might be considered more 'clerical' in nature, or less specific to the field of history, like polling calendars to find an available meeting time, or drafting a blurb to promote an upcoming talk. The primary author of this matrix was committee member and Digital History Specialist Eric Johnson.

Transcription (OCR image to text)

- **Comments:** AI can potentially allow OCR (a well-established technology) to be applied to new document source bases more easily – as it replaces bespoke software with more generalizable systems driven by training data.
- **Potential positives**: OCR can be the first step in efficiently creating digitized documents, suitable for searching, further AI-driven analysis, and online publication. AI can make OCR processing faster and more accurate, while potentially making OCR available for new languages and document types at a lower cost and time investment.
- **Potential negatives:** As with all OCR technology, the potential for errors remains significant. Manual cleaning will likely be needed for many purposes, which will often require expert knowledge along with substantial time investments. Naïve use of the outputs could be problematic. Despite the advantages offered by AI-driven OCR, significant time and technical expertise may still be needed to get a new system off the ground. In general, these systems will require substantial existing corpuses of digitized and cleaned data to serve as training data. Broadly speaking, therefore, it is unclear whether AI-driven OCR will diminish or widen the gap between the tools available to scholars working in commonly-used and modern collections, as opposed to those working in smaller, more marginal, and older collections. Additionally, currently-offered transcription platforms (e.g. Trankribus) have poor support for traditional diplomatics, which can flatten historical artifacts into disembodied texts without the material evidence for their transmission that is the core of historical work in manuscript studies.
- **Possible policies or guardrails:** Avoid unquestioning reliance on raw OCR outputs, particularly when used by non-experts. Teachers promoting use of automated transcription ought to model detailed validation procedures for students, including attention to diplomatic information loss when moving from object to transcription. Scholars ought to include significant paradata with any dataset indicating the nature of the transcription process, the extent of

validation and quantification of error rates, and thorough indication of diplomatic information excluded in the transcription process.

Literature search (identification of published sources)

- **Comments:** AI tools could be used to identify published literature relevant to a topic of interest, either in lieu of, or in addition to, traditional search methods (algorithmic web and library search, collection browsing, bibliographic aids). This use case presupposes that any draft bibliography thus produced would not be an end-product in itself, but raw material for the researcher's further reading and analysis. Note that complete avoidance of AI in this regard is becoming infeasible, as AI is being integrated seamlessly into search engines and other tools.
- **Potential positives:** In principle AI could greatly reduce the time necessary to assemble a viable draft bibliography on any given subject. Such a bibliography could, theoretically, end up being more complete than one assembled by more traditional methods. In addition, AI could potentially include literature in languages that the researcher does not possess.
- **Potential negatives:** It is unclear how accurate and comprehensive AI systems will be in selecting appropriate literature. AI systems trained for general use may lack the tuning needed to curate an appropriate academic reading list targeted to a particular research question. One obvious limiting factor is the source base made available to the AI system. Non-digitized, paywalled or excluded literature would potentially be excluded – and publishers have become increasingly adept at excluding chatbot crawlers from their databases. In short, much important scholarship remains electronically marginalized due to incomplete digitization and poor OCR. Relying on algorithmic search has the distinct capacity to cause researchers to overlook important studies, leading to reduplicative work at best, and at worst, historically indefensible outputs that would be avoided if traditional bibliographic methods were prioritized. Using AI to generate a secondary source base may undermine the serendipity of unexpected connection making that comes with more traditional perusals of bibliographies and collections. Hallucinations remain a distinct possibility, but tempered by the expectation that the researcher will use this list as a springboard for substantive work.
- **Possible policies or guardrails:** When conducting a literary search, it is recommended that a draft list be developed first using traditional means; after which explicit recourse to AI might be beneficial in helping to identify gaps and overlooked sources. This presupposes that the researcher has the skills to do this work without the assistance of AI. One risk that should be acknowledged is that as use of AI as a 'backstop' grows more common, the skills and will to do the work by traditional methods may naturally wane.

