

Teaching English Literature in the Age of Artificial Intelligence: New Pedagogies and New Challenges

Ainum Suhail

English Language Instructor, University College of AD Dayer,
Jizan University, Saudi Arabia

Published on 16-10-2026; DOI:10.36993/RJOE.2026.11.130

Abstract

This paper examines how the rapid diffusion of generative artificial intelligence (GenAI) is reshaping the teaching of English literature in schools and universities. Drawing on recent scholarship in AI literacy, academic integrity, and humanities pedagogy, it argues that literary study is distinctively positioned to respond productively to GenAI, precisely because interpretation, ambiguity, and cultural context resist algorithmic codification even as GenAI lightens mechanical text-processing work. The paper surveys five emerging pedagogical practices — AI-assisted close reading, prompt literacy as a form of critical literacy, adversarial or comparative reading exercises that set human interpretation against machine output, co-creative writing pedagogy, and AI-supported accessibility scaffolding — before turning to five interlocking challenges these practices raise: academic integrity and assessment validity, the erosion of interpretive authority and critical thinking, bias and cultural homogenization in AI-generated readings, unequal access to AI tools, and the practical difficulty of redesigning assessment. A review of empirical evidence, including a controlled study comparing AI-generated and student essays on Old English poetry, grounds the discussion in classroom-level findings rather than speculation. The paper closes with a practical framework for literature departments and individual instructors seeking to integrate GenAI without diminishing the discipline's core commitments to close reading, cultural and historical context, and independent critical judgment.

Keywords: Generative Artificial Intelligence (GenAI); English Literature Pedagogy; AI Literacy; Critical Thinking; Academic Integrity.

Introduction

The arrival of widely accessible generative artificial intelligence (GenAI) systems, from ChatGPT's public release in late 2022 through the more capable

multimodal and reasoning-oriented models that followed, has unsettled the teaching of English literature as thoroughly as any technological development since the spread of the internet itself. Literature classrooms have always depended on students producing their own written interpretation of a text: the essay, the seminar response, the annotated close reading.

GenAI systems can now produce a plausible, often fluent, version of that same interpretive writing in seconds, on demand, for a fraction of the effort a student would otherwise invest. This capability arrived without an accompanying consensus on how literature teachers should respond, and departments across secondary and higher education have spent the past several years improvising policy, pedagogy, and assessment design in real time.

This disruption is not uniform across disciplines. Fields with well-defined problem solution structures, such as mathematics or computer science, have faced their own version of the GenAI challenge, but English literature occupies a distinctive position. Interpretation of a poem, novel, or play depends on ambiguity, cultural and historical context, and the reader's own subjective judgment in ways that resist straightforward algorithmic codification. GenAI systems trained to predict plausible continuations of text can produce fluent description and generic thematic commentary, but a growing body of classroom research suggests that machine-generated literary analysis tends to describe rather than analyze, and to flatten the very ambiguity that literary study exists to explore. This resistance is not simply a limitation to be waited out; it is a pedagogical opening. Literature teaching, more than most disciplines, can use the visible gap between AI-generated and human interpretation as a teaching tool in its own right.

This paper has three aims. First, it surveys the pedagogical practices that literature educators are developing in response to GenAI, treating these as genuine innovations rather than stopgap accommodations. Second, it examines the challenges these practices raise, including but not limited to academic integrity, so that the discussion does not collapse into a narrow conversation about plagiarism detection. Third, it proposes a practical framework that literature departments and individual instructors can adapt to their own contexts. The paper draws on recent empirical studies of GenAI use in humanities classrooms, including controlled comparisons of AI-generated and student written literary analysis, alongside the broader literature on AI literacy in education. Throughout, the argument is that the goal of literature pedagogy in the age of AI is not to defend the discipline against a new technology,

but to clarify, through direct engagement with that technology, what literary reading and interpretation actually require of a human reader.

Background and Literature Review

Scholarship on AI in education has expanded rapidly enough that surveying even its recent output requires selectivity. Two strands are most relevant to literature teaching: research on AI literacy as a general educational competency, and research specific to the humanities and to English studies.

