HSSE Online 13(1) 24-37

Rethinking History Education in the Age of AI

Mathew Lim

National Junior College (Singapore)

Abstract

The rapid proliferation of generative artificial intelligence (AI) has raised questions about the relevance of history education. In response, this paper examines the limitations of AI, particularly its large language models (LLMs), and highlights the enduring educational value of historical thinking. While AI can generate plausible narratives, it often lacks empirical accuracy, interpretive depth, and contextual sensitivity—qualities essential to the discipline of history. Reaffirming history's epistemological foundations, the article argues that the rise of AI amplifies rather than reduces the importance of historical literacy. Historical literacy equips students to interrogate sources, evaluate bias, and navigate content increasingly shaped by algorithms. To support this, four pedagogical approaches are proposed: fostering critical engagement of AI-generated content, using AI tools to support source reading, developing AI literacy through inquiry-based projects, and revisiting historical source work with renewed disciplinary purpose. Cultivating critical, empathetic, and contextually grounded historical thinking is presented as an essential set of skills for preparing students to navigate an AI-mediated world.

Introduction

The emergence of generative artificial intelligence has led some to question the continued relevance of history education. If AI systems can produce sophisticated historical narratives and analyses instantaneously, why maintain traditional history teaching? Predictions that AI could replace human educators within a decade have further fueled these concerns. This perspective, however, fundamentally misunderstands both AI's limitations and history education's essential purpose. Current AI systems rely on pattern-matching rather than genuine understanding, often producing fabricated information, reflecting embedded biases, and presenting decontextualised content that lacks the nuanced interpretation essential to historical inquiry.

Consequently, far from diminishing its importance, the AI era makes history education more crucial than ever. The discipline's emphasis on critical thinking, source evaluation, and contextual understanding provides essential tools for navigating an information landscape increasingly populated by algorithmically generated content. Students need these interpretive skills to stay connected to the full range of human experience—our ability to derive meaning from the genuine experiences of others and to understand ourselves as part of ongoing human conversations and memory-making. This further requires us to distinguish authentic accounts of human experience from superficially authoritative AI outputs that lack the embodied understanding. This shift calls on history teachers to reflect on their

own historical literacy while integrating AI literacy into the history learning experience.

In response to the challenges posed by generative AI, it is important to reaffirm the value of fostering students' capacity for historical understanding. Four approaches are proposed to help strengthen the relevance of historical learning: encouraging critical engagement with AI-generated content; using AI tools to enhance historical reading and interpretation; developing students' AI literacy through inquiry-based projects; and revisiting history source work and historical thinking. It is hoped that, through improved understanding of AI as well as the disciplinary aims of history, teachers can better prepare students to be more ready participants in a digitally mediated world.

Why AI Cannot Replace History Teachers or the Learning of History

Since the emergence of generative artificial intelligence (Gen-AI) in 2023, considerable debate has arisen about its transformative implications for education. Salman Khan (2024: 39–41) optimistically envisions AI as a transformative teaching assistant, capable of creating personalised tutors that adapt to individual students' learning needs and provide specific, tailored feedback. This helps broaden educational opportunities for students who may not have access to qualified teachers or supportive learning environments. In Singapore, the Ministry of Education has integrated AI features into the Singapore Student Learning Space (SLS), an online platform used across schools. These features include the Learning Assistant, a customisable chatbot, a pilot adaptive learning system for Mathematics and Geography, a Feedback Assistant that provides timely, automated feedback and assigns marks based on teachers' 'context prompting', and a Data Assistant that analyses and summarises students' responses These tools offer multiple benefits, particularly in enhancing engagement and supporting personalised learning. The customisable chatbot, for instance, can enhance engagement by simulating historical figures for interactive role-play. Adaptive systems, though not yet available for history, could adjust content based on learners' needs, helping to close gaps and potentially support differentiated learning.[i]

In the long run, Bill Gates has predicted that AI could supplant many roles currently filled by human educators within the next decade, citing AI-driven tutoring systems as likely alternatives (Huddleston, 2025). A similar view was espoused earlier by historian Anthony Seldon, who anticipated that intelligent machines would begin replacing teachers in classrooms within a decade (von Radowitz, 2017).

However, the notion that AI can replace teachers underestimates the limitations of AI systems, particularly large language models (LLMs), when applied to the teaching of history. History is not simply a collection of facts to be retrieved and reframed by an algorithm. It is a discipline rooted in critical interpretation, contextual sensitivity, and empathetic engagement with the complexity of human experience. Unlike human historians, AI models do not truly understand the materials they process; they identify statistical patterns and assemble plausible narratives without grasping context, bias, or nuance—elements that are essential to the meaningful study of the past.

