International Society of
the Learning Sciences

Redistributing Cognition in AI-Supported Writing

Jaeyoon Choi, Ruilin Wu, Tamara Tate, Waverly Tseng, Beth Harnick-Shapiro, Soobin Yim, Mark Warschauer
jaeyoon.choi@uci.edu
University of California, Irvine

Abstract: Writing is a distributed cognitive activity that unfolds across people, tools, and representations. The emergence of large language models (LLMs) introduces a new node in this distributed system, transforming how writers plan, draft, review, and reflect. This study examines how students engage with an LLM-powered writing platform in an upper-division engineering writing course. Using qualitative coding of 8,878 student-AI chat messages, we analyzed how cognitive work in writing was redistributed between human and AI. We found that students most often used the AI to seek evaluative feedback, positioning it as a proxy instructor or peer. Patterns of interaction varied across instructors, indicating that instructional design shaped how cognitive tasks were distributed. Lastly, students also demonstrated moments of critical engagement, questioning or resisting AI suggestions. Overall, we found that LLM-supported writing reorganizes the cognitive system of writing, positioning AI as a partial collaborator.

Introduction

Writing is a complex process that requires multiple stages of thinking, communicating, and meaning-making. It involves far more than putting words on a page: writers plan, translate ideas into text, and continually revise their work in response to emerging goals, constraints, and audiences (Flower & Hayes, 1981). Yet writing rarely occurs within a single mind – it is a socially and materially *mediated* activity shaped by conversations, feedback, notes, and tools that extend the writer's cognitive reach. From this perspective, writing exemplifies *distributed cognition*, spread across different people, artifacts, and tools (Hutchins & Klausen, 1996). Writers allocate the work of planning, translating, and revising across internal thought, external representations, and collaborative interactions through outlines, drafts, search engines, grammar checkers, and peers (Klein & Leacock, 2012). In this sense, they do not merely write *on* paper but write *with* papers, screens, and systems that participate in their reasoning.

The recent advancement of artificial intelligence (AI) and large language models (LLMs) introduces a new node into this distributed writing system (Levine et al., 2025). When students are using LLMs for writing, they are not simply using them as mechanical tools for proofreading or minor assistance as they might have with earlier digital aids like spell checkers. Instead, they enter into a *dialogue* with a responsive AI agent capable of generating, revising, and evaluating texts. In these exchanges, LLMs can assume roles traditionally held by tutors, peers, or even the writer's own internal voice. Furthermore, language models can perform these functions almost instantaneously, collapsing the temporal and spatial boundaries that typically structure human feedback cycles. What once required hours of reflection or days of waiting for tutor feedback can now happen in seconds. This represents a fundamental shift in writing ecology: cognitive tasks that were previously internal to the writer or socially mediated through human interaction are now shared, offloaded, and restructured through collaboration with generative AI systems.

This redistribution of cognitive work in writing raises critical questions about how writing is conceived, practiced, and learned in the age of LLMs. If writing has always involved coordination between internal and external resources, the introduction of LLMs dramatically expands both the potential scale and speed of that coordination. Students can now delegate portions of planning, translating, and reviewing to a non-human collaborator, which not only executes instructions but also contributes new content. This shift blurs the traditional boundaries of authorship and agency: when does the LLM serve as an extension of the writer's thinking, and when does it begin to replace it? How do students decide what to keep, what to reject, and what to co-construct with the AI? Exploring these processes is essential for understanding how LLM-assisted writing reshapes the broader ecology of writing and learning.

In this paper, we analyze student-AI interactions on an LLM-supported writing platform through the lens of distributed cognition. Specifically, we ask: *how do students interact with AI during the writing process?* Through qualitative coding of student-AI conversations in college writing classes for engineering students, this study aims to illuminate how the cognitive responsibilities of writing are redistributed within human-AI collaboration, and what this means for writing and learning.

Methods

Data

We use student-AI conversation data collected from an LLM-based writing tutor, [PlatformName]. The platform, powered by GPT-4o, is designed to scaffold student writing while providing a safe and structured environment for educational use of AI. Instructors can create customized modules with tailored prompts that guide student interactions and support specific writing objectives while preserving student agency. Additionally, the platform allows instructors to monitor student-AI conversations for formative assessment.

