

INSTRUCTOR-GROUNDED AGENTIC AI FOR CONTINUOUS LEARNING AND CAREER ALIGNMENT

Samuel SETSOFIA*, Carolyn TANG KMET*, Samson QUAYE**

*Northwestern University, Chicago, USA

**Illinois Institute of Technology, Chicago, USA

Samuelsetsofia2025.1@u.north-western.edu

Abstract: The widespread adoption of large language models in higher education has introduced new opportunities for scalable academic support, but existing systems remain largely unaligned with instructional context and pedagogical intent. Current AI tools operate as general-purpose assistants, lacking awareness of course-specific materials, instructor-defined learning objectives, and continuity across student interactions. This limitation contributes to fragmented learning experiences and weak alignment between academic work and career outcomes. This paper introduces Instructor-Grounded Agentic AI, a framework for designing educational AI systems that integrate course-specific knowledge, pedagogical strategies, and longitudinal student interaction. The proposed approach models learning as a context-conditioned and continuous process, combining retrieval-based grounding, structured pedagogical guidance, and agentic reasoning capabilities. A dual-layer architecture is presented, consisting of a course-grounded assistant that extends instructor presence and a longitudinal learning companion that accumulates student work across courses to support career alignment. We formalize the learning interaction as a function of query, instructional context, pedagogical policy, and memory, and define an objective that captures learning depth, continuity, and career relevance. The paper outlines a mixed-methods research design and presents preliminary insights from an initial deployment, demonstrating the feasibility of the approach. The proposed framework advances a shift from general-purpose AI assistance toward instructionally aligned, continuous learning systems, with implications for educational design and workforce readiness.

Keywords: Artificial Intelligence, Agentic AI, educational science, continuous learning, career alignment

1. Introduction

The rapid adoption of large language models in higher education has fundamentally reshaped how students access information, solve problems, and engage with course material. Recent evidence suggests that a majority of university students now rely on general-purpose AI systems for academic support. While these systems provide immediate and often high-quality responses, they operate without awareness of instructor-specific pedagogical intent, course structure, or learning objectives. This creates a critical misalignment between AI-assisted interactions and the goals of formal

education. At the same time, higher education continues to face a longstanding structural challenge: the fragmentation of learning across courses, disciplines, and academic experiences. Students accumulate knowledge in isolated contexts but often struggle to integrate and apply that knowledge in meaningful ways, particularly in professional and career-oriented settings. This gap becomes especially visible during the transition from education to employment, where students must translate academic work into clear demonstrations of skills and competencies. Current AI systems exacerbate both challenges. First, they

generate responses that are contextually generic rather than instructionally grounded, often prioritizing answer delivery over guided reasoning. Second, they lack persistence across time, treating each interaction as independent rather than part of a continuous learning trajectory. As a result, AI-assisted learning today is largely optimized for short-term task completion rather than long-term knowledge development and career alignment.

This paper argues that effective AI in education must move beyond general-purpose assistance toward systems that are explicitly grounded in instructional context and capable of supporting continuous learning. To address this, we introduce Instructor-Grounded Agentic AI, a framework for designing AI systems that integrate instructor-specific knowledge, pedagogical strategies, and longitudinal student interaction. The proposed approach combines three key elements: (i) grounding in course-specific materials through retrieval-based mechanisms, (ii) incorporation of pedagogical strategies such as scaffolding and Socratic guidance, and (iii) agentic capabilities that enable adaptive, multi-step interaction over time.

Central to this framework is a dual-layer architecture. The first layer consists of a course-grounded assistant that extends the instructor’s presence beyond the classroom by providing context-aware academic support. The second layer introduces a longitudinal learning companion that accumulates student work across courses, enabling reflective learning and supporting the translation of academic experience into career-relevant representations. Together, these layers establish a continuous learning pipeline that connects immediate instructional support with long-term educational and professional outcomes.