Inherent Limitations of LLMs

Among current AI systems, large language models (LLMs) are the most prominent due to their capacity to generate

human-like text, which has a direct impact on how educational content is produced and utilised. These models are capable of generating historical texts and narratives by drawing on extensive datasets. However, they operate by identifying patterns and producing a "reasonable continuation" of existing text, rather than genuinely understanding context, nuance, or underlying meaning (Wolfram, 2023). This limitation is often described as the "stochastic parrot" phenomenon, where LLMs produce seemingly coherent text without true comprehension and lack the [ii]"communicative intent" characteristic of human beings (Bender et al., 2021: 616). Furthermore, these systems lack the capacity for causal reasoning. As Chomsky et al. (2023) argue, LLMs are constitutionally incapable of distinguishing between correlation and causation. They excel at describing and predicting based on data patterns but cannot explain the underlying causal mechanisms. These limitations pose significant challenges when LLMs are deployed in educational contexts, where an understanding of causality can play an important role in fostering critical thinking and supporting deeper learning, particularly in disciplines that focus on causal reasoning.

Because LLMs rely on probabilistic pattern-matching rather than grounded and causal understanding, they can generate inconsistent outputs and may even "hallucinate" information. Hallucinations arise when models predict plausible-sounding content without verifying it against any factual source. These outputs often take the form of fictitious statistics, studies, or historical events that appear authoritative but are entirely fabricated. This makes them particularly difficult to identify without rigorous verification (Dahl et al., 2024), which explains why such tools are unreliable for producing consistent or meaningful historical analysis.

Furthermore, the absence of genuine real-world understanding in LLMs not only leads to hallucinations but also contributes to the reproduction of biases embedded in their training data. These biases may subtly manifest as stereotyped portrayals, unequal treatment of demographic groups, or skewed historical interpretations. Emily Bender et al. (2021) warn that increasing the scale of these models without addressing underlying data biases and ethical concerns risks perpetuating and amplifying harmful inaccuracies.

Compounding these challenges is what Jeffrey Yost (2023) terms "dual decontextualisation," a phenomenon where AI systems lose both the historical context of the material they process and the data context of their training sources. This means that LLMs frequently present historical material stripped of essential cultural, temporal, and situational contexts, which can lead to oversimplified or distorted representations that undermine the richness and complexity of historical understanding. At the same time, the provenance of the training data, including details such as the time and place of creation, authorship, and the circumstances under which it was produced, is routinely obscured (Bender et al., 2021: 615). As observed by Huang and Chang (2024), training methodologies that aggregate vast volumes of text from diverse and sometimes incompatible sources contribute to a broader loss of traceability and source attribution in LLMs. Building on this, we can further establish that critical information such as original publication dates, authoritative source distinctions, and intended audience nuances is typically lost during pre-training. Furthermore, LLM-generated outputs may occasionally conflate temporally or thematically inconsistent material, which weakens the accuracy and integrity of historical representation.

A further problem of provenance arises when LLMs are explicitly asked to account for the origin of their sources. Because LLMs are probabilistic text generators, they produce outputs by predicting the most likely sequence of words based on statistical patterns learned from training data. When prompted to provide citations, the model attempts to generate text that resembles a citation, drawing on patterns of how references typically appear in its training corpus. As a result, it may fabricate references — complete with plausible-sounding author names, article titles, journal names, and publication dates — that are entirely fictional. In early 2023, I demonstrated this problem by prompting the model to generate historical sources on the Maria Hertogh riots, only to find that the citations were fabricated (Lim, 2023). This problem of fabricated sources has persisted to varying degrees, even with improvements in language models, as noted by several studies and news reports.

In response to the limitations of LLMs regarding provenance, attribution, and contextual accuracy, one popular method known as Retrieval-Augmented Generation (RAG) enhances the outputs of LLMs by incorporating external knowledge bases such as Wikipedia. This allows models to ground their responses in verifiable sources, typically presented as footnotes or hyperlinks. Popular chatbots like Perplexity.ai, Gemini, and ChatGPT all use RAG to match LLM outputs with online metadata, thereby helping to provide AI-generated content supported by traceable sources.

However, RAG systems do not fundamentally resolve the challenges associated with LLMs. Rather, they reflect and extend the underlying problems of statistical pattern-matching and probabilistic aggregation. Contextual errors can still arise when there are mismatches between the retrieved information, the model's internal representations, and the user's intent, resulting in inconsistencies between source material and generated output. Moreover, important contextual nuances may be omitted during the retrieval and integration process, distorting meanings in ways analogous to quoting sources out of context (Wong et al., 2025).

