

Co-Constructing AI Boundaries: Agency, Judgment, and Ethical Literacy in AI-Mediated Meaning-Making

W. Ian O'Byrne

Department of Teacher Education, College of Charleston

Literacy Research: Theory, Method, and Practice (2026), 1-22. SAGE Publications.

DOI: <https://doi.org/10.1177/23813377261476996>

ORCID: <https://orcid.org/0000-0003-0307-9538>

This page hosts the author accepted manuscript, reconciled line by line against the published article: Sage's copyediting, two citations added in production, the published reference list, and the author biography are all carried here. What remains is typesetting. Cite the version of record via the DOI.

Abstract

As generative artificial intelligence (AI) becomes embedded in everyday literacy practices, pre-service teachers (PSTs) must decide what work remains distinctly human in AI-mediated meaning-making. This qualitative case study investigates how PSTs enact boundary work with generative AI during a semester-long literacy ethnography. Drawing on Nexus Analysis and digital trace data from Google's NotebookLM, I analyze moments of agency, refusal, and ethical judgment within the human–AI evaluative loop. Findings identify two contrasting interaction profiles, the Orchestrator and the Outsourcer. Orchestrators retained epistemic authority by repeatedly constraining, correcting, and re-authoring the AI system's output, while Outsourcers delegated organizational and interpretive work to the system early and accepted its framing with little revision. These differences in how PSTs worked with the AI system were visible only through process-level analysis of interaction traces, as the final artifacts appeared deceptively similar. The study reconceptualizes human-in-the-loop as an interactional literacy practice rather than a technical safeguard. Implications highlight the need for teacher education to design "loops worth living in" that normalize productive friction and professional refusal, ensuring AI amplifies rather than replaces human judgment.

Keywords: AI literacy, nexus analysis, human-in-the-loop, teacher education

Introduction

The rapid integration of artificial intelligence (AI) into educational settings has introduced a new class of pedagogical "wicked problems," particularly around writing, inquiry, and professional judgment (Scalzer et al., 2025). Within literacy education, these tools and models are increasingly used to

summarize, synthesize, draft, organize, and reformulate text. As a result, they do not simply add efficiency to existing workflows. They also reshape where interpretation happens, how authority is distributed, and what counts as authorship in classroom work.

These questions are especially pressing in teacher education, where pre-service teachers (PSTs) occupy a distinctive dual positionality. As students, they complete graded academic work and, as emerging professionals, learn to exercise pedagogical judgment. Their decisions about AI use are consequential in two directions. PSTs must determine what constitutes acceptable support within coursework aligned with academic integrity expectations. They are simultaneously apprenticing into professional habits of responsibility, justification, and care that they will later carry into their own classrooms. For PSTs, AI adoption is not simply a matter of convenience. It is a rehearsal for future professional decision-making and a foundational part of how they learn to design learning for others. This makes teacher education a particularly revealing site for examining how the boundaries of AI-mediated literacy practice are negotiated.

Recent scholarship calls for AI literacy that moves beyond technical proficiency to encompass critical evaluation, transparency, and ethical judgment (Long & Magerko, 2020). In this study, AI literacy is defined as the capacity not only to operate AI tools and models but also to critically discern which tasks should or should not be delegated to them, and how those choices shape learning and professional responsibility (Chiu et al., 2024). Relatedly, this study also adopts a bounded definition of AI ethics. Rather than addressing macro-level ethical debates such as environmental costs, global labor exploitation, or platform governance, this study focuses specifically on ethical literacy within situated practice (Raffaghelli et al., 2026). It examines how learners justify AI use, negotiate authorship and responsibility, and determine what epistemic labor must remain uniquely human during composing and inquiry (Higgs & Stornaiuolo, 2024).

To be precise about terms, in this article, I use AI to refer to the broad category of computational systems that perform tasks associated with human cognition (Siemens et al., 2022), and generative AI to refer to systems that produce new text, images, or other content in response to prompts (Bandi et al., 2023).

From a sociocultural literacy studies perspective, these decisions are not best understood as matters of tool adoption alone. Classrooms are also social and interpretive spaces in which boundaries between human and machine work are negotiated in practice (Burnett, 2011). In this article, I use boundary work to refer to the interactional process through which learners decide what cognitive, interpretive, and ethical labor remains human and what may be delegated to AI (Kim et al., 2025). These boundaries are enacted through prompts, revisions, constraints, refusals, and reflective commentary rather than merely declared in policy statements. Understanding AI literacy, therefore, requires attention to both process and product, as final artifacts may appear polished and coherent while obscuring the judgments, interventions, and refusals that shaped them. For this reason, the analysis

examines interactional traces rather than relying solely on completed assignments or self-reported attitudes (Palmquist et al., 2025).

I investigate how PSTs engaged with generative AI during a semester-long literacy ethnography using Google's NotebookLM, a source-grounded generative AI environment. I ask how PSTs negotiated acceptable AI use during authentic literacy work, how those decisions became visible in interactions, and why they were professionally consequential. Using interaction logs, reflections, and final course artifacts, I trace how participants enacted agency through iterative correction, refusal of generic output, and evaluation of epistemic quality. By foregrounding boundary work as an interactional accomplishment, this study reframes human-in-the-loop (HITL) as a literacy practice that learners must actively construct rather than a technical safeguard they simply comply with.

Theoretical Perspectives

Generative AI complicates traditional approaches to assessing literacy learning because polished final products can obscure the judgments, negotiations, and revisions that shaped their production (Bozkurt et al., 2024). If the goal is to understand how future educators work responsibly with AI, final artifacts alone are insufficient. What matters is how students prompt, evaluate, revise, reject, and justify AI-generated contributions along the way. This study adopts a process-oriented lens, treating interaction traces as evidence of how agency, responsibility, and ethical judgment become visible in situated literacy activity (Hancock et al., 2020).

This study brings together four linked theoretical moves. Critical AI Literacy provides the normative orientation, Nexus Analysis provides the sociocultural analytic frame, boundary work and agency serve as the core analytic constructs, and interaction traces function as the empirical site where these negotiations become visible. Together, these perspectives support a process-level analysis of how PSTs enact ethical and epistemic judgment in practice.

Critical AI Literacy as a Normative Orientation

This study is grounded in Critical AI Literacy, understood not as technical competence alone but as an ethical and sociopolitical orientation toward AI-mediated meaning-making (Velandar et al., 2024). Here, Critical AI Literacy names the ethical and interpretive demands of working with AI in educational settings, extending posthuman approaches that emphasize ethical relationships among humans, technologies, and literacy practices (Burris & Leander, 2024; Leander & Burris, 2020). It asks not only whether learners can use AI effectively, but whether they can justify that use in ways consistent with professional responsibility, authorship, and pedagogical judgment, and why such use matters, for whom, and to whose benefit (Leander & Burris, 2020).