Archival search (identification of archival sources)

- **Comments:** AI tools could be used to winnow a large archival collection by identifying primary source documents relevant to a particular research question. This would be comparable to keyword searching an archival database, but potentially more robust.
- **Potential positives:** In principle AI could vastly compress the time needed to find relevant documents in the archives. It could theoretically identify such documents in a more comprehensive way than either keyword searching, or manual browsing that is directed by traditional finding aids. In addition, AI could potentially include literature in languages that the researcher does not possess. When AI search is used in an iterative fashion, it potentially has better flexibility for adapting 'mid-stream' as the research topic evolves, relative to more traditional archive perusal.
- **Potential negatives:** Such a use case assumes that the archive has been comprehensively digitized and OCRed. Although increasingly common, it remains the exception rather than the rule in many fields, and especially so in pre-modern collections. It is unclear how accurate and comprehensive AI systems will be in selecting appropriate documents. AI systems trained for general use may lack the tuning needed to make historical judgements tuned to a particular research question. Even more than in conducting secondary literature reviews, historians have traditionally profited from immersion in the documentary source base. AI-driven search risks alienating historians both from the cultural texture in which particular documents exist, and the institutional frameworks that drove their creation and retention. Likewise, the serendipity that comes from browsing the archives, and the resultant retuning of research questions, may be undermined by AI-driven searches. Hallucinations remain a distinct possibility, but tempered by the expectation that the researcher will use the output as a starting point for substantive work.
- **Possible policies or guardrails:** Historians might make a practice of never relying solely on AI-driven surveys of archival collections, but always pair it with more hands-on methods. In particular, browsing prior to any AI-driven survey is likely to lead to the formulation of better research questions and prompts. As above, teachers modeling use of AI for graduate students ought to model comprehensive validation procedures, and researchers ought to publish paradata indicating the nature of their discovery practice and the specific models used to parse datasets.

Literature review and summarization

- **Comments:** AI tools could be used to complete something like a traditional literature review – summarizing a large body of secondary sources to identify major schools of thought, defining debates, pivotal works or authors, and prevailing arguments.
- **Potential positives:** In principle AI could save time in reading and digesting large amounts of published material. In principle this could subsume a global academic literature base, including languages the researcher does not have. In principle, such a literature review might encompass a corpus more extensive than anything one researcher would have time to read.

- **Potential negatives:** It is unclear how much AI systems trained for general use could be relied on to paint an accurate interpretive picture of any particular historiography. By relying on a generic literature summary, the researcher's own particular analytical perspective on the historiography is likely to be elided. It is unclear what judgement an AI system will use in selecting sources, or how reliably that would conform to academic standards. One obvious limiting factor is the source base made available to the AI system. Non-digitized, paywalled or excluded literature would potentially be excluded. Hallucinations remain a distinct possibility, and the risk is great to the extent such a literature review is taken as authoritative.
- **Possible policies or guardrails**: In general, this use is not recommended. To the extent it was pursued, it might be employed as a backstop to augment the researcher's own takeaways from the literature. This presupposes that the researcher has the skills to do this work without the assistance of AI. One risk that should be acknowledged is that as use of AI as a 'backstop' grows more common, the skills and will to do the work by traditional methods may naturally wane.

Textual and thematic analysis of a primary source corpus

- **Comments:** AI tools could be used to do various sorts of abstract textual analysis across a very large document corpus. Such analysis might, for example, seek to identify prevalent topics and themes; patterns of change over time; suggestive linguistic motifs; formative events; or key individuals, institutions, and social networks. Documents serving as 'good examples' of particular thematic patterns might also be surfaced. Such analysis might be used both in the development of research questions or hypotheses, and the assembly of materials to support a given argument.
- **Potential positives:** In principle AI could save huge amounts of time in reading and digesting a large corpus of archival documents. In principle, such analysis could offer a more comprehensive and more quantitatively "objective" perspective on a large corpus than any one researcher could assemble in a reasonable amount of time. As such, AI might conceivably also turn up important or interesting patterns that the researcher did not detect. In principle this would allow researchers productive access to corpuses written in languages they do not have.
- **Potential negatives:** It is unclear how much AI systems trained for general use could be relied on to paint an accurate interpretive picture of a particular culturally-situated document base. Between the lack of cultural specificity and the risk of hallucinations, it is hard to see how such analysis could be vested with a great deal of trust. Likewise, much of the value that a historian brings to their work is their particular, "non-objective" point of view. Presumably much of this would be lost through the genericizing effect of AI. More fundamentally, the heart of history as an academic field is the researcher's work embedding themselves in the time, culture, and material artifacts of the field they are studying. It is hard to see how the unique value of the discipline would not be eroded as the work of grappling with primary sources is offloaded to a software system. This use case also presupposes that the archive has been

comprehensively digitized and OCRed. Although increasingly common, this remains the exception rather than the rule in many fields.

