Volume 35, 2025
http://journals.sfu.ca/dwr

Article

Where We're At, What We Must Know, and Where We Can Go: A Systematic Review of Research about Writing and Artificial Intelligence

Christopher Eaton
University of Toronto Mississauga

Isabella Belmonte
University of Toronto Mississauga

Talla Enaya
University of Toronto Mississauga

Sarah Flood
University of Toronto Mississauga

Zainab Khalil
University of Toronto Mississauga

Anthony Makwanda
University of Toronto Mississauga

Mian Muhammad Ahmed Shah
University of Toronto Mississauga

Alexia Toma
University of Toronto Mississauga

Tiffany Wang
University of Toronto Mississauga

Connor Yu
University of Toronto Mississauga

Volume 35, 2025
http://journals.sfu.ca/dwr

Abstract

This paper reviews AI discussions to develop a thematic map of scholarship on AI, writing, and writing pedagogies. The project, developed between a university instructor and nine undergraduate students, had two central aims. The first was to synthesize the early research conversations on AI in writing contexts. The second was to identify gaps that will point to important next steps as the scholarly record on AI and writing develops. The paper presents five prominent themes: AI literacy, evaluating AI outputs, rhetoric in Human-AI interactions, AI and bias, and academic integrity. Conversations on AI literacy and academic integrity represent conceptual level discussions around AI and writing. The conversations around evaluation, rhetoric, and bias align more with AI writing practices and how these practices affect teaching and learning. Together, they provide a useful snapshot of scholarship and inform future work in a rapidly developing research conversation.

Keywords: artificial intelligence; generative AI; writing and AI; writing pedagogies; systematic review

The sudden influence of generative artificial intelligence (AI), particularly large language models, has caused a flurry of research and scholarly conversation about AI's impact on writing, literacies, and pedagogies. The scholarly record has grown quickly as writing practitioners confront impacts of AI in their learning environments. The speed of both the technology and early scholarship has created an uneven understanding of AI in writing contexts. This rapid growth also highlights how practitioners have reacted in real time, laying a theoretical and empirical foundation that did not exist three years ago.

Given the rapid nature of the scholarly record's growth, it is prudent to develop a thematic map of what is known and what is not known about AI in writing practices. This map would help develop writing curricula and scholarship that respond to AI's influence. Our project set out to develop this map through a literature review that explored extant scholarship on writing, pedagogy, and literate practices in AI-mediated writing environments. Our collaboration of nine learners and a course instructor formed after a course at a Canadian University. Our project goal was two-fold: to synthesize existing research along the most prominent themes, and, in turn, analyze these themes to determine what future research would best expand the understanding of AI's influence.

Two questions formed the foundation for our discussions:

1. What are the most prominent gaps in current scholarly conversations on AI and writing?
2. How might these conversations inform next steps for writing researchers and practitioners?

The first question prioritized what was missing in the scholarly record. Focusing on gaps allowed us to narrow our lens such that we could consider suggestions that moved beyond the obvious. For example, since we know there is a lack of empirical evidence, it is important to consider what kinds of empirical evidence might make meaningful contributions to addressing those gaps. This is where the second question came in: identifying gaps offered a starting point, but pointing to next steps was integral to laying out the map for others to follow.

In doing so, we developed a snapshot of the current scholarly landscape while using that landscape to inform future research. This paper is the result. We begin with an overview of our project's origin and then describe our process for developing the systematic review. We then synthesize the literature around our major themes before ending with a discussion of future research steps that would add a layer to each theme discussed.

Project Origin

This project began in an upper-year undergraduate course at a large Canadian university. The course investigated generative AI's implications for writing and knowledge making. The course was capped at 20 students, making for a hands-on, small group, active learning experience. One major assignment asked learners to develop a literature review using scholarship on the intersection between AI and writing, rhetoric, and literacy. The assignment had three objectives: 1) help learners trace early patterns in research on AI in writing studies; 2) encourage learners to make connections between these conversations; 3) prompt learners to map gaps in the conversations and suggest next steps. Together, the objectives helped learners establish a knowledge foundation that could inform future class interactions, activities, and assignments. They also positioned learners to make their own contribution to scholarly conversations about AI.

Chris (the course instructor) thought that the reviews developed for class could contribute to scholarly literature. After the course, he invited learners who wanted to extend what they learned from that assignment to a new context. Nine learners agreed to participate, and this new collaboration occurred between April and July 2025. Participants came from diverse disciplinary backgrounds, such as economics, sociology, health sciences, and information technology. Having student participants added a useful perspective to the review because they offered a perspective that contemplated the articles in relation to their own experiences with AI. The gaps we identified and the suggestions we made could therefore account for both sides of the pedagogical relationship.

Methodology

Methodologically, our review occupied a hybrid position. The breadth of our research questions aligned with a scoping review of the literature about a field or discipline (e.g., Hassel, 2013). The nascent nature of AI scholarship prevented some important dimensions of a more advanced systematic review, such as meta-analyses or methodological reviews. However, we also realized that our collaboration asked more specific sub-questions and involved an implicit degree of assessing the quality of study/theory. Our shared history grappling with early AI conversations for class positioned us to view the literature through specific lenses, with some pressure points standing out naturally from others based on our shared history. For example, as much as we were interested in finding gaps in extant literature, our prior interactions had primed us to consider the role of assessing output, engaging rhetorically with outputs, and the problematic layers of bias in outputs. We found ourselves

asking: What is not known pedagogically about how learners engage with AI outputs? What may be overlooked in current conversations around bias and AI? What is known about the rhetorical implications of AI, and how does that affect learning? While we still searched widely in the literature, the types of gaps we contemplated drew us to specific types of conversations. Our approach, in this sense, aligned more with the representative sample that is common in a systematic review.

A systematic review synthesizes extant research in a field. It aims to pull together a "representative sample of the evidence" to influence "policy, practice and future research" (Torgerson, 2007, p. 288). As Lee and Kwon (2024) noted, synthesis leads to interpreting results that can then be disseminated, which places interpretation at the heart of the approach. Davis et al. (2012) raised questions about this approach to reviewing literature, noting that it can be reductionist and limited in its findings. However, Scott et al. (2018) emphasized the systematic review's ability to give relevance to a diversity of methodologies as part of the representative sample selected. In our situation, the systematic review allowed us to capture the array of approaches that have been used to grapple with generative AI's sudden, prominent, and growing influence on writing. We did not have to limit our scope to a specific methodological approach. The systematic review's flexibility also allowed us to account for the range of theoretical frameworks that inform research on generative AI, writing, and literacies.