For teacher education, this orientation matters because PSTs are learning to make judgments they will later institutionalize in classrooms. Decisions about delegation and authorship are, therefore, also

decisions about what kinds of pedagogical judgment they are learning to inhabit. Critical AI Literacy clarifies why questions of authorship, responsibility, transparency, and human judgment matter, while leaving open the empirical question of how such commitments are enacted in everyday literacy work (O’Byrne et al., 2024).

AI as Mediational Means and Nexus Analysis

If Critical AI Literacy clarifies the stakes of the problem, Nexus Analysis provides a way to study how those stakes are enacted in practice (Scollon & Scollon, 2004). It frames AI-mediated literacy activity not as isolated tool use, but as socially organized action shaped by prior histories, moment-to-moment interaction, and institutional conditions. Generative AI is treated here as a mediational means: a socially and historically shaped instrument that restructures the activity it mediates (Wertsch, 1991). Nexus Analysis examines mediation across three analytic cycles:

- **Historical Body:** learners’ prior knowledge, literacy histories, identities, and technological habitus that shape orientations toward authority and responsibility (Norouzi et al., 2019).
- **Interaction Order:** the moment-by-moment organization of the human–AI relationship, visible in patterns of prompting, latency, correction, and revision through which authority is asserted, deferred, or reclaimed (Lindgren, 2025).
- **Discourses in Place:** institutional and technological narratives that shape what counts as legitimate AI use, including constraints embedded in the tool environment (Nguyen & Hekman, 2024).

In this study, these cycles are inferred from patterned action across prompts, revisions, and reflections rather than from the AI system's internal states or properties.

HITL as Interactional Practice and Cognitive Friction

Within this framework, the study adopts a HITL stance, conceptualizing AI as a cognitive amplifier rather than a replacement for human judgment (Natarajan et al., 2025). Here, HITL is treated as an interactional literacy practice in which learners retain or cede epistemic authority as AI participates in composing and inquiry, foregrounding learners' critical agency within postdigital educational environments (Code, 2025). Human judgment remains active when learners set goals, curate inputs, evaluate outputs, and shape meaning, while AI supports retrieval, organization, and synthesis (Li et al., 2025).

Meaningful amplification becomes visible when it generates cognitive friction, the resistance that arises when machine output challenges, constrains, or falls short of human intent (Schuetz & Venkatesh, 2020). When interaction becomes too smooth, evaluative labor can disappear from view and agency can give way to algorithmic passivity. By contrast, moments of hesitation, redirection,

correction, and refusal provide evidence that human judgment remains active in the loop (Ackerman, 2025).

Boundary Work and Agency

To operationalize this focus, the analysis centers on two interrelated constructs: boundary work and agency. Boundary work functions as an analytic lens rather than a conclusion in itself. It refers to the interactional practices through which learners negotiate the division of cognitive, interpretive, and ethical labor between themselves and AI, determining what remains human and what may be delegated (Bearman & Ajjawi, 2023). In this study, boundary work is enacted through restriction, correction, refusal, modification, and articulation of AI's appropriate role.

Agency is examined through those same practices, especially in moments of interruption, correction, refusal, and re-authoring. Agency often appears not as maximal AI use but as intentional constraint and critical evaluation (Nopas, 2025). Together, boundary work and agency provide the study's primary analytic vocabulary for describing how PSTs negotiate control, responsibility, and accountability as AI participates in meaning-making.

Institutional and Affective Tensions

AI-mediated boundary work unfolds within contradictory institutional and professional discourses. For PSTs, these tensions are intensified by their dual positionality as students accountable to present academic norms and emerging professionals learning the pedagogical norms they will later be expected to enact (Schiera, 2019). Across educational contexts, AI is simultaneously framed as a necessary innovation and a threat to integrity, legitimacy, and professional judgment (Adillón et al., 2024). This mixed messaging can produce uncertainty, hesitation, and shame, shaping how learners publicly justify and privately negotiate AI use (McNealis, 2025). In this study, these tensions are treated as part of the discursive terrain within which boundary work becomes professionally consequential (Wise & O'Byrne, 2016).

Interaction Traces as Evidence of Social Action

Because agency and boundary work are enacted in interaction, this study treats AI interaction logs as traces of social action rather than neutral technical records. Prompts, revisions, pauses, and transformations provide process-level evidence of how learners accept, contest, constrain, or rework AI contributions. This makes visible forms of judgment that would be difficult to recover from final products alone. The analysis, therefore, attends to latency, revision, and reflection as evidence of how epistemic authority is negotiated moment by moment in AI-mediated literacy work (Wang & Wang, 2025).

Taken together, these perspectives provide a coherent basis for the study. Critical AI Literacy clarifies why AI-mediated authorship, judgment, and responsibility matter in teacher education. Nexus Analysis provides a sociocultural frame for tracing how those concerns unfold across histories,

interactions, and institutional conditions. Boundary work and agency supply the core analytic vocabulary for examining how learners negotiate control and responsibility. Interaction traces, in turn, provide the empirical record through which those negotiations become visible.

Methodology

This qualitative case study (Naeem & Thomas, 2025) was conducted at a small liberal arts institution in the southeastern United States in an undergraduate language and literacy course in the School of Education. The course also fulfilled General Education Social Sciences and Multicultural Understanding and Global Citizenship (MUGC) requirements, which accounts for the presence of one participant outside of teacher education. Participants were 23 undergraduate students. Most were PSTs preparing for middle grades or secondary education, with a small number (2) from elementary or early childhood education. Four participants were identified as male and the remainder as female. This mix reflects both the course's teacher education focus and its broader General Education enrollment.

The study was situated within a semester-long literacy ethnography sequence. For a midterm assignment, students conducted a brief literacy autoethnography of their own lives over five days. For the final project, they expanded outward to study a community of their choosing, collecting observations, artifacts, reflections, interviews, and other evidence to analyze how literacy, language, identity, culture, power, and access operated within that community. Students were expected to connect these local observations to broader course concepts related to equity, communication, assimilation, accommodation, and cultural identity.

As part of the final project, students were required to use Google's NotebookLM as a cognitive partner during data collection, reflection, and analysis. NotebookLM was intentionally selected because its retrieval-augmented, source-grounded structure enabled students to work with a bounded corpus of self-collected materials rather than an open web search. Pedagogically, the tool was framed not as a replacement for human analysis, but as a structured thinking partner that could help students organize observations, surface patterns, question assumptions, and test emerging interpretations. In a more conventional version of the assignment, students might have worked with a peer to discuss what they were seeing in their data. Here, NotebookLM was used to approximate that dialogic analytic support while leaving responsibility for interpretation with the student.