Additionally, RAG's effectiveness depends critically on the quality, reliability, and accessibility of external sources. Because it retrieves information rather than verifying or deeply interpreting it, the system inherits any biases, inaccuracies, or omissions present in those sources. When retrieved documents are unreliable, one-sided, or poorly verified, RAG may reinforce misinformation or distortions rather than correct them (Wong et al., 2025). Restricted access to subscription-based or proprietary databases further creates knowledge gaps, particularly in specialised academic, legal, or technical domains where authoritative information often lies behind paywalls.

Therefore, while RAG provides more traceable provenance and an impression of accuracy, it remains vulnerable to the quality and availability of the information it retrieves. It also suffers from semantic misalignment due to the inherently probabilistic nature of both LLMs and the retrieval process, as these systems rely on statistical associations and relevance rather than true semantic understanding.

Overall, despite advances in AI technologies such as LLMs and RAG, their inherent limitations reinforce the indispensable role of history education. These technologies present significant challenges for history education by altering how history is read, written, taught, and understood. These models increasingly adopt an authoritative tone that mimics

scholarly writing while lacking the accountability and factual grounding of genuine expertise. This veneer of authority can mislead readers into accepting information uncritically and anthropomorphising the models, promoting reliance on oversimplified, instant outputs, and fostering a 'crutch mentality' that discourages deeper engagement with primary sources. A recent MIT study warns that such overreliance may create cognitive debt, where users progressively outsource thinking processes to AI models and lose their capacity to critically evaluate or generate content independently (Kosmyna et al., 2025). This undermines the cognitive benefits of actively 'doing' history, which demands sustained attention, analytical and critical reasoning, and the integration of multiple perspectives and historical contexts over time. These activities engage executive functions such as working memory, cognitive flexibility, and evaluative judgement, all essential to higher-order thinking and strengthened through historical practice. Beyond these cognitive concerns, LLMs obscure the provenance and intent behind historical documents, reducing rich, contextualised sources to flattened patterns of language and weakening readers' capacity to grasp the situated intentions and inner thinking of historical actors. For learners, such decontextualised narratives provide only shallow representations of historical figures and their situated and lived experiences, eroding the development of critical and empathetic historical understanding of individuals' minds as shaped by the contexts of the past. The more profound consequence is that persuasive, but unreliable AI-generated content promotes passive information consumption over active inquiry. In this evolving landscape, history education becomes increasingly indispensable, serving not merely to preserve interpretive rigour but to cultivate the critical thinking, contextual awareness, and meaningful engagement with the past that are essential in an age of algorithmic knowledge.

The Enduring Value of History Education in the Age of AI

History education fosters essential skills for navigating this complex information landscape. Through historical thinking, students learn to question information critically and assess its provenance. They develop the ability to evaluate reliability within social and temporal contexts. As Sam Wineburg explains, this includes interpretive skills such as identifying bias, analysing motives, and comparing narratives to form informed judgments (1991: 498–499). These competencies are vital in an era where digital content can mimic authority while concealing distortion.

One important aspect of historical thinking involves assessing bias and veracity. It employs specific methods such as sourcing, corroboration, and contextualisation to interrogate not only the content of a source but also its origins, purpose, and the conditions under which it was produced (Fitzgerald, 1983; Wineburg, 1991: 510–512; 2018: 173–177). It should also be added that although these methods demonstrate rigorous scrutiny of evidence, they should not be confused with the processes of 'verification' as advanced by logical positivists or 'falsification' as popularised by Karl Popper in scientific reasoning (Popper, 2002: 20). While historical and scientific inquiry may share specific procedures, such as evaluating evidence and assessing claims, their underlying epistemologies differ significantly. Scientific reasoning, particularly in the positivist tradition, is focused on the nomothetic aim of seeking generalisable laws through rigorous examination of empirical evidence. Historical reasoning, by contrast, is

grounded in interpretivism and emphasises contextual understanding and the recovery of human meaning.

Historical inquiry, at its core, rests on an interpretivist foundation; historians view knowledge as shaped by temporal context and embedded meaning. It requires perspective-taking and what David Stockley terms "empathetic reconstruction"—the attempt to understand past beliefs, motives, and actions through their own frames of reference, grounded in evidence (Stockley, 1983: 53–55; Seixas, 2015: 9-10). This approach aligns epistemologically with the concept of *verstehen*. Here, understanding human action requires reconstructing the meanings individuals attached to their behaviour within their unique context; this idea was developed by Max Weber and discussed by Stockley (Stockley, 1983: 53–55).