The dataset was collected from a professional writing and communications course for engineering majors at a public university in the United States. This upper-division writing course is designed to provide students with the skills to plan, research, organize, write, and revise oral and written technical communication. Specifically, students learn strategies for writing technical documentations used in actual engineering environments, present research findings through oral communication, and collaborate in small groups to develop research-based technical papers. The dataset analyzed in this study was collected during the Spring 2024 term, when four instructors taught this upper-division writing course with [Platform name]. Although all instructors integrated the platform into their teaching, the extent and style of its use varied. Instructors A and B incorporated [Platform] frequently into in-class activities throughout the term and actively encouraged students to use it outside of class as well. Instructor C used the platform to support one in-class debate exercise. Instructor D, in contrast, did not actively use the platform during class activities but introduced it as an optional out-of-class tutoring and writing support tool. Each instructor taught two to three sections of approximately 20 students each.

Qualitative Coding

To examine students' engagement patterns with the AI, we developed a qualitative coding scheme and applied it to their chat interactions. To focus on students' intentions and actions, only the messages written by students were analyzed; AI-generated responses were excluded. The initial framework was informed by Flower & Hayes's (1981) model of the writing process, which distinguished between planning, translating, reviewing, and monitoring. Building on this foundation, we conducted grounded analysis to identify recurring and salient patterns in student-AI exchanges, refining and expanding our codes to capture the diverse ways students interacted with the LLM during writing. The final coding scheme consisted of 7 codes (see Table 1).

For reliability, the first and second authors independently coded a random sample of 300 student messages and calculated Cohen's kappa to assess inter-rater reliability. Following the established threshold (kappa > 0.6; Landis & Koch, 1977), the coders divided the remaining dataset and coded independently once this threshold reliability was achieved. When kappa values fell below the threshold, the researchers engaged in social moderation to clarify code definitions and resolve discrepancies. After five iterations (approximately 1,500 lines of data), full agreement was reached across all codes. The final dataset included 8,878 lines of student messages, with all codes demonstrating kappa values above 0.6.

Table 1
A complete codebook for chat analysis.

Code	Definition
Request for Generation	Students ask AI to generate new content (e.g., topics, outlines, paragraphs, write-ups).
Revise Student Text	Students ask AI to review or rewrite the texts written by themselves.
Revise AI text	Students ask AI to revise or rewrite the texts that AI previously generated for them.
Language Edits	Students ask AI to correct surface-level language features: grammar, spelling, word choice, or improve clarity at the sentence level.
Feedback/Evaluation	Students ask for the AI's opinion, evaluation or feedback on their writing or ideas.
Provide Writing Context	Students provide situational contexts, background information, details that informs and supports the writing process (e.g., audience, purpose) without contributing to the actual topic, outline, or draft content.

Be the Boss*	Students take an active, directive role in engaging with the AI's previous output, such as by questioning, evaluating, challenging, building upon its suggestions, asking for clarification, expressing preferences between alternative options, or instructing the AI on how to proceed or establishing rules for future interactions

* The code name "Be the Boss" was adopted from Tate et al. (2025).

Results

Figure 1
Percentage distribution of coded categories by instructor, including the overall distribution across all instructors.