Students first selected a community to study and began documenting observations, artifacts, questions, and reflections in Google Drive. They then uploaded those materials to NotebookLM and used guided prompts to identify patterns in language, culture, identity, power, access, assimilation, and accommodation. After reviewing NotebookLM's responses, students returned to their own notes to indicate what they agreed with, disagreed with, or wished to revise. This reflective cycle of collecting data, generating interpretation, engaging AI feedback, and reasserting human judgment was repeated across multiple stages of the project and culminated in a final presentation and written reflection. All

23 students shared their NotebookLM notebooks with the instructor-researcher, providing access to the full interaction corpus used in this study.

This research was guided by the following question: *How do pre-service educators negotiate boundaries of acceptable AI use during authentic literacy work?* To address this, the analysis examined (a) how boundary work decisions were enacted through patterned interactional practices with AI, and (b) how ethical tensions and professional legitimacy concerns surfaced as PSTs navigated AI use across their dual roles as students and emerging professionals.

AI System Context

Because AI-mediated boundary work unfolds in interaction, this study treats system-generated interaction logs as digital traces of social action. Students' work in NotebookLM produced time-stamped records of prompts, responses, and revisions that render otherwise invisible evaluative decisions visible for analysis. To clarify how interactional evidence was generated and captured, this section outlines the AI system's log structure used for analysis.

NotebookLM

Google's NotebookLM served as the AI system for this study. NotebookLM is designed for writing and inquiry with user-provided materials, allowing learners to upload or link sources and then interact with an AI system to generate summaries, explanations, and draft text grounded in those sources. Because the tool structures interaction around a curated source set, it foregrounds practices central to disciplinary literacy (selection, evaluation, synthesis, and verification) rather than open-ended generation. In this study, NotebookLM functioned as a constrained environment in which learners' epistemic decisions could be observed through their source choices, prompting moves, and responses to AI output.

Retrieval-Augmented Generation (RAG)

NotebookLM operates as a RAG system, which differs from open-ended large language models that generate responses primarily from a general training corpus. In RAG systems, AI output is produced in relation to a retrievable set of user-provided documents, enabling responses that are more context-specific and traceable to sources. This distinction matters methodologically as learners' knowledge-making depended in part on what they chose to include in the source set, and their evaluative work often took the form of verifying whether the AI system's claims were warranted by those materials.

Data Sources

Data were triangulated across multiple sources to support analytic rigor and transparency. Primary data consisted of digital interaction logs exported from NotebookLM in JSON format, supplemented by final student-authored ethnography projects, weekly and end-of-semester written reflections

documenting students' decision-making processes, and researcher field notes recording classroom interactions and moments of productive struggle. Each source is described below.

Interaction Logs (JSON) as Digital Trace Data

The primary data source was a set of time-stamped interaction logs exported from NotebookLM in JSON (JavaScript Object Notation) format. JSON is a structured, text-based format that stores entries as key–value records. In this dataset, each entry included message content and associated metadata, such as timestamps, speaker role, and interaction sequencing. Treating these logs as digital trace data, I analyzed not only the textual content of prompts and responses but also process markers such as prompt iterations and response timing. This provided evidence of how participants enacted agency through intervening, revising, refusing, and accepting AI output in the moment.

Participant Commentary and Ethical Context

Students were not treated as passive subjects of data extraction. Participants were aware that their AI-mediated work would be examined as part of the study, and they contributed reflective commentary through course reflections, class discussions, and end-of-course responses about what the tool surfaced, where it was useful, and where its interpretations felt insufficient or misleading. In this way, their reflections helped contextualize the interaction logs and informed the analysis's interpretive frame without positioning them as formal co-researchers in the coding process. All data collection followed Institutional Review Board-approved protocols.

Analysis

To interpret AI-mediated interaction, analysis was guided by Nexus Analysis (Scollon & Scollon, 2004), which examines social action across three interconnected dimensions. Interaction Order attends to the moment-by-moment organization of authority and responsibility within specific events, such as prompt–response–revision sequences. Historical Body captures the patterned habits and dispositions participants bring to these interactions over time (Wohlwend, 2020). Discourses in Place refer to the institutional, cultural, and technological conditions that shape what actions are possible or legitimate within the interaction, including constraints embedded in the AI system itself (Scollon & Scollon, 2004).

Unit of Analysis: The Trio

The primary unit of analysis was a tripartite interactional chain, referred to as the Trio: student prompt → AI response → student reflection. Following Nexus Analysis, the interactional chain names the *unit* segmented and coded in this study as a bounded sequence of connected social actions. This chain, in turn, renders visible what I term the *evaluative loop*: the iterative process through which a student prompts the AI system, judges the adequacy of its output, and reworks, redirects, or accepts it. The interactional chain is the observable structure, while the evaluative loop is the recurring cycle of

judgment that becomes visible within that structure. This distinction matters because agency in AI-mediated work is enacted less through any single prompt or response than through repeated passes through evaluation and revision.

The reflection component of each Trio refers specifically to the reflective note a student recorded and preserved through NotebookLM's *Convert to Source* function, which appended their interpretive commentary directly into the working corpus. This in-system reflection is distinct from the weekly and end-of-semester course reflections described in the Data Sources section; whereas the latter offered retrospective, contextualizing commentary, the Convert-to-Source reflection captured in-the-moment evaluative stances toward specific AI output. Because it sits inside the interactional chain rather than outside it, this reflection provides direct evidence of boundary work, ethical reasoning, and interpretive stance at the point of decision.

Treating the Trio as the primary unit of analysis ensured that claims about agency were grounded in observable interactional sequences rather than inferred solely from final artifacts or self-report data. This unit allowed distinctions to be drawn between qualitatively different patterns of engagement, which are elaborated through cross-case analysis.

Analytic Procedure

Analytic work proceeded in four stages: initial segmentation of the interaction logs, micro-interactional analysis of each Trio, iterative qualitative coding, and cross-case comparison. The overall goal was to make visible how agency, boundary work, and epistemic judgment unfolded in real time across AI-mediated literacy activity, while avoiding over-interpretation of final products alone.