Our review involved three main steps: coding our class literature reviews, selecting sources for the review, coding these sources, and developing a qualitative synthesis. Some steps, such as coding and synthesizing, began as individual practices for each team member. We then gathered to discuss coding results, compare findings, and establish themes and next steps. Our coding process included several tiers outlined by Adu (2019), moving from open coding, to developing categories, and then developing themes through constant comparison. More details on the steps and coding processes are provided in the subsections that follow.

As the instructor and project lead, Chris participated in every stage of the review and reviewed all scholarship. This added a layer of validity to each part of the coding process. Chris's review could be compared to the reviews from individual groups. Then, the reviews could be discussed in the group meetings to consolidate ideas and refine themes. The separate layers of review spurred collegial dialogue where no one's perspective was privileged. Rather, we worked together to develop ideas and make connections between conversation clusters.

Coding Timeline and Student Training

Students received training in coding and analyzing throughout the project. Our course had a 1.5-hour lesson in how to write a literature review. This lesson occurred when the assignment was introduced. It focused on organizing sources thematically, looking for overlapping conversations, and identifying and articulating gaps in that conversation. This lesson was complemented by learners working with sources during some of our evaluation activities related to AI outputs. Sample literature review publications were included on our university's LMS (e.g., Graham et al., 2018; Hassel, 2013; Scott et al., 2018) to provide learners with different models for organizing their conversations that could facilitate our training.

Our first training session was late April 2025. The session was dedicated to coding the anonymized class literature reviews to establish early themes. A shared folder was also created with additional resources, including a breakdown of coding processes, training materials, and information on scoping and systematic reviews. This folder supported the team as we split into smaller groups to complete the individual reviews.

The second training session occurred in May 2025 when we met to discuss the findings from the class literature reviews. We spent the first part of this meeting discussing codes and categories. Students then learned how to conduct a review and qualitative synthesis. They learned about developing codes around a publication's introduction, thesis, research question, discussion, and conclusions. The training then focused on how to move from codes and categories into themes, as well as how to cross reference findings with the earlier findings from the literature reviews.

A more detailed breakdown of our coding steps is provided in the following sections.

Step 1: Coding Literature Reviews

The first step in our process was to code the literature reviews conducted in class. The literature reviews were anonymized by Chris. Anonymizing the review placed emphasis on the content of the papers rather than who developed the reviews. This was important because our small class meant that learners developed varying degrees of relationships to each other, and these relationships had the potential to influence the coding. Only reviews from project participants (not everyone in the class) were included in the review (9 total).

Our review drew upon Adu's (2019) description of open coding. Adu modelled how to develop codes and how to categorize those codes for a more focused understanding of the relationships

between codes. We differed from Adu insofar as this review did not engage in iterative theoretical sampling to develop themes; instead, these early codes and categories offered a useful way to compare results from the systematic review portion of the project (see Step 3). After the initial training session, we met as a group to discuss our findings. We collaborated on a shared document to consolidate our individual categories and codes. Our discussion then zoomed out to consider early themes that might inform future analyses and/or that would be useful to keep in mind as we compare the literature surveyed in the previous fall to a larger and (at times) more recent corpus of publications.

Coding the literature reviews had one final benefit. Our discussions of categories and early themes helped us organize clusters of conversations. These discussions provided useful keywords that facilitated library searches (see Step 2) and established a general structure for organizing conversations. We anticipated revising this structure at a later coding meeting as we learned more about the conversations. However, this early review offered a structure to organize our work as a group and divide different parts of the conversation, which led to a more focused review. These categories included: writing skills/writing development; feedback/revision; genre and rhetoric with AI; critical thinking and AI; evaluative judgement and black boxes; academic integrity and ethics in writing classrooms; writing education and pedagogy. We realize that there are other ways to organize categories (see Risko, 2008; Scott et al., 2018); this organization emerged organically from our work together and made the most sense for our analysis.

Step 2: Selecting Sources for Review

The second stage involved establishing a list of sources to review and include in our representative sample of scholarship. What constitutes a representative sample varies by study. For example, Lee and Kwon (2024) reviewed 25 papers, Scott et al. (2018) reviewed 29 papers, and Risko et al. (2008) examined 82 publications. Our selection process was guided by our inclusion and exclusion criteria. These criteria were established with the aim of capturing the diversity and multiplicity of theories and methodologies that are used in writing and literacy research. This broad goal reflects processes undertaken by systematic reviews in education (e.g., Risko et al., 2008; Scott et al. 2018). From here, we established inclusion and exclusion criteria, which were adapted from Scott et al. (2018) for a new and emergent context. The inclusion criteria included:

- Studies related to writing and reading with AI
- Studies about AI pedagogy, with priority to pedagogical studies in literacy

- Published within the last 10 years
- Published in English
- Published in a peer reviewed publishing

The exclusion criteria included:

- Studies about AI older than 10 years—due to the recency and technological relevance of large language models
- Studies that discuss artificial intelligence without connection to reading, writing, literacy, or the pedagogies related to these areas
- Publications that discuss academic integrity and artificial intelligence without a clear connection to writing, reading, literacy, or pedagogies related to these areas

These criteria directed the searches and informed how the literature was organized into conversation clusters.

Chris began by organizing the sources listed in the class literature reviews into the categories listed in the previous section. He then consulted four special issues on AI, writing, and literacy from the following journals: *Composition Studies* (2023); *Double Helix: A Journal of Critical Thinking and Writing* (2023); *Computers and Composition* (2024); and *Reading and Research Quarterly* (2024). The next step was to conduct a library search using the key words listed in Table 1. The broad categories listed emerged from the course focus and aligned with the primary lenses that learners in the course used to complete their assignments. This way of establishing search terms is consistent with the search approach used by Davis et al. (2012), who focused on large areas and divided the searches across many keywords. The keywords listed in Table 1 derived from a combination of the publication keywords in the references added from the class reviews and from the codes/categories/early themes from step 1. The keywords are all cross-referenced with “artificial intelligence”, “AI”, “generative artificial intelligence”, “chatbot”, and “large language model” as alternate keywords.