Importantly, historical sources are not merely texts to be decoded for factual content; they serve as windows into the lives, beliefs, and conflicts of people in the past. Engaging with them becomes an active and interpretive process—a conversation across time that requires students to bridge their own perspective with that of historical actors. This interpretive encounter involves what Hans-Georg Gadamer describes as a "fusion of horizons" between the present-day reader and the historical source (Gadamer, 2004: 305). Such engagement requires treating sources not as inert empirical data, but as products of human authorship, shaped by the author's intentions, worldview, audience, cultural norms, and historical circumstances

Robin Collingwood reinforces this view, insisting that history is the "reenactment of past thought in the historian's own mind" (Collingwood, 2005: 215). Engaging with historical sources demands that the historian does not merely record external events, but actively "discern the thought of its agent" (Collingwood, 2005: 213). In this sense, Collingwood argues that the past is not something to be observed from a distance like a spectacle, but something to be understood from within. This requires the historian to reconstruct the intentions, reasoning, circumstances, and cultures that shaped a person's decisions and actions as part of their lived experience (Wineburg, 1991; 2018: 173–177). From this perspective, historians play a critical role as what William Sewell calls "theoreticians of temporality", analysing how temporal contexts shape the lives of people in the past rather than merely recounting events in chronological order (Sewell, 2005: 6). Historical knowledge, then, is inseparable from the interpretive process by which the historian reenacts and critically engages with past thought.

Drawing on the above insights, Richard Bernstein emphasises that reading historical sources and doing history is not only an interpretive act but also a moral practice of learning and self-reflection. For Bernstein, understanding history becomes a means of challenging one's assumptions and deepening judgement. It fosters responsible engagement with contemporary life and encourages self-reflection. This mode of inquiry functions as both a moral and intellectual exercise in openness (Bernstein, 1983: 143). Such an approach is particularly valuable in breaking through echo chambers and assumption bubbles reinforced by digital technologies, including AI, which can limit understanding.

Sustaining this moral and intellectual openness requires individuals to remain active agents in the interpretive process, engaging critically with historical sources rather than passively accepting pre-digested narratives. This kind of careful scrutiny

becomes especially vital in the age of generative AI, where content is increasingly produced through algorithmic processes that simulate authority but lack historical grounding. These outputs offer 'flattened' versions of the past, stripped of the contextual depth necessary for meaningful historical understanding.

This imperative to cultivate critical historical thinking resonates with Singapore's history curriculum, which defines history as a "thinking discipline" that fosters reasoning, empathy, and historical perspective (Afandi and Lim, 2022: 394; Ministry of Education, 2023: 10–12). Within this framework, students learn not only about the past but also how historical meaning is constructed and contested. They examine how these interpretive processes unfold within specific temporal and cultural contexts (Afandi and Lim, 2022: 395).

Suggested Approaches to Teaching History

Given the enduring importance of historical education, how should history be taught in a landscape increasingly shaped by artificial intelligence? History teaching must evolve not by abandoning traditional humanistic methods, but by integrating digital literacy with critical, contextual engagement. This requires rethinking pedagogy, assessment, and classroom practices to ensure that students develop the skills to interpret, analyse, and critique both historical content and digitally produced text. The following suggestions outline ways to enhance the teaching of history in the age of AI.

1. *Encouraging Critical Engagement with AI Outputs*

In approaching history as a discipline grounded in interpretation, evidence, and contextual understanding, students must be introduced to the notion that historical knowledge is constructed, not merely retrieved. This epistemological awareness is particularly crucial when engaging with AI-generated outputs. Rather than accepting such AI content at face value, students should learn to critically examine its origins, biases, and underlying assumptions. Questions such as "Whose perspective is being represented?", "What is the historical context of this interpretation?", "What are the assumptions?" and "Who is likely to present this view?" should become part of how students naturally think when reading historical texts. Students should understand how large language models (LLMs) generate their outputs, which are based on training data and algorithms and are probabilistic, which means that they are not authored by a person and are not consistent or accountable in the way human-written sources are. More importantly, the information generated is often decontextualised and 'de-provenanced' from their original sources. While generative AI tools may provide quick and convenient information, these should serve only as starting points for deeper inquiry. Teachers should train students to scrutinise AI responses using the methods of historians. Comparing AI-generated content with other print or digital sources would help students identify gaps, biases, assumptions, perspectives, as well as the contextual roots of certain viewpoints. Practical classroom activities could include asking students to evaluate AI-generated outputs by comparing them with a range of primary and secondary sources. An example of this approach is the use of Character.ai, which allows students to interact with AI-generated historical personas. While such tools may produce inaccuracies, they offer opportunities for students to practise source verification and deepen their historical understanding by

cross-referencing them with other sources (Lee, 2023). This would help students corroborate these outputs with established evidence or point out inaccuracies, biases, and missing perspectives.