Interaction Log Segmentation. First, all NotebookLM JSON logs were de-identified and segmented into Trios, the basic interactional unit used in this study. A Trio consisted of a student prompt, the AI system's response, and the student's reflective note, preserved through NotebookLM's *Convert to Source* function. This yielded a bounded interactional chain that could be read as a sequence of connected social actions rather than as isolated turns. I treated each Trio as the primary site of analysis because it captured the moment when students prompted, evaluated, revised, constrained, or accepted AI output.

Micro-Interactional Trio Analysis. To support a close reading of these sequences, I conducted micro-interactional analysis of each Trio. This meant examining the prompt-response-reflection sequence line by line to identify how students organized their interaction with AI, where they exercised judgment, and how they set boundaries for acceptable AI use. Within this process, I attended to three interpretive indicators: latency, prompt evolution, and edit distance. These three were selected because each draws on a different modality of evidence in the trace data (i.e., temporal, linguistic, and textual), allowing interactional judgment to be triangulated across the timing of a student's response, the wording of their prompts, and the distance between the AI output and the final text. Latency is the time between the AI system's output and the student's next action. Rather than treating pauses as fixed

numerical signals, I interpreted them qualitatively as possible signs of uptake, hesitation, reconsideration, or evaluative delay. Prompt evolution refers to how prompts changed over time in specificity, constraint, role assignment, and analytic demand. Edit distance refers to the degree of divergence between AI-generated language and the student’s final wording or structure. These were not used as stand-alone metrics, but as interpretive signals of deliberation, reframing, and epistemic control across AI-mediated interaction.

To organize the first pass of micro-interaction analysis, I used a five-part schema adapted from Mollick and Mollick's (2023) framework, chosen because it segments AI-supported academic work into discrete, observable phases and could therefore be repurposed from a pedagogical scaffold into an analytic lens: Inputs, Prompts, Outputs, Integration, and Reflection. Inputs captured what students uploaded, selected, or curated as context for NotebookLM. Prompts captured how students directed or constrained the AI system. Outputs captured what the system produced and how students responded to it. Integration captured how students transformed, rejected, or incorporated AI output into their own work. Reflection captured how students articulated their stance toward the interaction, including questions of authorship, responsibility, ethics, and legitimacy. In this study, the first three components primarily captured the interactional surface of the task, while Integration and Reflection captured the student’s boundary work around authorship and responsibility (Table 1).

Table 1. Coding Scheme for Human-AI Interactional Processes

CODING COMPONENT	NEXUS CONSTRUCT	ANALYTIC FOCUS
Inputs	Historical Body	Prior literacy habits, source curation, and professional background.
Prompts	Interaction Order	Constraint-setting, turn-taking, and the assertion of epistemic stance.
Outputs	Discourses in Place	Negotiation of tool constraints, citation accuracy, and "AI voice."
Integration	Boundary Work	Evidence of human transformation, rejection, or synthetic re-authoring.
Reflection	Societal Discourses	Ethical justification, legitimacy concerns, and professional identity.

Note. Adapted from Scollon and Scollon’s (2004) "Nexus of Practice" and Wertsch’s (1991) "Mediational Means."

Iterative Qualitative Coding. The five-part schema functioned as an initial organizing structure for micro-interactional analysis, but it did not stand alone. After this first pass, I conducted iterative qualitative coding to identify recurring forms of agency enactment and boundary work across cases. Coding focused on three analytic categories: epistemic posture, cognitive intent, and evaluative action. Epistemic posture described whether the student approached AI output deferentially, dialogically, or skeptically. Cognitive intent was captured as exploratory, clarifying, transformational, or metacognitive. Evaluative action captured what the student did with the output: accept, correct, refuse,

or re-author. These categories were used to identify patterned interactional differences across Trios and to distinguish sustained evaluative engagement from moments of delegation or offloading.

Cross-Case Analysis

Cross-case comparison followed within-case coding. Here, I compared recurring patterns across all 23 cases to determine whether similar interactional forms clustered into broader profiles. Memo writing, repeated close reading of logs, and comparison across final artifacts and reflections helped refine these patterns into the two analytic profiles presented in the findings: the Orchestrator and the Outsourcer. These profiles are not participant types or personality labels. They are analytic abstractions that describe contrasting ways of organizing the evaluative loop traced through the Trio.

Trustworthiness and Analytic Limits

Trustworthiness was supported through triangulation across interaction logs, participant reflections, and researcher field notes (Deggs & Hernandez, 2018). This process included constant comparison, analytic memoing, and transparency about coding decisions. To trace the transition from AI contribution to human-authored meaning, analysis also attended to artifacts from the Trio's evaluative loop, including draft iterations and documented instances of refusal or re-prompting, which provided concrete evidence of interpretive control across iterative cycles of collaboration. Interpretive claims were advanced only when interactional patterns recurred across multiple Trios and data sources.

Because the analysis is grounded in sociocultural theory, agency, boundary work, and ethical reasoning are treated here not as internal traits, but as interactionally visible practices that emerge through patterned action over time (Clarke, 2022). Although digital trace data provide fine-grained access to interactional structure, they cannot independently capture participants' intentions or affective states (Keusch & Kreuter, 2021). Accordingly, interpretations are constrained to documented patterns of action over time. The analytic goal is not to evaluate individual competence or generalize behavior, but to document how ethical and epistemic responsibility is negotiated in practice under conditions of institutional and technological ambiguity.

Findings

This study examines how pre-service educators in a literacy course decide what work remains human when generative AI is present, how those decisions are enacted in interaction, and why they are ethically and professionally consequential. Rather than evaluating whether students used AI appropriately, the analysis traces how boundaries around acceptable AI use were constructed, challenged, and revised through moment-by-moment interaction during a semester-long literacy ethnography.

Distinct interactional patterns emerged in how authority, epistemic responsibility, and authorship were distributed between human and machine. These patterns were not visible in the final artifacts

alone but became evident through interactional traces such as prompts, corrections, refusals, and reflective commentary. Across the dataset, similar configurations of interactional organization recurred, giving rise to two contrasting interaction profiles. These profiles did not reflect fixed dispositions or skill levels, but rather different ways of organizing the evaluative loop captured in the Trio between the student and the AI system under conditions of institutional ambiguity (Jocius et al., 2023b).

One profile was characterized by active negotiations of boundary work and sustained human oversight, while the other involved early delegation of authority and minimal intervention once AI output appeared sufficient. Cross-case analysis revealed that these differences hinged not on technical proficiency, but on how learners managed friction, legitimacy, and responsibility within the interaction order.

Together, these findings demonstrate that agency in AI-mediated literacy work is enacted through interaction rather than tool choice, and that ethical engagement depends on how learners negotiate boundary work, refusal, and accountability as meaning-making unfolds.