- **Possible policies or guardrails:** In general, this use is not recommended. To the extent it was pursued, it might be employed very selectively, in pursuit of particular targeted inquiries, as an additional window onto a source base that has been thoroughly explored by more traditional means. (Having said this, it is reasonable to suppose that the findings from such queries might prompt reassessment of the researcher's own analysis, which points more toward an iterative process but also serious interpretive questions.) This presupposes that the researcher has the skills to do this work without the assistance of AI. One risk that should be acknowledged is that as use of AI as a 'backstop' grows more common, the skills and will to do the work by traditional methods may naturally wane.

Translation (known language)

- **Comments:** Machine translation is an established technology, but AI tools promise to do translation across a greater swath of languages, with greater accuracy and usability. As such, AI could be used to translate entire corpuses of text from a language the researcher knows but lacks rapid fluency, into their primary working language.
- **Potential positives:** In principle machine translation, including AI-facilitated translation, could save the researcher large amounts of time by translating documents into a language they can read with greater speed and facility.
- **Potential negatives:** It is unclear that AI translation is living up to its promises of accuracy and universality. All translation is a deeply nuanced exercise, and history is a field particularly imbricated with those nuances. Moreover, all machine-translation systems, including AI-driven ones, are likely to lack the specific temporal and cultural context of specific historical settings to do reliable translations. As such, the quality and accuracy of AI-driven translation for historical purposes remains highly suspect.
- **Possible policies or guardrails:** In general, this use is not recommended. To the extent it is used, such translations should only be considered a 'rough draft,' perhaps serviceable for navigating and mapping out a text corpus at the first approximation. All serious use of the materials, however, should refer back to the original documents to confirm the accuracy of the translation. This presupposes that the researcher has the skills to accurately translate the texts themselves. One risk that should be acknowledged is that as use of AI as a 'backstop' grows more common, the skills and will to do the work by traditional methods may naturally wane.

Translation (unknown language)

- **Comments:** Machine translation is an established technology, but AI tools promise to do translation across a greater swath of languages, with greater accuracy and usability. As such, AI could be used to translate entire corpuses of text from a language the researcher does not know into one they do know.

- **Potential positives:** In principle machine translation, including AI-facilitated translation, could enable the researcher to access documents that they otherwise lacked the language skills to use – thus broadening their source base, offering a more comprehensive perspective, and perhaps even enabling more extensive cross-community analysis, free from the silos of linguistic ghettoization.
- **Potential negatives:** It is unclear that AI translation is living up to its promises of accuracy and universality. All translation is a deeply nuanced exercise, and history is a field particularly imbricated with those nuances. Moreover, all machine-translation systems, including AI-driven ones, are likely to lack the specific temporal and cultural context of specific historical settings to do reliable translations. As such, the quality and accuracy of AI-driven translation for historical purposes remains highly suspect.
- **Possible policies or guardrails:** In general, this use is not recommended. The researcher simply lacks any ability to confirm the quality of such translations.

Textual to tabular data

- **Comments:** Regularized and tabular data can be useful for doing quantitative analysis. Historically, regularized data often appears in textual form. AI can be used to extract such data into a form suitable for quantitative analysis. The extent of the processing sought from AI might vary widely in such cases. At one extreme, it could simply take data already in tabular form, and convert into a data file format (e.g. CSV) for easy processing. At the other extreme, it could parse highly-free-form textual data, and propose a schema and classification scheme for regularizing it.
- **Potential positives:** AI-driven data processing could save researchers substantial time in terms of data entry. In principle, the AI could also help a researcher less familiar with data analysis methods to uncover proposedly-implicit data structures and classifications to aid them in making their sources amenable to quantitative analysis.
- **Potential negatives:** Hallucinations and errors remain a concern, and cross-checking would be essential. Development of effective schemas and classification systems requires a deep understanding both of the sources and the research questions. Far from adding clarity, a spurious classification scheme can easily muddy the waters, creating findings that appear authoritative but are in fact dubious.
- **Possible policies or guardrails:** Such use may be legitimate, but care should be taken in how much of the analytical work is ceded to the AI. The determination would depend on factors such as how closely the text is to already being organized into a tabular, regularized form; and how readily the AI outputs can be cross-checked for accuracy.