2. Reading Sources with AI Assistance

As discussed, historians rely on disciplined interpretive practices to interrogate sources. In his later work, Wineburg introduced the method of *lateral reading*, which involves leaving a website or source to consult other materials, including digital sources, allowing readers to evaluate credibility by drawing on a wider informational context (Wineburg, 2018: 150–151). Lateral reading further supports a deeper understanding of context by encouraging readers to situate individual claims within the wider networks of knowledge, perspective, and evidence. Such habits of mind are essential for navigating a digital landscape where misinformation often presents itself in the guise of authority (Wineburg & Caulfield, 2023: 221–222).

In today's digital and AI-rich environment, students and teachers can read sources with the support of generative AI, which can also enhance lateral reading practices. AI's semantic search capabilities—used in tools like ChatGPT, Perplexity.ai, and Google Gemini—facilitate the discovery of relevant contextual information, streamlining what once required extensive searching and cross-referencing. These tools support lateral reading by helping users locate relevant materials more efficiently, though they must be used with discernment. Additionally, AI could serve as a form of "co-intelligence", a notion proposed by Ethan Mollick (2024), suggesting that AI can act as a highly capable collaborator that enhances human thinking and writing by offering suggestions and explanations while leaving interpretation and judgment to the human reader. In historical education, AI can clarify terminology, provide background context, and suggest related sources, thus reinforcing students' understanding of both content and context.

Nevertheless, history teachers must ensure AI does not replace the interpretive work of historical reading. While AI can provide suggestions, such as clarifying terminology or suggesting sources, it does not grasp meaning or context in the way a human reader does and may oversimplify complex events or reflect biases from its training data. History teachers remain vital in guiding students through the careful reading process, fostering their ability to engage with historical sources critically, attentively, and with contextual awareness.

3. Developing AI Literacy Through History Investigation Projects

In the Singapore history curriculum, students are required to undertake History Investigation Projects guided by the principles of inquiry-based learning (Ministry of Education, 2021). These projects are designed to move students beyond rote memorisation, encouraging them to view history as a discipline of interpretation and evidence.

At the beginning of the project, teachers should explicitly communicate to students the learning objectives. The students should understand that through the inquiry process, they can explore and comprehend the lived experiences of people in the past, as well as the political, economic, and social conditions that shaped those experiences. Achieving this requires meaningful engagement with the research process, including examining primary and secondary sources, understanding their context, evaluating their credibility, and drawing informed conclusions. Teachers

should emphasise to students that the value of the project lies not only in meeting its intellectual demands, but also in the fulfilment and personal growth that come from engaging in authentic historical inquiry. To illustrate this, I previously created and showed a video titled *'Learning History in the Age of ChatGPT'*, using an AI avatar and voice to prepare my students for their Historical Investigation Project. [iii]The video highlights the limitations of relying solely on AI for historical understanding and underscores the importance of developing critical thinking through direct engagement with historical sources and perspectives (Lee, 2023).

Teachers should guide students to be mindful of the limitations of AI tools during the inquiry process. While generative AI tools may offer helpful background information or clarify terms, they often lack the capacity to interpret historical nuance and context. Overreliance on AI can lead to shallow or misleading conclusions. This includes students using AI to write extensively, often bypassing critical stages of the thinking and composition process. It thus deprives students of the opportunity to develop independent thinking, analytical rigour, historical literacy, and a deeper insight into the past. To use AI effectively, students should be guided through a structured writing process that preserves ownership of their work. For example, they could be advised to produce an initial draft independently, without using any AI tools. Once completed, they might use AI for suggestions or feedback. Students should then reflect on this input and discuss it with their teachers or peers. Finally, they should repeat the process by writing a second draft without AI assistance. This ensures that AI remains a supportive tool rather than a substitute for critical thinking or authentic writing. This "Brain-to-LLM" approach is supported by the aforementioned MIT research, which suggests that the greatest cognitive gains occur when learners first engage in tasks independently, using only their own thinking, before supplementing their efforts with AI support (Kosmyna et al., 2025). These findings underscore the pedagogical value of encouraging students to think and write unaided before turning to AI for refinement and feedback.

Figure 1. Image generated by ChatGPT that was shown to students

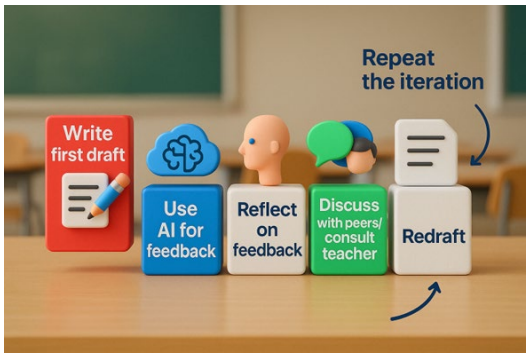

In addition, students must be cautioned against using AI to generate citations. Contrary to the advice of some institutions and publications, tools like ChatGPT should not be cited as sources due to their probabilistic nature and the absence of verifiable authorship or provenance. More specifically, as outputs generated by algorithms rather than authored by individuals, AI-generated content lacks identifiable origin, context, intentionality, and any basis for accountability or redeemability. Students should instead be guided to consult, understand, and reference credible, non-AI sources directly.