Two Distinct Interaction Profiles

To make these interactional differences visible, the findings are organized around two analytically distinct interaction profiles that emerged through cross-case comparison. In this study, an interaction profile refers to a recurring pattern in how responsibility, authority, and agency are organized across human–AI interaction, inferred from sustained interactional traces rather than from individual characteristics or isolated actions. These profiles, the Orchestrator and the Outsourcer, are not typologies of individuals, nor do they represent levels of competence or ethical disposition. Instead, they function as analytic lenses for examining how PSTs structured the evaluative loop between themselves and the AI system.

Each profile traces a patterned configuration of interaction order, boundary work, and epistemic responsibility observed across multiple participants over time. These profiles did not map neatly onto every participant, but both patterns recurred across the dataset. While each profile is illustrated through a focal participant for analytic clarity, the emphasis remains on recurring interactional organization rather than on the individual as a case. The findings below trace each profile using Nexus Analysis, with particular attention to interaction order, observable boundary work, and the dispositions and prior habits that appeared to shape participation in the evaluative loop traced through the Trio.

Profile 1: The Orchestrator - Sustained Agency Through Active Boundary Work

The Orchestrator profile reflects a pattern of AI-mediated literacy work in which responsibility for meaning-making remains firmly with the human actor. Across interactions, students exhibiting this profile treated the AI system as a subordinate analytic resource rather than an authoritative voice,

enacting agency through sustained boundary work. Agency was expressed not through efficiency or automation, but through ongoing oversight, constraint, and intervention within the interaction order.

Interaction Order: AI as Junior Analyst. In this profile, the interaction order positioned the AI system as a junior analyst rather than a co-author or expert. JSON interaction logs revealed repeated short-latency cycles in which AI output was immediately evaluated, redirected, or constrained. Rather than accepting initial responses, the student engaged in iterative re-prompting to demand greater specificity, contextual nuance, and closer alignment with epistemic goals.

These stop-and-fix patterns indicate a repair-oriented dialogue. Authority remained with the student, who consistently acted as senior author, using the AI system to surface patterns or organize material while retaining responsibility for interpretation and framing.

Historical Body: Critical and Resistant Orientations. This interactional stance was reflected in prompt phrasing and task framing. Prompts were high-constraint and persona-driven, explicitly instructing the AI system to prioritize political resistance, emotional context, or interpretive nuance while suppressing generic historical summaries. Neutrality was treated not as a default to be accepted, but as a stance to be interrogated.

Across interactions, the student demonstrated discomfort with conceptually flattened representations. This resistance suggested a disposition toward interpretive depth and skepticism of ostensibly objective summaries.

Boundary Work in the Evaluative Loop. This negotiation was most clearly visible in the evaluative loop captured through the Trio (student prompt → AI response → student reflection). In one representative sequence, the AI system generated a source-grounded summary using neutral academic language. The student appended a reflection explicitly rejecting this framing, noting that it erased urgency and political stakes present in the original sources.

Rather than discarding the AI system's output entirely, the student retained its organizational structure while replacing descriptors, examples, and interpretive framing with their own language. One reflective note makes the stance especially clear: the student worried that the AI system would only reflect the resources already provided, said they needed to "lock in on only one part," and admitted, "I worry about this notebook and how I am choosing to use it." These statements matter because they show the student not simply resisting AI, but actively narrowing scope and testing whether the tool's synthesis was too broad or insufficiently aligned with the project's emerging purpose.

Later notes from the same case show that this hesitation did not lead to abandonment but to a more selective interpretive use. The student's synthesis moved toward a more theory-linked analysis in which some content "meets people where they are" while other content "challenges them to grow and change." This movement from skepticism to narrowing to recontextualized synthesis illustrates the

correction loop that characterized the Orchestrator pattern. In one reflection, the student also explicitly noted the move toward revision and voice: when the output was superficial, they revised the prompt to be more specific and rewrote the AI-generated text in their own voice. AI support was used, but only after the student reasserted control over scope, framing, and interpretive emphasis.

Analytic Summary. The Orchestrator pattern was characterized by sustained epistemic control. The AI system was used to support organization and synthesis, but never to define stance or meaning. Agency was enacted through effort, friction, and deliberate constraint rather than speed or delegation. Ethical engagement in this profile was inseparable from the ongoing boundary work and negotiations within the interaction order.

Profile 2: The Outsourcer - Attenuated Agency Through Delegation

The Outsourcer profile reflects a pattern of AI-mediated literacy work in which responsibility for meaning-making is delegated early and remains largely with the AI system. Across interactions, students exhibiting this profile positioned the AI system as a primary organizer and epistemic authority rather than as a subordinate analytic resource. Agency in this profile was enacted through efficiency and task completion rather than through sustained boundary work or interpretive intervention.

Interaction Order: AI as Structural Authority. In this profile, the interaction order positioned the AI system as a structural authority that set the agenda for the work. JSON interaction logs revealed low-latency integration, with extended gaps between AI output and human intervention and minimal evidence of repair or re-prompting. Initial AI responses frequently became the project's structural backbone, shaping organization, genre, and scope with little subsequent modification.

Rather than functioning as a dialogic partnership, interaction took the form of reference-oriented delegation. The AI system's grounded output was treated as sufficient for task completion, and authority over meaning-making was implicitly ceded once responses aligned with assignment expectations.

Historical Body: Efficiency-Oriented Orientations. This interactional stance was reflected in prompt phrasing and task framing. Prompts were low-constraint and generic (e.g., "Write an introduction based on the uploaded files"), signaling deference to the AI system's default interpretive frames. Across interactions, the student routinely accepted default genres generated by the AI system, including study guides, quizzes, glossaries, and summaries.

These patterned actions are consistent with habits often reinforced in school settings that prioritize efficiency, compliance, and genre conformity. Grounded, source-linked output was treated as inherently adequate, reducing the perceived need for interrogation, reinterpretation, or stance-taking.

Absence of Boundary Work in the Evaluative Loop. In contrast to the Orchestrator profile, boundary work was largely absent from the evaluative loop. Trio analysis showed that reflections, when present, were brief and affirmative, such as “This looks good and saved me a lot of time.” There was little evidence of correction, refusal, or revoicing.

As a result, AI-generated language and structure were frequently lifted into final artifacts with minimal transformation. This pattern was analytically coded as cognitive offloading, in which responsibility for interpretation and framing remained with the AI system rather than being reclaimed by the student.