Table 1. A list of search terms used to find scholarship on generative AI

Broad Category	Keywords for search
Literacy	Academic literacies, reading, writing, disciplinary literacy, language arts, literacy instruction, reading/writing connection
Writing	Composition, composing, writing skills, writing in the disciplines, writing across the curriculum, writing instruction/pedagogy, writing to learn, rhetoric, genre
Reading	Reading skills, source engagement, content-area literacy, literature review, science of reading, reading to learn

The completed list was uploaded to the shared folder, and an email was sent to four individual teams of student collaborators (2-3 people per team). Each team was required to review one or two categories or conversations, with each team reviewing approximately 15-20 publications. In total, 76 sources were included in the review.

The sources selected were not screened for methodological quality. This is a common procedure in systematic literature reviews (Scott et al.,2018; Torgerson, 2007). The nascent nature of research on AI and writing made a methodological review impractical because there is insufficient empirical research to select sources along methodological lines. Instead, we selected papers that best fit our inclusion and exclusion criteria and that best fit the conversation clusters that emerged in our early discussions. This is also why our representative sample tends to be consistent with a study like Risko et al. (2008) rather than the more selective syntheses of other systematic reviews. We aimed to generate a snapshot of emergent conversations around AI, and doing so required being open to exploring the early conversations, theoretical and conceptual publications, and the early empirical research available in the scholarly record as of spring 2025.

Step 3: Qualitative Synthesis

The final step was to provide a qualitative synthesis for each study explored. Student teams started with reviewing the literature independently. They then met in small groups (2-3 people) and consolidated their findings within the categories they explored. These results were then uploaded to the shared folder. Concurrently, Chris reviewed and developed a qualitative synthesis for each of the 76 publications. These findings then informed the conversation we had when we discussed results in mid-June.

This June meeting was informed in large part by the synthesis created by each group. Following the methods of qualitative synthesis outlined by Scott et al. (2018), the syntheses sought to consolidate the research purpose, questions, and findings from each publication. We differed from Scott et al. (year) because we did not develop themes from these syntheses alone. Each group followed the process of constant comparison detailed by Adu (2019), where they developed codes and categories from the syntheses. This made it easier to make comparisons between categories.

Results and Discussion

In the paragraphs that follow, we discuss core themes that emerged from our investigation. We begin each subsection with an overview of key dimensions of the conversations. We then use the overview to offer a discussion of what next steps might add to the conversations that are already in place. Occasionally, we use this discussion to challenge or prompt more critical thought around the conversations.

AI Literacy Frameworks

A common discussion point is the need to develop AI literacy. This need applies to both sides of the pedagogical relationship—students and teachers must develop AI literacy (e.g., Cardon et al., 2023; Lin et al., 2025). Many models exist for how this can occur. Su and Yang (2023) introduced the IDEE framework, which prioritizes four features: identifying the desired outcome, determining the appropriate level of automation, considering ethics, and evaluating effectiveness. Lo (2023) created the CLEAR framework for effective AI prompting, essentially breaking AI requests into five core principles: concision, logic, explicitness, adaptability, and reflexivity. More specific to learners, Lin et al. (2025) proposed the IDEA framework for multicultural second-language classrooms. The model has four components: interpret, design, evaluate, and articulate. These components model how AI might be integrated into digital multimodal compositions. Beck and Levine (2024) developed a framework that centered human imagination and agency to support writing development. Finally, Burriss and Leander (2024) drew attention to posthuman considerations of AI pedagogies, wherein human-machine knowledge making practices are constantly influenced by machine/human inter and intra-actions. Intra-actions highlight not just how different AI practices may interact but also the effects of interactions on literate activities.

There are also AI literacy frameworks more directed at teachers. Warschauer et al. (2023) and Tseng and Warschauer (2023) developed a pedagogical framework consisting of five elements: understand, access, prompt, corroborate, and incorporate. Ng et al. (2021) offered an early look into what AI literacy may look like by drawing upon the TPACK framework built around three types of knowledge: technological, pedagogical, and content. These dimensions connect to learning artefacts, pedagogical approaches, and subject matter and have proven useful to help teachers integrate technologies into their pedagogies. Moreover, Burriss and Leander (2024) encouraged educators (particularly those in writing and literacy) to adapt their posthuman AI framework so that they could use it proactively in environments influenced by AI. Moreover, Tzirides et al.'s (2024) case study showed how students may develop their AI literacy skills using a curriculum the authors developed. One further framework was developed by Cummings et al. (2024) who proposed "DEER" for developing teaching approaches that respond to generative AI's influence: clearly define project stages, evaluate specific AI tools that pair with the learning opportunity, encourage student exploration of and with AI, and build in space for learners to reflect.

These frameworks share a few common threads. They emphasize a need to understand the intersection between humans and the AI tools they use. The intersection, however, was emphasized more for teachers and pedagogical design, most prominently in Ng et al. (2021) and Cummings et al. (2024). However, Su and Yang (2023) and Burriss and Leander (2024) addressed these connections for both learners and teachers. An important next step will be to understand more about these intersections for learners and writing. These conversations exist in the scholarly conversation (see the Human-AI Rhetorical Collaborations section below), but they are not as prominent in AI literacy frameworks.

Evaluation gets more attention with student AI use. Evaluating AI output features in frameworks by Lo (2023), Su and Yang (2023), and Warschauer et al. (2023). It is also applied in Beck and Levine's (2024) discussion, particularly as a means for writers to connect output to wider community expectations and influences. Associated with evaluation is a common emphasis on reflexivity. It is explicit in Cummings et al. (2024) and Lo (2023), and more implicit in Burriss and Leander (2024). Student reflexivity is thought to provide learners with an advantage in meeting the evaluation requirements that are built into many frameworks. More on evaluation is available in the next section.

However, the diversity of AI literacy frameworks available brings into question the value of the term AI literacy in general. Johnson (2023) highlighted the risk of writing instructors who feel the need to develop entirely new knowledge foundations to respond to AI's influence. Johnson noted the

importance of grounding pedagogical and scholarly practices on existing research. The eclectic mix of AI literacy frameworks means that some frameworks build on prior knowledge better than others. For example, Burriss and Leander (2024) build from posthuman knowledge foundations. Similarly, Girdharry and Khachatryan (2023) used Meaningful Writing and de Matas (2023) drew upon Writing about Writing to inform their scholarship. Conversely, Lin et al. (2025) and Warschauer et al. (2023) depended less on prior knowledge to develop their frameworks. The variety of frameworks points to a challenge in determining what would be best suited for a given writing situation or classroom. Whether any frameworks are substantiated as useful tools for writing and pedagogy remains an open question. If AI literacy becomes a catch-all term for any activity related to AI use, then it—along with the frameworks associated with it—could lose utility and relevance.