AI literacy frameworks

Early frameworks for AI literacy, developed before the generative AI boom, focused on how AI systems work technically: what a decision tree or neural network is, how training data shapes output. The arrival of large language models requires substantial reframing, since generative systems behave less like a fixed algorithm to be understood and more like an unpredictable conversational partner to be evaluated. Recent frameworks accordingly define AI literacy as the capacity to critically engage with AI tools, including understanding their functions, limitations, and ethical implications, rather than simply knowing how they are built. Institutions have adopted established models of teacher knowledge, such as the TPACK framework linking content, pedagogy, and technology, to this generative AI context, and reviews of the literature consistently report that professional development access, rather than technical sophistication, is the strongest predictor of confident and appropriate classroom use.

A recurring finding across this literature is that AI literacy is not a single, agreed-upon skill set. It ranges from young children learning that a chatbot is not a person, to university students practicing prompt engineering, to instructors learning to redesign assessment.

For literature teaching specifically, AI literacy needs a further, discipline-specific layer: understanding what a language model can and cannot do with figurative language, ambiguity, unreliable narration, and historical or cultural context, since these are the very features that define literary reading.

Epistemic authority and the humanities

A second and more theoretically oriented strand of scholarship asks a harder question: what happens when students begin treating AI-generated readings of a text as authoritative, rather than as one interpretation among many to be tested against the text itself. Every learning tool carries an implicit pedagogy, a set of assumptions about what counts as knowledge and whose interpretations are trustworthy. Generative AI systems, despite their apparent neutrality, are no exception; their

output reflects patterns in training data that can reproduce cultural and linguistic bias, and research has found that AI-generated content can subtly reinforce dominant framings that students may reproduce in their own writing without recognizing the source. For a discipline built on interrogating whose voice is heard and whose is marginalized in a text or a canon, this is not a peripheral concern but a central one.

Humanities scholars have also noted that literature departments are, paradoxically, well equipped to respond to this problem, because interrogating the authority and situatedness of any text or claim is already the discipline's core method. Several scholars go further, arguing that the interpretive skills literature teaches, distinguishing description from analysis, reading against the grain of a text, attending to what is left unsaid, are exactly the skills students now need to evaluate AI output responsibly. On this view, GenAI does not threaten the value of literary study so much as make the case for it more urgent.

New Pedagogies

Across the studies and institutional guidelines surveyed for this paper, five pedagogical approaches recur often enough to be treated as an emerging consensus, even though no single institution has adopted all five in a coherent program.

AI-assisted close reading

Rather than banning AI outright, some instructors now use it as a deliberate foil for close reading instruction. Students are given a passage and asked to generate an AI reading of it, then to annotate that output against the primary text, identifying moments where the AI has described rather than analyzed, misread a tone or an allusion, or missed a formal feature such as meter, enjambment, or free indirect discourse. This inverts the traditional model of close reading instruction, where the instructor models expert reading and the student imitates it; here, the student is positioned as the more careful reader from the outset, evaluating a fluent but shallow interlocutor. Early classroom reports suggest this framing increases student engagement with formal literary features, since spotting what an AI system gets wrong requires attending closely to exactly the details close reading is meant to teach.

Prompt literacy as critical literacy

A second approach treats the crafting of prompts to a language model as an extension of critical literacy rather than a separate technical skill. Asking a model to "analyze the imagery in this poem" typically produces generic, low-value output; asking it to consider a specific critical lens, a particular formal feature, or a named counterargument produces something worth evaluating. Teaching students to write increasingly precise, theoretically informed prompts, and then to critique the

resulting output, functions as an indirect way of teaching them what a good literary argument actually requires: specificity, a defensible claim, and attention to textual evidence. Framed this way, prompt writing becomes a diagnostic exercise in argumentative literacy rather than a shortcut around it.

Adversarial and comparative reading

Building on the first two approaches, some courses now stage explicit human-versus machine comparisons as a recurring classroom activity: students produce their own analysis of a passage independently, then compare it against an AI-generated analysis of the same passage, in class, in groups. This comparative structure surfaces disciplinary values that are often left implicit, such as the expectation that a literary argument should notice ambiguity rather than resolve it prematurely, or that historical and biographical context changes what a passage means. It also gives instructors a natural occasion to discuss where the AI's reading was, in fact, useful, which helps keep the exercise from becoming a simple demonstration that AI is inferior, and instead builds a more accurate, mixed picture of what the tool is and is not good for.