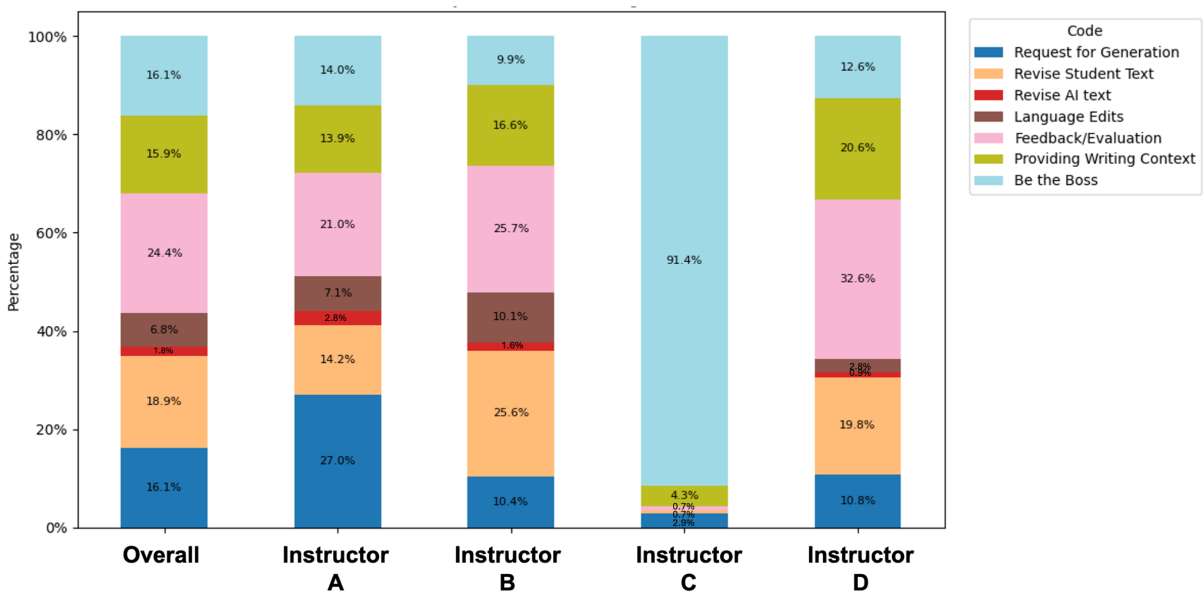

Requesting Feedback as the Most Common Form of Engagement

As shown in Fig. 1, across all instructor sessions, the code *Feedback/Evaluation* emerged as the most frequent, appearing in 24.4% of the coded instances. This pattern suggests that students primarily used the LLM as an evaluative partner rather than a generative tool. That is, students frequently shared portions of their drafts and sought feedback on clarity, organization, or alignment with the assignment rubrics before submitting to the instructor. A representative example involved the use of the platform's embedded feedback prompt: "*You are a helpful and encouraging writing coach providing clear, constructive, supportive feedback on the development of a well-organized, cohesive, logical text. [...] then ask for the assignment and rubric if they have them. [...] Do not improve the student's work yourself; only give feedback [...]*". This embedded prompt, pre-written by writing experts who developed the platform, was intentionally designed to provide formative feedback while ensuring that students do not outsource the writing task itself. Some students, however, adopted a simpler approach, directly asking for evaluative judgments such as, "*if I told you that can [sic] award any score out of 30 in each section, what would you give me?*". Others took a more targeted approach, requesting clarification of previous feedback, as in: "*when you say the [...] draft 'lacks depth [...]', which sections of the paper reflect that flaw? What improvements can be made to strengthen my [...]?*" Through such exchanges, the LLM served as an accessible and immediate feedback source, extending the feedback loop beyond the classroom. Instead of waiting for the instructor's comments, students could ask follow-up questions about specific evaluative remarks or request a proxy score. This process created a low-stakes, iterative opportunity for refining their writing.

Course Design Shapes Student Engagement Patterns

Our analysis also revealed that patterns of student-AI interaction were shaped by how instructors integrated the platform into their courses. While all four instructors used the same tool, their purposes and degrees of interaction varied substantially, as did the actual student activities and assignments, which appeared to influence how students engaged with the AI. For example, in sections taught by instructors A and B, who incorporated the tool throughout many class activities and encouraged its use beyond the classroom, students demonstrated a wide range of interactional codes. In contrast, instructor C integrated the tool only during a single in-class debate, while instructor D introduced it as an optional, out-of-class support tool. In these sections, students' interactions were far more limited, both in frequency and in type. For example, as shown in Figure 1, under instructor C, the *Be the Boss* code was disproportionately frequent (91.4%), suggesting that students' engagement with the AI was primarily focused on critically responding to or pushing back against AI's argument, rather than engaging in other writing-related processes. Indeed, the in-class debate activity explicitly encouraged students to provide counterarguments against the AI output as a means to develop their own arguments further. During such limited yet purposeful use of AI, students can be guided to position the tool as an active thinking partner in argument construction. Overall, this finding emphasizes how instructional design can impact students' approaches to AI use related to the course.