Formulation of research questions or arguments

- **Comments:** A researcher might engage with an AI, perhaps in iterative fashion, to formulate and refine their research questions, or arguments.

· **Potential positives:** In principle, an AI system could augment other mechanisms that historians use to develop their research agendas, such as conference discussions and roundtables, peer-engagement, writing groups, etc. As a system well versed in the academic literature, an AI could theoretically give a researcher a very clear sounding board informed by the historiography of their particular sub-field. Even if, perhaps, AI systems are not generative in the way the human mind is (a disputed proposition), the researcher's creativity might be stimulated through the dialogue.
· **Potential negatives:** A cornerstone of historical scholarship is the process of creative interpretation of the past, from the viewpoint of our current time and individualized perspectives, in earnest dialogue with the artifacts (material, textual, oral, etc.) of the past. AI systems, by definition, encompass the crystallized, amalgamated perspectives of the past historians incorporated in them. As such, they are designed to perpetuate, rather than create. Philosophically, one can argue all scholarship consists of reworking what already exists. Regardless of how one views that premise, however, it may be the case that because AI systems are simply not geared toward the invention of something new, they are more likely to redirect thinking into old channels than to foster creativity.
· **Possible policies or guardrails:** In general, this use is not recommended. The benefits to be gained seem unlikely to outweigh the drag on intellectual creativity.

Writing

· **Comments:** Under the researcher's direction and prompting, an AI system could be used to compose conference presentations, articles, and chapters.
· **Potential positives:** In principle, AI could save the researcher a lot of time in crafting well-written, readable historical analysis. In principle, this would allow non-native speakers to write fluently and persuasively in other languages.
· **Potential negatives:** Writing and analysis are tightly coupled, and it's not clear that the quality of the analysis will survive the outsourcing of the writing work to an AI system. The result might be to create papers that read as finished drafts, but which have the analytical quality of a rough draft. The net effect could be to sap the intellectual value of our work. History is also a field that benefits from a multiplicity of viewpoints. By outsourcing the writing work, there is a risk that the genericized sensibility of the AI will (perhaps in a creeping fashion) overwhelm the researcher's situated perspective.
· **Possible policies or guardrails:** In general, this use is not recommended.

Copy editing

· **Comments:** The researcher might submit a completed work of historical writing to the AI system for the correction of typos, misspellings, grammatical errors, citation formatting mistakes, etc.

- **Potential positives:** This could save the researcher time and increase their confidence that the work they do is technically proficient. In principle, this would allow non-native speakers an added degree of confidence in the technical proficiency of their writing.
- **Potential negatives:** AI hallucinations and errors remain an issue. It is hard to draw a bright line between copy-editing and 'suggestions' or 'corrections' of a more substantive nature.
- **Possible policies or guardrails:** This use seems reasonable, though the AI output should never be taken without review.

Substantive editing / rewriting

- **Comments:** The researcher might submit a rough first draft of their own authorship to the AI system, with the mandate to do a substantial rewrite for style, flow, readability, engagement, mechanics, etc.
- **Potential positives:** This could save the researcher time and increase their confidence that the work they do is readable and professional. In principle, this would allow non-native speakers an added degree of confidence in the quality of their writing.
- **Potential negatives:** AI hallucinations and errors remain an issue. There is considerable risk that an AI rewrite would modify the author's original arguments substantively, elide their personal style in a detrimental manner, and/or flatten their analytical perspective.
- **Possible policies or guardrails:** In general, this use is not recommended.

Feedback on writing

- **Comments:** The researcher might draft of their own authorship to the AI system, not for a rewrite, but rather with a request for substantive and stylistic feedback, as raw material for a manual rewrite.
- **Potential positives:** This could provide the researcher with another channel of feedback on their work, one that might be more responsive and timelier than consultation with a peer.
- **Potential negatives:** There would be some residual risk of the AI steering the author toward a more genericized perspective, though mitigated by the fact that the author is driving the process. It's unclear how much utility this would offer.
- **Possible policies or guardrails:** This use seems reasonable, inasmuch as it might be considered comparable to eliciting the feedback of a colleague.