It may be helpful to consider "literacy" as too narrow of a term. It is increasingly common to think of literacies, which accounts for the multiplicity of literate practices that can be enmeshed in the uptake of a technology. A more pluralistic view would align with Burriss and Leander's (2024) and Wang and Wang's (2025) discussions of the posthuman and AI, where human and AI intra-actions demand a variety of literate practices and competencies. These practices would not account for literacy in any singular sense; rather a framework on AI literacies would account for the diversity of practices and insights that may inform such a framework.

If an AI literacy framework is a desired outcome for scholarly and pedagogical practice, then it will be necessary to both acknowledge the diversity of literate practices and consolidate existing frameworks. It would be useful to develop a synthesis of extant AI literacy frameworks to understand common patterns in the literature. Our three common threads (human-AI connections, evaluation, and reflexivity) may offer a starting point for synthesis, but they are only the start of a bigger conversation. A deeper synthesis would allow scholars to identify gaps in and across frameworks. It would also offer practitioners a foundation for translating theoretical frameworks into pedagogical practice.

Evaluating Output

Evaluating AI outputs is central to many discussions around AI use as a knowledge making tool. Mollick and Mollick (2023), Su and Yang (2023), and Markauskaite et al. (2022) raised the need for evaluating AI outputs to effectively use them in knowledge creation. Evaluating the validity of AI outputs was central to early case studies about AI use in composition practices (e.g., Fyfe, 2022). Evaluation has also been framed as central to writing processes with AI. Pigg (2024) and Graham

(2023) highlighted the iterative, evaluation-heavy nature of using AI to support writing. Graham (2023) provided a model for how writing processes may evolve given AI's influence, with the major shift occurring at the level of evaluation to ensure validity of AI output. Within this model, evaluation is a catalyst for future iterations of prompting and idea generation.

Evaluative judgement is an important concept within this discussion. Bearman et al. (2024) highlighted that good evaluative judgements require a "contextualised understanding of quality" (p. 7) to assess the quality of an AI output relative to the disciplinary context. As Tai et al. (2018) emphasized, this practice requires an understanding of what quality is in the first place. As such, Bearman et al. (2024) explained that 1) learners do not always have the skills to make good judgments and 2) learners could (paradoxically) use AI to facilitate strong judgements if they have the appropriate training. Insights from Anson (2024), while not directly related to evaluative judgment, complicated these ideas. Anson (2024) drew attention to the need for learners to recognize what language choices used are appropriate for disciplinary contexts. So, it is not simply the evaluation that is necessary; making strong evaluations depends on an understanding of epistemology and power within a discipline both with and without AI. This is something that learners need time to develop.

The ability to evaluate validity is obscured, however, by the nature of AI itself. This is what is called the "black box." Scholarship has identified the black box because of the opacity of AI systems (Bearman & Ajjawi, 2023; Zednik, 2021). Moulaison-Sandy (2025) defined the black box as a model produces a "situation where the internal logic is obscured" (p. 9). While the lack of transparency is problematic, it is a fundamental function of generative AI because it allows users to focus on the outputs. Zednik (2021) proposed a framework that accounts for the realities of the black box, and Bearman and Ajjawi (2023) detailed how pedagogies could help learners work within the confines of the black box.

One argument, however, brings into question how plausible it is to work within the black box. Flenady and Sparrow (2025) argued that outputs can never be fully verified, pointing to the reality that AI outputs do not have an agency or imperative to report the truth about the world. Focusing solely on outputs omits many important dimensions that contribute to evaluating AI's validity which leads to a "perverse outcome of making individual students entirely responsible for the truth of the claims in the work they submit" (p. 5). These truths cannot be entirely validated, leaving learners responsible for constructions that they cannot trace.

Some of the insights, however, are at odds with realities in writing classrooms. There tends to be a culture of resistance to current AI practices, especially in the insights of scholarship from Flenady and Sparrow (2025), and—to a lesser degree—Anson (2024). While caution is important, it is equally important not to back instructors into a corner where AI is antithetical to learners being able to make strong evaluations of knowledge within their discipline. Notions that learning must occur without AI first so that learners develop a foundation in important knowledge-making practices are valid, albeit idealistic. They tend to be at odds with the practices learners employ in writing classrooms, where AI is often being used without important knowledge foundations and where learners are making evaluations without these foundations.

How can pedagogies be designed to address evaluating AI outputs? This area is where data will be important going forward. If making strong evaluations is central to strong AI use, then understanding how evaluations work in different contexts will be useful. This will help teachers and scholars understand where and how learner evaluations go awry. Case studies from classrooms where the quality of evaluations is assessed will allow future pedagogies to respond to the patterns that have emerged from learners. Moreover, it is important to consider how learners can develop stronger evaluative capabilities. It may be true that this is a skill with which learners struggle (Bearman et al., 2024), and it may be true that they can never fully validate an AI output (Flenady & Sparrow, 2025). However, the right skill development should, presumably, improve the quality of the evaluations so that the claims learners make with and without AI can accurately meet the demands and conventions of their disciplines. By understanding what prerequisite knowledge is required to make strong evaluations across disciplines, writing pedagogies will be better positioned to adapt assessments to the reality of AI in writing classrooms.

Human-AI Rhetorical Collaborations

Much has been made AI's potential as a rhetorical helper. Hart-Davidson (2018) experimented with AI tools to understand their capacity for rhetoric. He concluded that AI could not independently do rhetoric but could support humans' rhetorical practices. More recently, Omizo and Hart-Davidson (2024) found that AI could mimic genre patterns but lacked the depth required to adapt those patterns to the requirements of the rhetorical context and of audiences. Wang (2024), Moulaison-Sandy (2025), and Van Poucke (2024) have also drawn attention to a chatbot's lack of engagement with the affective dimensions of writing, which hinders its ability to make judgements that meet the needs of people within the rhetorical context. These are important notes given that AI tools tend to

mimic even when the context and the affective dimensions of the context call for a different rhetorical approach.