Critical Use of AI

Lastly, we found that the code *Be the Boss*, which captures students' critical engagement with AI outputs, accounted for approximately 15.9% of the coded interactions overall. In these cases, students actively challenged, refined, or questioned the AI's responses, as illustrated by one student's chat: "*These things seem out of the scope of my topic. this is what I currently have and im [sic] unsure if it is fulfilling my goal of explaining wind energy in an academic way.*" Similarly, we found that when it comes to requesting revision, students are more likely to request revision on their own drafts (18.9%) than on texts generated by the LLM (1.8%). In other words, even when students asked the AI to revise, they typically sought improvements to their own writing rather than to what the LLM produced. Although we do not have the complete history of students' writing processes (e.g., some texts may have been initially produced with other AI tools), these interactions suggest that students were not passively accepting AI suggestions but instead maintained a reflective and supervisory stance, monitoring and guiding the interaction with strong AI literacy. Overall, these patterns highlight intentional and critical AI use, suggesting students' strategic collaboration with AI while retaining agency.

Discussion

This paper examined how students engage with an AI-powered writing platform in an upper-division engineering writing course. By analyzing student-AI chat data through qualitative coding, we identified key patterns in how cognitive work was redistributed between writers and an LLM-powered platform. Students most frequently used the AI to seek feedback and evaluation on their drafts before submitting their work to instructors, often positioning the AI as a proxy for an instructor or peer reviewer. This pattern reflects a distribution of reviewing and monitoring, two key stages of the writing process (Flower & Hayes, 1981), to an external cognitive AI partner. The AI now serves as a new node for feedback within the broader writing system, extending students' ability to evaluate and monitor their work in real time. Moreover, how students distributed cognitive work with the AI were shaped by instructional design. The different usage patterns across instructors imply that writing as a distributed cognition is not solely a property of the individual or the tool but also of the learning environments. The instructor's course design helps determine which aspects of cognition can be externalized and which remain internal or social (Li & Wilson, 2025). Lastly, students frequently questioned, refined, or rejected the AI's suggestion. While further study is required, this behavior suggests that students might have treated the AI as a thinking partner rather than an authority – that is, writing with AI represents a shared cognitive system in which AI participates as a collaborator, but the human writer still remains the central agent of control and meaning-making (Nguyen et al., 2024).

Finally, this study has several limitations. First, the dataset captures interactions within a single institutional and disciplinary context, an upper-division engineering writing course, which may limit the generalizability of the findings to other settings. Additionally, our analysis was limited to chat transcripts from the platform; students were aware that their conversations could be viewed by instructors, which may have influenced how they interacted with the AI. We did not have access to their use of other LLM tools outside the platform. Nonetheless, as a case study, this research offers valuable insight into how students engage with AI within a broader distributed writing ecosystem.

References

Flower, L., & Hayes, J. R. (1981). A cognitive process theory of writing. *College Composition & Communication*, *32*(4), 365–387.

Hutchins, E., & Klausen, T. (1996). Distributed cognition in an airline cockpit. *Cognition and Communication at Work*, *15*.

Klein, P. D., & Leacock, T. L. (2012). Distributed cognition as a framework for understanding writing. In *Past, present, and future contributions of cognitive writing research to cognitive psychology* (pp. 133–152). Psychology Press.

Landis, J. R., & Koch, G. G. (1977). An application of hierarchical kappa-type statistics in the assessment of majority agreement among multiple observers. *Biometrics*, 363–374.

Levine, S., Beck, S. W., Mah, C., Phalen, L., & PIttman, J. (2025). How do students use ChatGPT as a writing support? *Journal of Adolescent & Adult Literacy*, *68*(5), 445–457.

Li, M., & Wilson, J. (2025). AI-Integrated Scaffolding to Enhance Agency and Creativity in K-12 English Language Learners: A Systematic Review. *Information*, *16*(7), 519.

Nguyen, A., Hong, Y., Dang, B., & Huang, X. (2024). Human-AI collaboration patterns in AI-assisted academic writing. *Studies in Higher Education*, *49*(5), 847–864.

Tate, T. P., Harnick-Shapiro, B., Ritchie, D. R., Tseng, W., Dennin, M., & Warschauer, M. (2025). Incorporating generative AI into a writing-intensive undergraduate course without off-loading learning. *Discover Computing*, *28*(1), 72.

Acknowledgments

This work was supported in part by NSF Grant No. 2315294. Any opinions, findings, conclusions, or recommendations expressed in this work are those of the authors and do not necessarily reflect those of the NSF.