While today's more advanced and popular chatbots, such as Perplexity.ai, ChatGPT, Gemini, and Microsoft Copilot, now include hyperlinks or citations drawn from internet-based content, students must take full responsibility for locating, reading, and critically evaluating the original materials themselves, and should cite only those primary or secondary sources directly,

not the chatbot's summarised or synthesised versions.

To help teachers and students navigate the appropriate boundaries of AI use, the AI Assessment Scale (AIAS), introduced by Perkins et al. (2024), provides a structured framework that classifies student engagement with generative AI across five levels—from minimal to extensive involvement. The AIAS outlines five levels: starting with no AI use at all; then using AI to generate ideas and structures; followed by AI-assisted editing; then AI-generated content evaluated by students; and finally, full AI-generated work with minimal student input and transparency.

Frameworks like AIAS are not prescriptive, but serve to guide teachers in communicating to students how much AI use is acceptable, and the rationale behind these expectations. Teachers could provide these AI guidelines to students when working on their History Investigation Projects:

- Use AI to brainstorm, but not for generating full responses.
- Read and make references to all sources. Avoid using AI to generate references and do not cite AI as a source.
- Use AI to correct grammar and spelling errors, but not to write on your behalf.
- Keep a log of your AI use, including screenshots of prompts and outputs.
- Explain whether and how you have used AI in your reflection.

These suggested guidelines help ensure that AI is integrated ethically and effectively in the learning process, while preserving the integrity, purpose, and critical engagement that historical inquiry demands.

4. Revisiting Source Work and Historical Thinking

The rise of artificial intelligence is rapidly reshaping the way knowledge is produced, distributed, and consumed. In this evolving landscape, students and teachers are no longer just recipients of information but co-producers of knowledge, often in collaboration with AI tools. This transformation raises urgent questions about what it means to think historically, and why such thinking matters in an already digitised and informationalised world that is becoming increasingly saturated with AI-generated content.

In Singapore, historical source analysis is often taught with a heavy emphasis on exam preparation, relying on rigid and formulaic scaffolds. While these strategies may seem to be a stop-gap to better performative outcomes, they lack interpretive depth and reduce students' inquiry of sources to a mechanistic checklist. Such 'pedagogies' have perpetuated a narrow and instrumental view of history teaching (Afandi & Lim, 2022: 394).

Various institutional and systemic factors that have contributed to these reductionistic pedagogical practices among teachers, but there is now an urgency for teachers to revisit their disciplinary knowledge and renew their commitment to fostering reasoning, interpretation, and judgment, to prepare students to navigate an AI-infused world. Reconnecting with these disciplinary foundations is essential—not only to honour the intellectual integrity of history but to ensure our students are genuinely future-ready.

More than four decades ago, Peter Lee

(1983) observed that many teachers in Britain were not sufficiently engaged in the philosophical foundations of history education because they perceived it as a time-consuming pursuit that lacked practicality (Lee, 1983: 20). Yet his argument that teachers must be thinkers who critically examine the nature of historical knowledge and guide students accordingly remains relevant. Today's epistemic challenges, shaped by algorithmic mediation, disinformation, and content saturation, make Lee's call for deeper disciplinary engagement feel more urgent than ever.

Teachers must critically re-examine the philosophical foundations of history education: What is history for? Why and what does it mean to analyse a source? How do we teach students to engage with the past? Historical thinking must be anchored in deeper epistemic reflection about evidence, interpretation, and perspective, without which classroom practice risks becoming uncritical and detached from the essence of the discipline (Lee, 1983: 20–21, 28–29). In this age of AI, the onus is on teachers not simply to transmit content or to expose students to historical approaches in a superficial manner, but to cultivate the cognitive and interpretive habits that enable students to interrogate information, discern meaning from it, and situate it within appropriate historical contexts. The challenge, then, is to move beyond procedural proficiency and foster in students a historically grounded disposition that resists superficial thinking and embraces deeper disciplined inquiry. This calls for a renewed focus on the essence, foundations, and purpose of history as a discipline, along with the corresponding theoretical underpinnings of historical pedagogy. Such a focus strengthens teachers' disciplinary understanding and, in turn, enables them to better equip students with the critical faculties and information literacy necessary to navigate an AI-infused information landscape