One focal case centered on a municipal fire department. In this notebook, the dominant pattern was rapid uptake of broad, organized synthesis rather than visible struggle with scope or framing. NotebookLM generated structured analyses under headings such as “How Literacy Practices Differ From or Resemble Other Communities,” “External Factors Influencing Literacy Practices,” and “Power Dynamics Affecting Access to Literacy Practices.” The resulting prose quickly moved into polished claims, suggesting that the notebook’s organizing structure had become the backbone of the analysis, with minimal visible resistance. Compared with the statewide advocacy case, there was little evidence of narrowing, disagreement, or explicit reframing of the AI system's categories. In analytic terms, this is where delegation becomes visible. The student remains present in the process, but much of the organizational and interpretive labor is ceded to the AI system.

Analytic Summary. Across interactions, the Outsourcer profile consistently deferred epistemic authority to the AI system. While this approach increased efficiency and reduced cognitive friction, it also limited opportunities for interpretive depth and ethical deliberation. Final artifacts were conceptually safe but shallow, reflecting the AI system’s neutral framing more than the student’s own epistemic stance. In this profile, the absence of boundary work rendered agency largely invisible within the interaction order.

Boundary Work as the Mechanism of Agency

To synthesize the two interaction profiles, I conducted a cross-case comparison focused on how evaluative boundary work operated across the interaction loop. The defining distinction was not whether AI was used, but whether participants sustained evaluative boundary work while interacting with it. See Table 2 for a summary of these differences.

Table 2. Cross-Profile Comparison: Boundary Work as the Mechanism of Agency

DIMENSION	ORCHESTRATOR (ACTIVE BOUNDARY WORK)	OUTSOURCER (ATTENUATED BOUNDARY WORK)
Interaction order	Human-led repair dialogue	AI-led delegation
Prompt structure	High-constraint, iterative	Low-constraint, generic
Response latency	Short, evaluative pauses	Long, deferential waits
Integration strategy	Revoicing, revising, synthesizing	Lifting and adopting output
Epistemic stance	Provisional treatment of AI text	Acceptance of AI as authority
Reflection	Explicit articulation of boundaries and limits	Minimal or affirming commentary

Across profiles, agency emerged not from prompt generation alone but from acts of interruption, correction, and refusal. Boundary work operated as the mechanism through which epistemic authority was either asserted or ceded within the interaction order. Looked at this way, the contrast between Orchestrator and Outsourcer was less about whether students used the tool and more about how long they remained actively accountable for what the tool produced.

Discussion

This study examined how PSTs decided what work remained human when generative AI was built into an authentic literacy task. Across the dataset, the central difference was not whether students used AI, but how they regulated its role in the evaluative loop. Rather than centering the tool itself, the findings foreground boundary work: the interactional process through which participants accepted, constrained, revised, or refused AI contributions. Ethical AI literacy thus appeared less as a matter of compliance than as a situated practice of judgment.

What is often framed as “AI use” was, here, more accurately observed as a negotiation of legitimacy, responsibility, and epistemic authority. These negotiations were largely invisible in final artifacts and became legible only through process-level analysis of prompts, revisions, uptake patterns, and reflections. One contribution of this study, then, is methodological and conceptual: it shows the value of tracing interaction rather than inferring judgment from polished final products alone.

Although grounded in a single course, the analytic framework developed here, combining Nexus Analysis with digital trace evidence, may be useful in other settings where interactional processes can be examined alongside final products. The discussion that follows focuses on three implications of the findings: a shift from verification to valuation, treating boundary work as an ethical literacy practice, and designing teacher education experiences that make evaluative judgment visible.

Shifting From Verification to Valuation

A central implication of this study is the need to move beyond verification-oriented models of digital literacy toward valuation-oriented practices. In earlier digital contexts, literacy instruction emphasized

whether information was accurate, credible, or relevant. While these questions remain necessary, they are insufficient in generative AI environments, particularly those using RAG systems that often produce technically accurate but conceptually flattened output.

The contrast between the Orchestrator and the Outsourcer illustrates this shift. Both students worked with source-grounded AI outputs. However, only the Orchestrator engaged in valuation, rejecting responses not because they were incorrect but because they lacked nuance, voice, or epistemic stance. The Outsourcer’s evaluative process often concluded once outputs appeared acceptable and efficient.

Across cases, human agency did not hinge on producing correct outputs but on judging whether AI contributions were sufficient for meaning-making. Moments of agency clustered around a recurring evaluative routine, which I conceptualize as a four-step Agency Check (see Table 3).

Table 3. The Agency Check: Evaluative Moves Through Which Preservice Teachers Enacted Epistemic Authority

STEP	GUIDING QUESTION	WHAT IS BEING EVALUATED	BOUNDARY WORK ENACTED
Credible	Is the information accurate and source-grounded?	Factual reliability, traceability, alignment with trusted knowledge	Rejecting hallucinations; reasserting human verification
Relevant	Does it address the specific task or inquiry?	Fit to purpose and contextual alignment	Steering AI away from generic output
Acceptable	Is it cohesive and serviceable as a draft?	Usability, coherence, draftability	Treating AI as a provisional collaborator rather than an authority
Nuanced	Does it reflect a particular epistemic stance, contextual sensitivity, and human voice?	Voice, stance, judgment, ethical positioning	Reclaiming authorship and interpretive agency

Together, these checks illustrate how participants treated AI output as provisional material to be evaluated, revised, or refused rather than accepted as authoritative. This shows the enactment of epistemic boundaries in practice.

Enacting Boundary Work as Ethical Practice

The interaction patterns suggest that ethical AI literacy, in this context, was not primarily demonstrated through compliance or disclosure, but through boundary-setting actions enacted within the evaluative loop. In the Orchestrator profile, agency was expressed through constraints, corrections, and refusals: actions that required effort, pause, and interpretive judgment. In contrast, delegation without evaluative interruption resulted in cognitive offloading, even when AI use was transparent and technically appropriate.

These findings reframe HITL not as a technical safeguard, but as a literacy practice. AI functioned as a cognitive amplifier only when humans actively regulated the loop by curating inputs, interrogating outputs, and re-authoring meaning. Where such regulation was absent, AI shifted from amplifier to structural authority. Boundary work in this study closely mirrors longstanding literacy practices, including evaluating sources, synthesizing across texts, shaping voice, and making interpretive claims. What changes in AI-mediated contexts is not the nature of literacy, but the visibility and urgency of boundary decisions.

Navigating Moral Ambiguity and the “Cheating Tension”

An emergent finding of this study was the persistent moral ambiguity surrounding AI use, even within a pedagogical context that explicitly sanctioned it. For some participants, especially those who engaged most actively in evaluative boundary-setting, AI support appeared to intensify questions about authorship, legitimacy, and professional integrity. In this dataset, that tension surfaced not as policy confusion alone, but as discomfort with how much interpretive labor could be shared without feeling that something important had been ceded.