Co-creative and generative writing pedagogy

In creative writing components of literature courses, GenAI is increasingly used not to produce finished work but to generate raw material for students to revise, reject, or redirect: alternative plot continuations, character voice experiments, or stylistic pastiche that a student then critiques and rewrites in their own voice. This approach treats the AI output as a draft to argue with rather than a draft to submit, and asks students to articulate, in their revisions, exactly what the AI version got wrong about voice, pacing, or characterization. Some instructors extend this into explicit imitation exercises, having students generate AI text in the style of a studied author and then identify what the model reproduced successfully (surface diction, syntax) against what it failed to capture (structural ambition, thematic coherence), turning stylistic analysis into an active, comparative exercise rather than a passive one.

Accessibility and personalized scaffolding

A less contested but still significant use of GenAI in literature classrooms is as an accessibility and scaffolding tool: generating glossaries for archaic or dialect vocabulary, producing simplified paraphrases of dense syntax for students who are building reading fluency or working in a second language, or offering background historical and cultural context on demand. Used this way, GenAI functions less as an interpretive shortcut and more as an extension of resources literature teachers have

always provided, such as annotated editions or study guides, but delivered adaptively and on demand. The pedagogical risk here is different from the risk of interpretive outsourcing: it lies in students substituting an AI paraphrase for engagement with the primary text's actual language, which is frequently where a literary work's meaning is most concentrated.

New Challenges

The pedagogical opportunities above do not resolve the difficulties GenAI introduces; if anything, each new use case introduces a corresponding new risk. Five challenges are especially persistent in the current literature and in departmental policy debates.

Academic integrity and assessment validity

The most visible concern is the oldest one: that students will submit AI-generated writing as their own. This is a genuine problem, but the research literature complicates the simple detection-and-punishment framing that dominated early institutional responses.

Detection tools remain unreliable, particularly against lightly edited AI text, and their error rates raise due process concerns of their own. More fundamentally, a well-designed controlled study comparing AI-generated and first-year undergraduate essays on Old English poetry at a research university found that experienced human markers could distinguish AI from student writing correctly only around four times out of five, and that the AI essays, while more formally structured, scored close to human essays overall, differing mainly by describing rather than analyzing textual features and occasionally missing key details in the passage. This suggests that surface-level detection is an unstable foundation for policy, and that the more durable response is to redesign assessment so that AI-generated text is a poor substitute for what is actually being tested, for instance by assessing in-class annotation, oral defense of a written argument, or drafts that document a student's evolving thinking, rather than only a finished essay.

Erosion of interpretive skill and critical thinking

A second, harder to measure concern is that habitual reliance on AI-generated interpretation may weaken students' own capacity for sustained, independent reading before that capacity has fully developed. Literary interpretation is a skill built through repeated struggle with ambiguity, and if students routinely outsource the first pass at that struggle to a model, they may never build the tolerance for productive confusion that literary study depends on. This concern is distinct from academic integrity, since it can occur even when no rule is broken and no assignment is technically outsourced,

for instance when a student reads an AI summary before reading the primary text and consequently never encounters the text on its own terms. Evidence on this question remains preliminary, but the concern has enough theoretical grounding in reading pedagogy that most of the frameworks discussed in this paper prioritize scaffolded, monitored AI use over unrestricted access, particularly for younger or less experienced readers.

Bias, cultural homogenization, and epistemic authority

Because language models are trained on large, unevenly representative bodies of existing text, their output tends to reproduce dominant interpretive frames, and research has documented cases in which AI-generated academic content reinforces gender and cultural bias that students then reproduce, sometimes without recognizing it as bias rather than neutral analysis. In literature classrooms, where a significant part of the discipline's recent development has involved recovering and legitimizing marginalized voices and unconventional readings, an over-reliance on AI-generated interpretation risks quietly re-centering more conventional, canonical readings that happen to dominate the training data. This is a subtler risk than outright misinformation, since the AI's output is often plausible and well-written, which makes it harder for students to recognize as one interpretation among many rather than as a neutral synthesis of scholarly consensus.

Unequal access and the digital divide

As GenAI tools become embedded in course design, disparities in access to the most capable models, whether due to cost, institutional licensing, device access, or reliable internet connectivity, risk becoming a new axis of educational inequality layered on top of existing ones. Students without home access to premium AI tools, or attending underresourced schools that restrict AI use for infrastructure or policy reasons, may enter courses that assume a baseline of AI fluency they have not had the chance to develop, while wealthier or better-connected peers arrive already comfortable with the tools. Any departmental policy that leans heavily on AI-integrated pedagogy needs an explicit plan for equitable access, or it risks widening exactly the gaps that literature education, with its long tradition of expanding access to cultural capital, has historically tried to close.