The patterned, algorithmic rhetorical nature of AI output is believed to have not only an effect on how people write but also on the way that writing is taken up and circulated in public discourse. Wang (2024) called this a "post-rhetoric condition" where "creating discourses through symbolic means in order to induce change loses its centrality in our public life" (p. 161). This possibility, for Wang (2024), opens the possibility that "the boundaries between human discourse and machine discourse no longer matter, are further blurred, or even cease to exist" (p. 162). Under this condition, rhetoric could become disconnected from important elements of human meaning making and in the dissemination of knowledge. When the limitations for chatbots to perform rhetoric and respond to genre contexts erode, what is left is a mimicry of patterns that are common in human interactions.

This makes determining where humans and AIs entangle and interact important. These entanglements are central to the conversations Burriss and Leander (2024) and Wang and Wang (2025) put forth about the posthuman knowledge maker. Determining where human intervention fits with machine mediation was implicit in Graham's (2023) model for writing processes that use AI. From a more critical perspective, Joksimovic et al. (2023) outlined that the "division of roles between humans and machines is rather static and is usually performed on arbitrary assumptions about the expected performance and the notion of who does what better" (p. 2). They pointed to a human's ability to think abstractly, reason, and make associations, which is not necessarily unique to humans. This is where Knowles's (2024) delineation of human-in-the-loop writing processes and machine-in-the-loop has become useful. Knowles (2024) presented these "loops" as two ends of a spectrum along which activities in the writing process may lie. Some activities, such as copyediting, may see more machine mediation, hence the human would be in the loop. Other activities, such as engaging with source material, require a more hands-on approach but could get an assist from AI (hypothetically), so the machine is in the loop for a more human-driven activity.

Each use of AI has the potential to affect rhetorical strategy and its effects on audiences. It is also easy to use AI without thinking about the way one is entangled with the technology and the rhetorical effects of that entanglement (see Burriss & Leander, 2024; Wang, 2024). Future scholarship might benefit from exploring how learners' abilities to engage rhetorically within their composition context is affected by AI use. The value of embracing learner experimentation becomes prominent here, because the more that scholarship can understand how learners take up AI to support rhetorical action, the better practitioners can understand the effects of the technology on learners' abilities to

act. To support this understanding, pedagogies and assessments that help students learn where different writing activities may exist on the spectrum outlined by Knowles (2024) will be useful to help them understand the effects of AI on rhetoric and genre.

Furthermore, it is important to look further into Wang's (2024) post-rhetoric condition to determine the implications of shifting boundaries between human and machine. There is an opportunity to understand the effect of AI on rhetoric in wider social contexts, moving from our classrooms and into more public-facing rhetorical strategies (and vice versa). The knowledge that would come from these investigations would move usefully from research to pedagogy and position teachers to better account for how learners encounter AI in their daily lives.

AI and Bias

The connection between bias and AI has a prominent place in scholarly conversations. Early discussions highlighted the potential for biases in AI programming to privilege certain ontologies and epistemologies at the expense of the diversity of knowledge makers who will use AI and contribute to scholarship (e.g., Kim & Kim, 2022; Lo, 2023). Others have noted the effect on bias in evaluating AI outputs (Warschauer et al., 2023) and the influence that AI biases may have on human work (Majdik & Graham, 2024). What is clear is that bias works both ways: AI can perpetuate biases from the training data, and it is difficult to detect user biases in AI outputs, especially when the user is evaluating those outputs (Robinson & Hollett, 2024; Su & Yang, 2023). This concern is amplified when humans take AI outputs for granted, as they do with many technologies. Because human errors and biases are inevitable, humans as arbiters of quality, validity, and oversight are essential, especially in pedagogical contexts (Markauskaite et al., 2022; Robinson & Hollett, 2024).

It is one thing to acknowledge and even promote humans as central agents in ensuring quality and validity. It is another to understand how. This understanding is still nebulous in the literature. The literature suggests the potential for AI outputs to contain discriminatory, racialized, or gendered outputs (e.g., Baba et al., 2024; Chan & Hu, 2023; Smith et al., 2025). There is less information about how to effectively intervene in the biases, especially when they are biases that users bring to the chatbot. There is currently a lack of research that can help understand the nature and effect of these biases on writing, knowledge making, and pedagogies. Writing studies scholars are well-positioned to contribute to these areas because of their expertise in analyzing language and engaging with corpus-based methodologies. Once the understanding of these biases becomes more specific, it will be easier to adapt pedagogies that can support students in navigating these biases.

Moreover, it is just as necessary to look at how user perceptions of AI will impact their use of these technologies because they influence both human trust and expectations. Learners and teachers have preconceived notions of AI that influence their use of the technology (Van Poucke, 2024) and how they navigate biases. If learners go in viewing chatbots as mostly credible and valid knowledge makers, then they are at risk of putting too much trust in outputs. On the other hand, if learners dismiss AI as a flawed tool, they will disregard building an applied understanding of the versatility of its potential. These preconceived notions are biases that determine how AI systems are interpreted by users. How learners perceive AI produces a host of biases in how they take up the technology in their writing processes, craft and input prompts, and evaluate and adopt features of AI outputs. These dimensions of bias are theorized and understood little in the literature to date and warrant more examination.

Academic Integrity

Closely connected with the idea of AI and bias is the idea of academic integrity. Both topics have developed traction in conversations related to writing, literacy, and pedagogy. Cardon et al. (2023) pointed to the potential for AI use to promote academic dishonesty and the need to adhere to ethical standards for knowledge making. Johnston et al.'s (2024) survey results problematize this idea, showing that many learners were confused about shifting AI policies across their classrooms that made it difficult for them to be certain about what constituted ethical AI use in different learning contexts. This confusion is understandable given Marcel and Kang's (2024) analysis of AI policies which showed the "downloading" of AI policy development and enactment to instructors. When instructors are left to their own devices for developing AI guidelines in their learning environments, it is likely that learners will encounter a moving target for what academic integrity looks like in different learning environments.