While this commitment to examining disciplinary history and its philosophical foundations may have previously seemed exhausting to teachers, it is now made more feasible due to the ready availability of AI-powered tools. With these tools, the barriers to accessing complex ideas are reduced, thereby augmenting the reading and learning experiences of teachers. They can upload documents and texts to platforms such as Google NotebookLM, Microsoft Copilot, ChatGPT, or Claude by Anthropic (not exhaustive and ever-increasing), and then explore complex historical and philosophical ideas through natural language dialogues. These tools can provide definitions, explanations, contextual insights, challenge assumptions, and more, allowing users to *read laterally* as they interrogate, probe and reflect on, and acquire knowledge with speed, clarity, and discernment. Such support empowers and equips teachers to deepen their disciplinary knowledge, which would translate into more meaningful and purposeful teaching of history.

Conclusion

It is undeniable that the advent of Generative Artificial Intelligence will shape the education landscape for the foreseeable future. Yet, the rise of artificial intelligence does not signal the end of history education, nor does it mean the loss of its importance. As this article has demonstrated, AI's inherent limitations, including hallucination, bias, and decontextualisation, make the critical thinking skills fostered by historical education more essential than ever. Students navigating information landscapes increasingly populated by algorithmic content require the interpretive tools that only rigorous historical training can provide.

However, recognising these challenges is insufficient. This article seeks to explore the relationship between AI and historical literacy, in the hope of helping teachers understand their convergence, overlaps, and how the two seemingly distinct spheres could be mediated. At the heart of it, history educators must actively embrace their role as both guardians of disciplinary integrity and architects of pedagogical innovation. The four suggestions outlined in the article offer possible pathways forward, but there could be other areas and aspects to explore and examine.

Finally, it has to be stressed that in this age of AI, the distinctly human capacities for critical thinking, empathetic understanding, and informed judgment that history education cultivates are not going to be relics of a pre-digital past but essential tools for an uncertain future.

Author's note and acknowledgements

Given the rapid evolution of AI technology and its transformative impact on education, I have chosen to examine the fundamental principles of LLMs and history and provide some broad suggestions, rather than concrete pedagogical examples in this paper. It is hoped that, in times of such technological change, theoretical and generalised discussions may provide more enduring value.

I would like to thank the Principal of National Junior College, Lucy Toh, as well as the Head of the Humanities Head of Department, Valerie Cher, for supporting my exploration of AI. This includes my participation and sharing on AI across various school, History Networked Learning Communities (NLCs), and Educational Technology Division (ETD) platforms. I would also like to thank my colleague, Jane Choong, with whom I have had professional conversations about various issues concerning history as a discipline, as well as for our collaboration in facilitating an 'Unconference' on a similar topic during the History Symposium 2024. These experiences have deepened my research and reflection on the topic and have been instrumental in shaping the writing of this paper.

References

Afandi, S., & Lim, I. M. (2022). History education in Singapore: Development and transformation. In Lee, Y.-J. (Ed.), *Education in Singapore: People-making and nation-building* (pp. 379–401). Singapore: Springer.

Bender, E. M., Gebru, T., McMillan-Major, A., & Shmitchell, S. (2021). On the dangers of stochastic parrots: Can language models be too big? In: *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency. (FAccT '21).* New York: Association for Computing Machinery, pp.610–623.

Bernstein, R. J. (1983). *Beyond objectivism and relativism: Science, hermeneutics, and praxis.* Philadelphia: University of Pennsylvania Press.

Chomsky, N., Roberts, I., & Watumull, J. (2023). The false promise of ChatGPT. The New York Times. Retrieved from https://www.nytimes.com/2023/03/08/opinion/noam-chomsky-chatgpt-ai.html

Collingwood, R. G. (2005). *The idea of history: Revised edition with lectures 1926–1928* (J. Van Der Dussen, Ed.). Oxford: Oxford University Press.

Dahl, M., Magesh, V., Suzgun, M., & Ho, D. E. (2024). *Hallucinating law: Legal mistakes with large language models are pervasive.* Stanford Human-Centered Artificial Intelligence. Retrieved from

https://hai.stanford.edu/news/hallucinating-law-legal-mistakes-large-language-models-are-pervasive

Fitzgerald, J. (1983). History in the curriculum: Debate on aims and values. *History and Theory*, 22(4), 81–100.

Gadamer, H.-G. (2004). *Truth and method* (2nd rev. ed., Weinsheimer, J., & Marshall, D. G., Trans.). London: Continuum.

Huang, J., & Chang, K. C.-C. (2024). *Citation: A key to building responsible and accountable large language models*. arXiv. Retrieved from https://arxiv.org/html/2307.02185v3#S3

Huddleston, T. Jr. (2025). Bill Gates: Within 10 years, AI will replace many doctors and teachers—humans won't be needed 'for most things'. *CNBC News*. Retrieved from https://www.cnbc.com/2025/03/26/bill-gates-on-ai-humans-wont-be-needed-for-most-things.html

Khan, S. (2024). *Brave new words: How AI will revolutionize education (and why that's a good thing)*. New York, NY: Viking.