This tension was not distributed evenly. It surfaced most strongly among students who engaged in the highest levels of boundary work. As these participants moved beyond functional use toward genuine cognitive partnering, their concerns about authorship and legitimacy intensified. This pattern was visible in extended response latencies, recursive prompting, and reflective commentary that treated AI assistance as both useful and suspect. Rather than reading this discomfort as a contradiction, the findings suggest it may be a sign that the evaluative stakes of collaboration had become newly visible.

In teacher education, this finding underscores the need to normalize ethical negotiation rather than attempt to resolve it through rigid rules alone. One instructional goal, then, is to help PSTs articulate their own internal red lines by distinguishing between the selective offloading of mechanical tasks and the non-delegable nature of interpretive, relational, and community-representational work.

Designing Loops Worth Living In

The findings point toward a reframing of AI integration in teacher education as a matter of interaction design rather than tool adoption (Jocius et al., 2023a). Preparing PSTs to work ethically with AI requires designing loops worth living in: learning environments that make evaluation, refusal, and reflection visible and expected.

Such designs emphasize pause-and-probe moments, structured reflection on the division of labor, and assessment practices that value process alongside product. When students are supported in embracing cognitive friction, they develop professional habits essential for ethical responsibility in K–12 classrooms. Ultimately, this study suggests that the central question of AI literacy is not whether students use AI, but how they learn to decide when and how not to delegate judgment.

Conclusion

Future educators are entering a professional landscape in which generative AI is already embedded in everyday writing, planning, and meaning-making. This study suggests that the quality of their learning and the ethical integrity of their work do not hinge on whether AI is used, but on how the boundaries between human judgment and machine contribution are actively constructed.

Across cases, differences in outcomes were driven by interactional practice rather than technical skill. When students engaged in active resistance, correction, and selective integration, AI served as a cognitive amplifier. When AI output was accepted without evaluative interruption, cognitive labor was offloaded, resulting in flattened representations and diminished epistemic responsibility. These differences were visible only through process-level analysis of interaction traces, underscoring the limits of evaluating AI-mediated work based solely on final artifacts.

For teacher preparation, these findings suggest that AI literacy must be reframed as boundary work. Preparing literacy educators is not primarily about teaching prompt strategies or enforcing compliance rules, but about cultivating the professional judgment required to decide what work remains human. Ethical AI integration depends less on smarter tools than on educators who can design and inhabit loops worth living in.

ORCID iDs

Ian W. O'Byrne — <https://orcid.org/0000-0003-0307-9538>

Funding

The author received no financial support for the research, authorship, and/or publication of this article.

Declaration of Conflicting Interests

The author declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

References

- Ackerman, L. (2025). The workflow as medium: A framework for navigating human-AI co-creation. arXiv preprint arXiv:2511.18182. <https://doi.org/10.48550/arXiv.2511.18182>
- Adillón, M. V., Bellón, E. M. E., Cantero, J. M. M., & Forgas, R. C. (2024). Academic integrity in pre-service teacher education: a review of the literature. *Práxis Educativa*, 19, 1–21. <https://doi.org/10.5212/PraxEduc.v.19.23273.026>

- Bandi, A., Adapa, P. V. S. R., & Kuchi, Y. E. V. P. K. (2023). The power of generative AI: A review of requirements, models, input–output formats, evaluation metrics, and challenges. *Future Internet*, 15(8), 260. <https://doi.org/10.3390/fi15080260>
- 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>
- Bozkurt, A., Xiao, J., Farrow, R., Bai, J. Y., Nerantzi, C., Moore, S., & Asino, T. I. (Eds). (2024). The manifesto for teaching and learning in a time of generative AI: A critical collective stance to better navigate the future. *Open Praxis*, 16(4), 487–513. <https://doi.org/10.55982/openpraxis.16.4.777>
- Burnett, C. (2011). Shifting and multiple spaces in classrooms: An Argument for investigating learners’ boundary-making around digital networked texts. *Journal of Literacy & Technology*, 12(3), 2–23.
- 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>
- Chiu, T. K., Ahmad, Z., Ismailov, M., & Sanusi, I. T. (2024). What are artificial intelligence literacy and competency? A comprehensive framework to support them. *Computers and Education Open*, 6, 100171. <https://doi.org/10.1016/j.caeo.2024.100171>
- Clarke, A. E. (2022). Situating grounded theory and situational analysis in interpretive qualitative inquiry. In *Situational analysis in practice* (pp. 47–96). Routledge. <https://doi.org/10.4324/9781003035923-6>
- Code, J. (2025). The entangled learner: Critical agency for the postdigital era. *Postdigital Science and Education*, 7(2), 1–23. <https://doi.org/10.1007/s42438-025-00544-1>
- Deggs, D., & Hernandez, F. (2018). Enhancing the value of qualitative field notes through purposeful reflection. *The Qualitative Report*, 23(10), 2552–2560. <https://doi.org/10.46743/2160-3715/2018.3569>
- Hancock, J. T., Naaman, M., & Levy, K. (2020). AI-mediated communication: Definition, research agenda, and ethical considerations. *Journal of Computer-Mediated Communication*, 25(1), 89–100. <https://doi.org/10.1093/jcmc/zmz022>
- Higgs, J. M., & 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>
- Jocius, R., Albert, J., Joshi, D., Blanton, M., & O’Byrne, W. I. (2023a). Creating sustainable computational thinking infusion: An analysis of teacher-led practitioner inquiry projects. *Journal of Technology and Teacher Education*, 31(4), 492–522. <https://doi.org/10.70725/583931ipnfup>
- Jocius, R., Joswick, C., Albert, J., Joshi, D., & Blanton, M. (2023b). Towards pedagogical content knowledge learning trajectories: Tracing elementary teachers’ infusion of computational thinking. *Professional Development in Education*, 51(1), 815–834. <https://doi.org/10.1080/19415257.2023.2228813>
- Keusch, F., & Kreuter, F. (2021). Digital trace data: Modes of data collection, applications, and errors at a glance. In *Handbook of computational social science*, volume 1 (pp. 100–118). Routledge. <https://doi.org/10.4324/9781003024583-8>
- Kim, Y. J., Zhang, C. X., & Fang, C. (2025). Navigating the human-AI divide: Boundary work in the age of generative AI. *Computers in Human Behavior: Artificial Humans*, 6, 1–11. <https://doi.org/10.1016/j.chbah.2025.100214>
- Leander, K. M., & Burriss, S. K. (2020). Critical literacy for a posthuman world: When people read, and become, with machines. *British Journal of Educational Technology*, 51(4), 1262–1276. <https://doi.org/10.1111/bjet.12924>