Others have questioned the plausibility of detecting AI use rather than adapting to its influence as a knowledge building tool. Eaton (2023) offered a prominent introduction to these ideas, and the subsequent web resource developed through the Werklund School of Education (Anselmo & Kendon, 2023) offers useful tips on bringing AI tools to knowledge making endeavours. The resource has become useful to understand assessment, ethics, bias in knowledge making, and establishing guidelines for AI use in postsecondary pedagogies. The website resonates with Johnson's (2023) emphasis on the need to meet students where they are and to develop pedagogies that are responsive to the reality that is AI. It also resonates with Luo's (2024) finding that AI policies often have narrow

interpretations of originality that are at odds with the tools learners are using. As Vetter et al. (2024) explained, one of the problems may indeed be top-down policy making. Although there is uncertainty in individual teachers being responsible for developing their own AI policies (Marcel & Kang, 2024), AI-ethics are negotiated in local pedagogical contexts (Vetter et al., 2024). A more collaborative, local approach to AI policy development may respond better to how learners are taking up AI in different writing contexts because it will respond to the ethical considerations of individual contexts.

Central to Eaton's (2023) argument is that AI will re-wire knowledge making practices. AI and new innovations like neurotechnology will become increasingly embedded into knowledge making practices, making it imprudent or impractical to detect AI use. As such, it becomes important to reconsider how knowledge is built and the ethical guidelines around that knowledge making so that pedagogies can be responsive to digital influences.

At odds in these conversations is how people understand ethical knowledge making. AI challenges ontological beliefs about how writing is used, how knowledge is created, and how society integrates technologies into day-to-day life. Markauskaite et al. (2022) said that "the question of how we prepare people for an AI world is fundamentally a question of values" (p. 9). This question of values, especially where writing is concerned, is multifaceted. Some conversations focus on intellectual property and copyright (Higgs & Stornaiuolo, 2024; Kostopoulos, 2024; Luo, 2024; Pandey et al., 2025). Others center on who assumes responsibility for AI outputs. These often consider student responsibilities for what they submit (Flenady & Sparrow, 2025; Fyfe, 2022; Kumar et al., 2024). It is important to recognize that learners are aware of these ethical conversations and willing to engage with them. For example, Johnston et al.'s (2024) respondents seemed to question their own values along with ethical AI uses. Fyfe's (2022) learners spent considerable energy trying to identify an ethical "line" along which AI could be used while still maintaining the author's authenticity. Similarly, Higgs and Stornaiuolo (2024) found that youth were concerned about creativity and authenticity, as well as AI's influence on issues of equity and access.

What these conversations show is that the question of values is shared by teachers and learners alike. Within this exists common ground to open discussions and co-develop policy, which is something that Higgs and Stornaiuolo (2024) and Vetter et al. (2024) have promoted. The conversations on academic integrity are fluid, and determining a single "ethic" for AI use in pedagogies seems unrealistic. Building upon common ground, co-developing policy, and having teachers and students evaluate policies together would offer a strong platform for understanding AI's role and usage in writing related classes.

While conversations and collaborations are useful and necessary, an important next step will be to see values in action. Conversations around academic integrity in writing-related contexts remains a largely theoretical exercise, and there is a need to see how faculty-student integrity collaborations play out. There is equally a need to see policies that engage with the values-oriented conversations in their local contexts. From here, other teachers can adapt policies to suit their own classrooms. This will also make it easier to make links between policy and practice. How teachers and learners enact policies in writing classrooms is a looming question that requires attention. In answering this question, it will become easier to understand the ethical conversations in all their dimensions and the different pressure points that exist as teachers and learners establish values and ethics within their contexts. These are messy conversations, and scholarship that dwells in, engages with, and accounts for the mess will lead to a more nuanced conversation around academic integrity in writing classrooms.

Limitations

The scholarly record about AI, writing, and literacy is evolving quickly. We approached this review at a time when publications on AI were becoming more frequent. No sooner would we discuss a theme than we would become aware of more articles on that theme in follow-up searches or in consultations with colleagues. Other times, a category like "critical thinking" appeared frequently, but the uses of the term were disparate and entangled with other themes. As conversations become more established, some may gain firmer footing and would benefit from review. This review offers a snapshot of conversations that is limited in two primary ways: 1) a rapidly evolving publishing environment, and 2) a corpus of publications that has not yet established a foundation or found an equilibrium in its conversations. We have pulled together clusters of promising future conversation pathways, and we have inferred next steps. The study of what held true and what evolved differently than we anticipated will be both prudent and necessary.

Returning to the suggestions will also allow for a more selective review. We included a large sample because it was difficult to determine what conversations were representative of an emerging field. We tried to wrap our intellectual arms around as much of the conversation as we could. Future reviews would benefit from curating sources using a methodological appraisal (see Scott et al., 2018; Torgerson, 2007) or meta-analysis (Graham et al., 2018). Appraising methodological quality will add a dimension of validity that is only possible with a more established record.

Conclusions

This review set out to develop a thematic map of what is known and not known about AI in writing and writing pedagogies. We considered gaps in the early conversations on AI, with an eye towards what may be useful next steps for researchers and practitioners. AI literacy and academic integrity stood out as prominent themes that address conceptual issues related to AI. The conversations on evaluation, rhetoric, and bias exist more on the level of praxis, both in terms of writing itself and in the pedagogies that focus on writing and literate practices. The conceptual conversations require two things: synthesis and evidence in action. Synthesis will enable a more focused understanding of how research on applies to specific writing contexts. Evidence in action will help to refine pedagogies that respond to AI's influence. For academic integrity, this means developing a larger corpus on how policies are enacted in writing classrooms. For AI literacy frameworks, a better understanding of how they are enacted will help to synthesize the most important dimensions that will affect pedagogies going forward.

Much of this understanding can come from better understanding specific elements of AI use. This involves being more aware of the role unconscious bias play in working with AI outputs. It also requires that pedagogies can help learners to account for these biases, not just the built-in biases of AI systems. Developing knowledge about how learners can become better evaluators of AI outputs—connecting the outputs to the requirements of the writing context—will be equally important. If evaluation is central to AI use, then it is imperative that pedagogies can provide a scaffold for learners to make better evaluations. Similarly, pedagogies that can model for learners how to engage rhetorically within their composition context if they are using AI will be valuable. So too will pedagogies that can draw attention to AI's effect on rhetoric in wider public discourses, which will allow learners to make connections between their own AI use and their everyday experiences.

References

Adu, P. (2019). *A step-by-step guide to qualitative data coding*. Routledge.