The practical difficulty of assessment redesign

Finally, and less discussed in the theoretical literature but acute in practice, is the sheer workload of redesigning assessment to be resilient to GenAI. Oral assessment, in-class writing, portfolio-based grading, and iterative drafting all produce better evidence of a student's own thinking than a single take-home essay,

but each is more time-intensive to administer and grade than the assignments they replace, particularly in large lecture courses common in introductory literature surveys. Departments adopting these approaches report significant front-loaded costs in redesigning rubrics and retraining teaching assistants, and institutions vary widely in whether they provide the release time or funding to support that redesign, meaning that the burden of adaptation often falls unevenly on individual instructors rather than being absorbed at the institutional level.

Evidence from Practice

Much of the debate around GenAI and literature teaching remains speculative, but a smaller set of controlled, classroom-grounded studies offers firmer footing.

The most directly relevant is a study conducted with first-year English Language and Literature undergraduates at the University of Oxford, examined on a module covering early medieval literature. Researchers had AI systems generate close-reading stylistic commentaries on passages of Old English poetry and compared them, blind, against genuine student commentaries, using experienced human markers under the standard UK grading system. The AI-written essays averaged a score of 60.46 against 63.57 for human essays, a difference that was not statistically significant. Markers correctly identified the origin of an essay in roughly four out of five cases, distinguishing human work chiefly by its sustained, close attention to poetic form and its use of cultural and critical context, and, somewhat ironically, by surface markers such as spelling slips and looser structure. AI essays were more formally polished but tended to describe rather than analyze, and sometimes missed key features of the passage entirely. The study also tested three AI-detection tools and found substantial variation in their accuracy, underscoring that automated detection is not yet a reliable basis for integrity enforcement in literary assessment specifically.

Broader systematic reviews of GenAI in higher education corroborate several of this study's implications at scale. Reviews of the literature on AI literacy report that professional development access and community-based peer support, more than formal training alone, predict confident and appropriate instructor use of AI tools, and that emotional barriers such as anxiety and discomfort remain a significant, under-addressed obstacle to adoption even among willing instructors. Reviews focused specifically on K-12 contexts find that the large majority of published studies call for sustained teacher professional development as a precondition for effective classroom integration, rather than treating AI literacy as

something students alone need to acquire. Across both K-12 and higher education, researchers converge on the point that GenAI systems, unless deliberately prompted with context, have no knowledge of a specific student's background, prior learning, or classroom culture, which keeps the teacher, rather than the tool, as the necessary interpretive authority in the room.

Taken together, this evidence supports a cautious middle position: GenAI-generated literary analysis is good enough to complicate simple detection-based integrity enforcement, but it remains distinguishable from skilled human interpretation on exactly the dimensions, sustained analysis, contextual grounding, and attention to ambiguity, that literature instruction aims to cultivate. This gap, rather than closing it entirely, is where much of the emerging pedagogy described in Section 3 does its work.

Toward a Framework for Practice

Drawing on the pedagogies and challenges surveyed above, this section proposes six practical commitments for literature departments and individual instructors, organized from policy-level to classroom-level decisions.

1.Publish a discipline-specific AI use policy, not a generic institutional one. Blanket university-wide AI policies rarely address what is distinctive about literary analysis, namely that some AI use (background research, accessibility scaffolding) is pedagogically benign while other use (outsourcing interpretation) undermines the assignment's purpose. Literature departments benefit from articulating, assignment by assignment, where on that spectrum a given task sits, rather than relying on a single yes-or-no rule.

3.Redesign assessment around process, not only product. Given the evidence that finished AI-generated essays can be difficult to distinguish reliably from student work, assessment should increasingly capture the process of interpretation: in-class annotation, drafts with visible revision, reading journals, and oral discussion of a written argument. These formats are more resistant to AI substitution and, as a secondary benefit, give instructors better insight into a student's actual reading development.

3.Teach AI evaluation as a literary skill, not a separate technology unit. Rather than a one-off AI literacy workshop, the skills needed to evaluate AI-generated literary analysis, distinguishing description from argument, checking claims against textual evidence, noticing what has been flattened or omitted, are close cousins of the skills already taught in close reading instruction. Embedding AI evaluation directly

into existing close reading and rhetoric units makes it durable rather than a novelty exercise.

4. Build in explicit attention to bias and canon. Given the documented tendency of AI output to reproduce dominant interpretive frames, instructors should treat AI-generated readings as an occasion to discuss whose interpretations are represented and whose are not, connecting this directly to existing coursework on canon formation, marginalized voices, and critical theory rather than treating it as a separate technology-ethics topic.