Kosmyna, N., Hauptmann, E., Yuan, Y. T., Situ, J., Liao, X.-H., Beresnitzky, A. V., Braunstein, I., & Maes, P. (2025). *Your brain on ChatGPT: Accumulation of cognitive debt when using an AI assistant for essay writing task*. arXiv. Retrieved from https://arxiv.org/abs/2506.08872

Lee, P. J. (1983). History teaching and philosophy of history. *History and Theory,* 22(4), 19–49.

Lee, S. (2023). This teacher's got his eye on AI. *Schoolbag*. Ministry of Education, Singapore. Retrieved from https://www.schoolbag.edu.sg/story/this-teacher-s-got-his-eye-on-ai/

Lim, M. (2023). Sources generated by ChatGPT. Medium. Retrieved from https://medium.com/@historylabsg/caveat-emptor-fake-historical-sources-generated-by-chatgpt-a49a82a44d22

Ministry of Education, Singapore (2023). *Upper secondary history syllabus.* Curriculum Planning & Development Division. Retrieved from https://www.moe.gov.sg/-/media/files/syllabus/2023-upper-sec-exp-na-history.pdf

Ministry of Education, Singapore (2021). *History syllabus: Lower secondary.* Curriculum Planning & Development Division. Retrieved from https://www.moe.gov.sg/-/media/files/secondary/syllabuses/humanities/archive---2021-history-lower-secondary-syllabus.pdf

Mollick, E. (2024). *Co-intelligence: Living and working with AI.* New York: Portfolio/Penguin.

Perkins, M., Furze, L., Roe, J., & MacVaugh, J. (2024). The Artificial Intelligence Assessment Scale (AIAS): A framework for ethical integration of generative AI in educational assessment. *Journal of University Teaching and Learning Practice*, 21(6).

Popper, K. R. (2002). *The logic of scientific discovery.* London and New York: Routledge.

von Radowitz, J. (2017). Intelligent machines will replace teachers within 10 years, leading public school headteacher predicts. *The Independent*. Retrieved from https://www.independent.co.uk/tech/intelligent-machines-replace-teachers-

classroom-10-years-ai-robots-sir-anthony-sheldon-wellington-college-a7939931.html

Seixas, P. (2015). A model of historical thinking. *Educational Philosophy and Theory,* 47(6), 614–622.

Sewell, W. H. Jr. (2005). *Logics of history: Social theory and social transformation.* Chicago: University of Chicago Press.

Stockley, D. (1983). Empathetic reconstruction in history and history teaching. *History and Theory,* 22(4), 50–65.

Wineburg, S. (1991). On the reading of historical texts: Notes on the breach between school and academy. *American Educational Research Journal,* 28(3), 495–519.

Wineburg, S. (2018). *Why learn history (when it's already on your phone).* Chicago: University of Chicago Press.

Wineburg, S., & Caulfield, M. (2023). *Verified: How to think straight, get duped less, and make better decisions about what to believe online*. Chicago: University of Chicago Press.

Wolfram, S. (2023). What is ChatGPT doing … and why does it work? Retrieved from https://writings.stephenwolfram.com/2023/02/what-is-chatgpt-doing-and-why-does-it-work/

Wong, A., Jiang, J., Sang, J., Shen, J., Vyas, M., Sadat, M., Zhu, Y., Bacchus, W., Shaikh, A., Rajani, N., & Arnaout, R. (2025). *Retrieval-augmented systems can be dangerous medical communicators*. arXiv. Retrieved from https://arxiv.org/abs/2502.14898

Yost, J. (2023). Exploring AI's dual-decontextualization, and why metaphors matter. Presented at the American Historical Association Congressional Briefing, Dirksen Senate Office Building, Washington, D.C., 20 October. Retrieved from https://www.historians.org/wp-content/uploads/2024/04/AI-Handout.pdf

[i] The full list of AI features for the SLS, along with instructions, can be found here: https://www.learning.moe.edu.sg/teachers/teaching-and-learning-on-sls/aied-features

[ii] With recent developments, Large Language Models (LLMs) are becoming multimodal, giving rise to Multimodal Large Language Models (MLLMs). These systems process and generate outputs across different types of data, such as text, images, and audio, leading some to suggest they offer deeper comprehension. However, despite this broader input, MLLMs still rely on statistical correlations rather than true understanding.

[iii] The video is available at the following link: https://www.youtube.com/watch?v=5796waYGsvc