Anselmo, L., & Kendon, T. (10 October, 2023). Exploring artificial intelligence and assessments: A guide for instructors considering artificial intelligence in their assessments. *Taylor Institute for Teaching and Learning*. https://taylorinstitute.ucalgary.ca/resources/exploring-artificial-intelligence-and-assessments

Anson, D.W.J. (2024). The impact of large language models on university students' literacy development: A dialogue with Lea and Street's academic literacies framework. *Higher Education Research & Development, 43*(7), 1465-1478. https://doi.org/10.1080/07294360.2024.2332259

Baba, K., Faddouli, E., & Cheimanoff, N. (2024). Mobile-optimized AI-driven personalized learning: A case study at Mohammed VI Polytechnic University. *International Journal of Interactive Mobile Technologies*, *18*(4), 81-96. https://doi.org/10.3991/ijim.v18i04.46547

Bearman, M., & Ajjawi, R. (2023). Learning to work with the black box: Pedagogy for a world with artificial intelligence. *British Journal of Educational Technology, 54*(5), 1160-1173. https://doi.org/10.1111/bjet.13337

Bearman, M., Tai, J., Dawson, P., Boud, D., & Ajjawi, R. (2024). Developing evaluative judgement for a time of generative artificial intelligence. *Assessment and Evaluation in Higher Education*, 1-13. https://doi.org/10.1080/02602938.2024.2335321

Beck, S.W., & Levine, S. (2024). The next word: A framework for imagining the benefits and harms of generative AI as a resource for learning to write. *Reading Research Quarterly, 59*(4), 706-715. https://doi.org/10.1002/rrq.567

Burriss, S.K., & Leander, K. (2024). Critical posthumanist literacy: Building theory for reading, writing, and living ethically with everyday artificial intelligence. *Reading Research Quarterly, 59*(4), 560-569. https://doi.org/10.1002/rrq.565

Cardon, P., Fleischmann, C., Aritz, J., Logemann, M. & Heidewald, J. (2023). The challenges and opportunities of AI-assisted writing: Developing AI literacy for the AI age. *Business and Professional Communication Quarterly, 86*(3), 257–295. https://doi.org/10.1177/23294906231176517

Cummings, R.E., Monroe, S.M., & Watkins, M. (2024). Generative Ai in first-year writing: An early analysis of affordances, limitations, and a framework for the future. *Computers and Composition, 71*, 1-11. https://doi.org/10.1016/j.compcom.2024.102827

Davis, D., Jindal-Snape, D., Collier, C., Digby, R., Hay, P., & Howe, A. (2012). Creative learning environments in education systematic literature review. *Thinking Skill and Creativity, 8*, 80-91. https://doi.org/10.1016/j.tsc.2012.07.004.

De Matas, J. (2023). ChatGPT and the future of writing about writing. *Double Helix, 11*, 1-7. https://doi.org/10.37514/DBH-J.2023.11.1.09

Eaton, S.E. (2023). Postplagiarism: Transdisciplinary ethics and integrity in the age of artificial intelligence and neurotechnology. *International Journal for Educational Integrity, 19*(23). https://edintegrity.biomedcentral.com/articles/10.1007/s40979-023-00144-1

Flenady, G., & Sparrow, R. (2025). Cut the bullshit: Why GenAI systems are neither collaborators nor tutors. *Teaching in Higher Education*, 1-10. https://doi.org/10.1080/13562517.2025.2497263

Fyfe, P. (2022). How to cheat on your final paper: Assigning AI for student writing. *AI & Society, 38*, 1395-1405. https://doi.org/10.1007/s00146-022-01397-z

Girdharry, K., & Khachatryan, D. (2023). Meaningful writing in the age of generative artificial intelligence. *Double Helix: A Journal of Critical Thinking and Writing, 11*. https://doi.org/10.37514/DBH-J.2023.11.1.04

Graham, S., Liu, X., Bartlett, B., Ng, C., Harris, K.R., Aitken, A., Barkel, A., Kavanaugh, C., & Talukdar, J. (2018). Reading for writing: A meta-analysis of the impact of reading interventions on writing. *Review of Educational Research, 88*(2), 243-284. https://doi.org/10.3102/0034654317746927

Hassel, H. (2013). Research gaps in *Teaching English in the Two-Year College*. *Teaching English in the Two-Year College, 40*(4), 343-363. https://doi.org/10.58680/tetyc201323601

Higgs, J., & Stornaiuolo, A. (2024). Being human in the age of generative AI: Young people's ethical concerns about writing and living with machines. *Reading Research Quarterly, 59*(4), 632-650. https://doi.org./10.1002/rrq.552

Johnson, G.P. (2023). Don't act like you forgot: Approaching another literacy "crisis" by (re)considering what we know about teaching writing with and through technologies. *Composition Studies, 51*(1), 169-175. https://compstudiesjournal.com/wp-content/uploads/2023/06/johnson.pdf

Johnston, H., Wells, R.F., Shanks, E.M., Boey, T., & Parsons, B.N. (2024). Student perspectives on the use of generative artificial intelligence technologies in higher education. *International Journal for Educational Integrity, 20*(2), 1-21. https://doi.org/10.1007/s40979-024-00149-4

Joksimovic, S., Ifenthaler, D., Marrone, R., De Laat, M., & Siemens, G. (2023). Opportunities of artificial intelligence for supporting complex problem-solving: Findings from a scoping review. *Computers and Education: Artificial Intelligence, 4*, 100138.

Kim, N., & Kim, M. (2022). Teacher's perceptions of using an artificial intelligence-based educational tool for scientific writing. *Frontiers in Education, 7,* 1-13. https://doi.org/10.3389/feduc.2022.755914

Kostopoulos, E. (2025). Student use of generative AI as a composing process supplement: Concerns for intellectual property and academic honesty. *Computers and Composition, 75*. https://doi.org/10.1016/j.compcom.2024.102894

Kumar, P.C., Cotter, K., & Cabrera, L. (2024). Taking responsibility for meaning and mattering: An agential realist approach to generative Ai and literacy. *Reading Research Quarterly, 59*(4), 570-578. https://doi.org/10.1002/rrq.570

Lee, S. J., & Kwon, K. (2024). A systematic review of AI education in K-12 classrooms from 2018 to 2023: Topics, strategies, and learning outcomes. *Computers and Education: Artificial Intelligence*, 100211.