- Li, T., Krajcik, J. S., & Spiro, R. (2025). The transformative collaboration of human intelligence and artificial intelligence in designing knowledge-in-use science assessment for learning. *Journal of Science Education and Technology*, 34(5), 1–27. <https://doi.org/10.1007/s10956-025-10275-4>
- Lindgren, H. (2025). Emerging roles and relationships among humans and interactive AI systems. *International Journal of Human–Computer Interaction*, 41(17), 10595–10617. <https://doi.org/10.1080/10447318.2024.2435693>
- Long, D., & Magerko, B. (2020, April). What is AI literacy? Competencies and design considerations. In *Proceedings of the 2020 CHI conference on human factors in computing systems* (pp. 1–16). Association for Computing Machinery. <https://doi.org/10.1145/3313831.3376727>
- McNealis, R. (2025). Shame in the machine: Affective accountability and the ethics of AI. *AI & SOCIETY*, 40(3), 1–11. <https://doi.org/10.1007/s00146-025-02472-x>
- Mollick, E., & Mollick, L. (2023). Assigning AI: Seven approaches for students, with prompts. arXiv preprint arXiv:2306.10052. <https://doi.org/10.2139/ssrn.4475995>
- Naeem, M., & Thomas, L. (2025). Case study research and artificial intelligence: A step-by-step process to using ChatGPT in the development of case studies in qualitative research. *International Journal of Qualitative Methods*, 24, 1–21. <https://doi.org/10.1177/16094069251371478>
- Natarajan, S., Mathur, S., Sidheekh, S., Stammer, W., & Kersting, K. (2025). Human-in-the-loop or AI-in-the-loop? Automate or collaborate? *Proceedings of the AAAI Conference on Artificial Intelligence*, 39(27), 28594–28600. <https://doi.org/10.1609/aaai.v39i27.35083>
- Nguyen, D., & Hekman, E. (2024). The news framing of artificial intelligence: A critical exploration of how media discourses make sense of automation. *AI & Society*, 39(2), 437–451. <https://doi.org/10.1007/s00146-022-01511-1>
- Nopas, D. S. (2025). Algorithmic learning or learner autonomy? Rethinking AI’s role in digital education. *Qualitative Research Journal*, 25(1), 45–62. <https://doi.org/10.1108/QRJ-11-2024-0282>
- Norouzi, B., Kinnula, M., & Iivari, N. (2019). Interaction order and historical body shaping children’s making projects —A literature review. *Multimodal Technologies and Interaction*, 3(4), 71. <https://doi.org/10.3390/mti3040071>
- O’Byrne, W. I., Crandall, B. R., Price-Dennis, D., Witte, S., Goering, C. Z., & Dail, J. S. (2024). When you know better, do better: Developing antiracist, digitally literate educators through critical media literacy. In M. E. Barnes & R. Marlatt (Eds.), *Teaching for equity, justice, and antiracism with digital literacy practices: Knowledge, tools, and strategies for the ELA classroom* (pp. 69–86). Routledge. <https://doi.org/10.4324/9781003430209-7>
- Palmquist, A., Sigurdardottir, H. D. I., & Myhre, H. (2025). Exploring interfaces and implications for integrating social-emotional competencies into AI literacy for education: A narrative review. *Journal of Computers in Education*, 12(4), 1–37. <https://doi.org/10.1007/s40692-025-00354-1>
- Raffaghelli, J. E., Vartiainen, H., Bower, M., Ronci, M., Shelton, C., MacCallum, K., Lee, J., Webb, M., Cthouki, Y., & Smith, D. (2026). Situated AI ethics: A cultural-historical and ecological framework for education. *Computers and Education Open*, 11, 1–8. <https://doi.org/10.1016/j.caeo.2026.100368>
- Scalzer, C., Pokhrel, S., Hunt, S., & Nelson, G. L. (2025). Designing for learning with generative AI is a wicked problem: An illustrative longitudinal qualitative case series. arXiv preprint arXiv:2507.17230. <https://doi.org/10.48550/arXiv.2507.17230>
- Schiera, A. J. (2019). Justice, practice and the ‘real world’: Pre-service teachers’ critically conscious visions for teaching amid the complexities and challenges of learning to teach. *International Journal of Qualitative Studies in Education*, 32(7), 929–946. <https://doi.org/10.1080/09518398.2019.1609125>

- Schuetz, S., & Venkatesh, V. (2020). The rise of human machines: How cognitive computing systems challenge assumptions of user-system interaction. *Journal of the Association for Information Systems*, 21(2), 460–482. <https://doi.org/10.17705/1jais.00608>
- Scollon, R., & Scollon, S. W. (2004). *Nexus analysis: Discourse and the emerging Internet*. Routledge. <https://doi.org/10.4324/9780203694343>
- Siemens, G., Marmolejo-Ramos, F., Gabriel, F., Medeiros, K., Marrone, R., Joksimovic, S., & De Laat, M. (2022). Human and artificial cognition. *Computers and Education: Artificial Intelligence*, 3, 100107. <https://doi.org/10.1016/j.caeai.2022.100107>
- Velander, J., Otero, N., & Milrad, M. (2024). What is critical (about) AI literacy? Exploring conceptualizations present in AI literacy discourse. In *Framing futures in postdigital education: Critical concepts for data-driven practices* (pp. 139–160). Springer Nature Switzerland. https://doi.org/10.1007/978-3-031-58622-4_8
- Wang, Z., & Wang, C. (2025). A posthumanist approach to AI literacy. *Computers and Composition*, 76, Article 102933. <https://doi.org/10.1016/j.compcom.2025.102933>
- Wertsch, J. V. (1991). *Voices of the mind: Sociocultural approach to mediated action*. Harvard University Press.
- Wise, J. B., & O’Byrne, W. I. (2016). Social scholars: Educators’ digital identity construction in open, online learning environments. *Literacy Research: Theory, Method, and Practice*, 65(1), 398–414. <https://doi.org/10.1177/2381336915617607>
- Wohlwend, K. (2020). *Literacies that move and matter: Nexus analysis for contemporary childhoods*. Routledge. <https://doi.org/10.4324/9780429265754>

Author biography

Ian W. O’Byrne is an associate professor of literacy education at the College of Charleston in South Carolina. His research focuses on the dispositions and literacy practices of individuals as they read, write, and communicate in online and/or hybrid spaces. His work can be found on his website (<https://wiobyne.com/>) or in his weekly newsletter (<https://digitallyliterate.net/>).