5. Address access inequity explicitly, at the department level. Departments should not assume uniform student access to premium AI tools and should either provide access through institutional licensing or design assignments that do not depend on tools some students cannot reach, so that AI-integrated pedagogy does not become a new source of disadvantage.

6. Invest in sustained, peer-supported instructor development. The strongest single predictor of confident, pedagogically sound AI integration in the literature reviewed here is not formal training alone but ongoing peer community and hands-on practice with the tools. Departments adopting AI-integrated pedagogy should budget for recurring, informal peer exchange, not only a single onboarding session, and should treat instructor confidence with these tools as a resourcing question rather than an individual failing.

Conclusion

Generative AI has not made English literature teaching obsolete, nor has it left the discipline untouched. What the evidence assembled here suggests instead is a discipline in the process of clarifying its own purpose under pressure. When an AI system can produce fluent, well-structured, plausible commentary on a poem in seconds, the value of literary education can no longer rest on the production of commentary alone; it must rest on the quality of attention, the interpretive judgment, and the willingness to sit with ambiguity that produce commentary worth reading. The controlled evidence reviewed in this paper suggests that human interpretation, at its best, is still distinguishable from machine output on exactly these grounds, but that gap is not guaranteed to persist without deliberate pedagogical effort.

The five pedagogies surveyed here, AI-assisted close reading, prompt literacy, adversarial comparison, co-creative writing, and accessibility scaffolding, share a common logic: they use the AI system as an object of critical scrutiny rather than as a silent collaborator or a banned technology, treating the encounter with machine-generated interpretation as itself a literary and critical exercise. The five challenges,

integrity, skill erosion, bias, access, and assessment workload, are not fully solvable, and any department claiming to have solved them should be regarded with some skepticism. They are, however, manageable through the kind of concrete, discipline-specific policy and assessment redesign proposed in Section 6.

The deeper claim of this paper is that literature departments are unusually well positioned, among academic disciplines, to lead this adaptation rather than simply endure it. A field built on interrogating whose voice is authoritative in a text, on distinguishing description from analysis, and on tolerating ambiguity rather than resolving it prematurely already possesses the interpretive habits that responsible AI use requires. The task facing literature educators is not to invent these habits from scratch, but to make them explicit, and to teach them as deliberately in relation to a language model's output as they have long been taught in relation to a printed page.

Works Cited

- Annapureddy, R., et al. "Generative AI Literacy: Twelve Defining Competencies." *Digital Government: Research and Practice*, vol. 6, 2025. <https://doi.org/10.1145/3685680>.
- "ChatGPT versus Human Essayists: An Exploration of the Impact of Artificial Intelligence for Authorship and Academic Integrity in the Humanities." *International Journal for Educational Integrity*, 2024, <https://doi.org/10.1007/s40979-024-00161-8>.
- Cotton, D. R. E., et al. "Chatting and Cheating: Ensuring Academic Integrity in the Era of ChatGPT." *Innovations in Education and Teaching International*, vol. 61, no. 2, 2024, pp. 228–239.
- "Epistemic Authority and Generative AI in Learning Spaces: Rethinking Knowledge in the Algorithmic Age." *Frontiers in Education*, 2025, <https://doi.org/10.3389/feduc.2025.1647687>.
- Fleischmann, C., et al. "Fostering GenAI Literacy in Higher Education for Future Workplace Preparedness." *Proceedings of ECIS 2024*, 2024.
- "Generative AI Use in K-12 Education: A Systematic Review." *Frontiers in Education*, 2025, <https://doi.org/10.3389/feduc.2025.1647573>.
- Kalantzis, M. "Literacy in the Time of Artificial Intelligence." *Reading Research Quarterly*, 2025, <https://doi.org/10.1002/rrq.591>.

- Michel-Villarreal, R., et al. "Challenges and Opportunities of Generative AI for Higher Education as Explained by ChatGPT." *Education Sciences*, vol. 13, no. 9, 2023, article 856.
- Mishra, P., et al. "TPACK in the Age of ChatGPT and Generative AI." *Journal of Digital Learning in Teacher Education*, vol. 39, no. 4, 2023, pp. 235–251.
- Moorhouse, B. L., et al. *Generative AI Tools and Assessment: Guidelines of the World's Top-Ranking Universities*. Working paper, 2023.