Lin, C., Zhou, K., Li, L., & Sun, L. (2025). Integrating generative Ai into digital multimodal composition: A study of multicultural second-language classrooms. *Computers and Composition, 75*. https://doi.org/10.1016/j.compcom.2024.102895

Lo, L. (2023). The art and science of prompt engineering: A new literacy in the information age. *Internet Reference Services Quarterly*, 1-8. https://doi.org/10.1080/10875301.2023.2227621

Luo, J. (2024). A critical review of GenAI policies in higher education assessment: A call to reconsider the 'originality' of students' work. *Assessment & Evaluation in Higher Education*, 1-14. https://doi.org/10.1080/02602938.2024.2309963

Majdik, Z.P., & Graham, S.S. (2024). Rhetoric of/with AI: An introduction. *Rhetoric Society Quarterly, 54*(3), 222-231. https://doi.org/10.1080/02773945.2024.2343264

Marcel, F., & Kang, P. Examining AI guidelines in Canadian universities: Implications on academic integrity in academic writing. *Discourse and Writing, 34*. https://doi.org/10.31468/dwr.1051

Markauskaite, L., Marrone, R., Poquet, O., Knight, S., Martinez-Maldonado, R., Howard, S., Tondeur, J., De Laat, M., Buckingham Shum, S., Gašević, D., & Siemens, G. (2022). Rethinking the entwinement between artificial intelligence and human learning: What capabilities do learners need for a world with AI? *Computers and Education. Artificial Intelligence*, *3*, 1-16. https://doi.org/10.1016/j.caeai.2022.100056

Mollick, E.R., & Mollick, L. (2023). Using AI to implement effective teaching strategies in classrooms: Five strategies, including prompts. *The Warton School Research Paper*. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4391243

Moulaison-Sandy, H. (2025). *AI and authorship in scholarly communication: Writing with intelligence*. Routledge.

Ng, D., Leung, J., Chu, S., & Qiao, M. (2021). Conceptualizing AI literacy: An exploratory review. *Computers and Education: Artificial Intelligence, 2,* 1-11. https://doi.org/10.1016/j.caeai.2021.100041

Omizo, R., & Hart-Davidson, B. (2024). Is genre enough? A theory of genre signaling as generative AI rhetoric. *Rhetoric Society Quarterly, 54*(3), 272-285. https://doi.org/10.1080/02773945.2024.2343615

Pandey, H.L., Bhusal, P.C., & Niraula, S. (2025). Large language models and digital multimodal composition in the first-year composition classrooms: An encroachment and/or enhancement dilemma. *Computers and Composition, 75*. https://doi.org/10.1016/j.compcom.2024.102892

Risko, V. J., Roller, C. M., Cummins, C., Bean, R. M., Collins, Block, Anders, C., et al. (2008). A critical analysis of research on reading teacher education. *Reading Research Quarterly, 43*, 252-288. https://doi.org/10.1598.RRQ.43.3.3

Robinson, B., & Hollett, T. (2024). Literacy in the age of AI. *Reading Research Quarterly, 59*(4), 555-559. https://doi.org/10.1002/rrq.581

Scott, C.E., McTigue, E., Miller, D.M., & Washburn, E. (2018). The what, when and how of preservice teachers and literacy across the disciplines: A systematic literature review of nearly 50 years of research. *Teaching and Teacher Education, 73*(2), 1-13. https://doi.org/10.1016/j.tate.2018.03.010

Smith, B.E., Shimuzu, A.Y., Burriss, S.K., Hundley, M., & Pendergrass, E. (2025). Multimodal composing with generative Ai: Examining preservice teachers' processes and perspectives. *Computers and Composition, 75*. https://doi.org/10.1016/j.compcom.2024.102896

Su, J., & Yang, W. (2023). Unlocking the power of ChatGPT: A framework for applying generative AI in education. *ECNU Review of Education, 6*(3), 355-366. https://doi.org/10.1177/20965311231168423

Tai, J., Ajjawi, R., Boud, D., Dawson, P., Panadero, E. (2018). Developing evaluative judgement: Enabling students to make decisions about the quality of work. *Higher Education, 76*, 467-481. https://doi.org/10.1007/s10734-017-0220-3

Torgerson, C.J. (2007). The quality of systematic reviews of effectiveness in literacy learning in English: A 'tertiary' review. *Journal of Research in Reading, 30*, 287-315. http://refhub.elsevier.com/S0742-051X(17)31085-5/sref83

Tseng, W., & Warschauer, M. (2023). AI-writing tools in education: If you can't beat them, join them. *Journal of China Computer-Assisted Language Learning*. https://doi.org/10.1515/jccall-2023-0008

Tzirides, A. O., Zapata, G., Kastania, N. P., Saini, A., Castro, V., Ismael, S., You, Y., Afonso dos Santos, T., Searsmith, D., O'Brien, C., Cope, B., & Kalantzis, M. (2024). Combining human and artificial intelligence for enhanced AI literacy in higher education. *Computers and Education Open*, *6*, Article 100184. https://doi.org/10.1016/j.caeo.2024.100184

Van Poucke, M. (2024). ChatGPT, the perfect virtual teaching assistant? Ideological bias in learner chatbot interactions. *Computers and Composition, 73*. https://doi.org/10.1016/j.compcom.2024.102871

Vetter, M.A., Lucia, B., Jiang, J., Othman, M. (2024). Towards a framework for local interrogation of AI ethics: A case study on text generators, academic integrity, and composing with ChatGPT. *Computers and Composition, 72*. https://doi.org/10.1016/j.compcom.2024.102831

Wang, Z. (2024). Post-rhetoric: A rhetorical profile of the generative artificial intelligence chatbot. *Rhetoric Review, 43*(3), 155-172. https://doi.org/10.1080/07350198.2024.2351723

Wang, Z., & Wang, C. (2025). A posthumanist approach to AI literacy. *Computers and Composition, 76*. https://doi.org/10.1016/j.compcom.2025.102933

Warschauer, M., Tseng, W., Yim, S., Webster, T., Jacob, S., Du, Q., & Tate, T. (2023). The affordances and contradictions of AI-generated text for writers of English as a second or foreign language. *Journal of Second Language Writing*, *62*, 101071. https://doi.org/10.1016/j.jslw.2023.101071

Zednik, C. (2021). Solving the black box problem: A normative framework for explainable artificial intelligence. *Philosophy & Technology, 34*(2), 265-288. https://doi.org/10.1007/s13347-019-00382-7