The contributions of this paper are threefold. First, we formally define the concept of Instructor-Grounded Agentic AI and position it within the broader landscape of AI in education. Second, we propose a

unified architectural framework that links course-level assistance with longitudinal learning and career alignment. Third, we present a research design and preliminary insights from an initial deployment, demonstrating the feasibility and potential impact of the approach. By reframing AI systems as extensions of instructional design rather than standalone tools, this work contributes toward a new paradigm in educational technology, one that emphasizes alignment, continuity, and the integration of learning with real-world application.

2. Related Work

Research on intelligent tutoring systems (ITS) has shown that adaptive and interactive guidance can significantly improve student learning outcomes. Early studies demonstrate that systems incorporating feedback, hints, and step-by-step reasoning can approach the effectiveness of human tutors [1]. However, most ITS are domain-specific and require extensive manual design, limiting their scalability across subjects and courses.

Recent advances in large language models have enabled more flexible AI support in education. General-purpose systems such as ChatGPT are widely used by students for explanations, problem-solving, and writing assistance. While these models provide high-quality responses, they are not grounded in course-specific content or instructor intent, often leading to misalignment with instructional goals [2], [3]. Retrieval-Augmented Generation (RAG) has emerged as a method for grounding language models in external knowledge sources. By retrieving relevant documents at query time, RAG improves factual accuracy and contextual relevance [4]. In educational settings, this approach allows AI systems to reference course materials such as lecture notes and readings. Existing implementations, however, primarily focus on information retrieval and do not incorporate structured pedagogical strategies such as scaffolding or guided

reasoning. Parallel work on reasoning-enhanced language models has explored the integration of multi-step reasoning and tool use within AI systems. Approaches such as chain-of-thought prompting and agentic reasoning frameworks enable models to produce more structured and interpretable responses [5], [6]. These methods improve reasoning capability but are not typically aligned with instructional design or instructor-specific context.

A related body of research addresses the gap between academic learning and career readiness. Studies consistently show that students struggle to articulate their skills and experiences in professional contexts [7], [8]. Existing AI tools in this space focus on resume generation or keyword matching rather than leveraging structured academic work as a knowledge source.

3. Problem Formulation

Current AI-assisted learning systems operate primarily as general-purpose response generators. Given a student query, these systems produce outputs based on broad pre-trained knowledge without explicit conditioning on instructional context or pedagogical intent. This interaction can be represented as a function:

$$y = f(x)$$

where x denotes the student query and y denotes the generated response. While effective for general explanation and information retrieval, this formulation does not incorporate course-specific knowledge, instructional strategies, or prior learning context.

In formal educational settings, learning is inherently structured and context-dependent. Effective instructional support requires alignment with course materials, adherence to pedagogical strategies, and continuity across interactions. To capture these requirements, the learning process can be expressed as:

$$y = f(x, C, P, M)$$

where C represents course-specific knowledge, P denotes pedagogical strategy,

and M captures longitudinal student memory across interactions. The absence of these components in general-purpose AI systems results in a misalignment between generated responses and instructional objectives. In particular, responses generated without instructional grounding differ from those conditioned on structured educational context. To address this limitation, educational AI systems must incorporate three key properties: grounding in instructor-defined knowledge, integration of pedagogical guidance, and persistence through longitudinal memory. These requirements motivate the instructor-Grounded Agentic AI framework proposed in this work.

4. Proposed Framework

4.1 Definition

We define *Instructor-Grounded Agentic AI* as an AI system that is explicitly conditioned on instructor-specific knowledge and guided by pedagogical strategies, while maintaining the ability to adapt over time through continuous interaction. Unlike general-purpose AI systems, which generate responses based on global knowledge, the proposed framework incorporates structured instructional context and learning objectives into the response generation process. Formally, the system models interaction as a policy:

$$y_t = \pi(x_t, C, P, M_{t-1})y_t$$

where x_t denotes the student query at time t , C represents course-specific knowledge, P denotes pedagogical strategy, and M_{t-1} captures prior interactions. The policy π produces a response y_t that is conditioned on both the current input and accumulated learning context.

4.2 System Architecture

The proposed framework adopts a dual-layer architecture that supports both immediate instructional interaction and long-term learning continuity. The first layer is a *course-grounded assistant*. This component is designed to extend the instructor's presence beyond the classroom by providing

responses grounded in course materials such as lecture notes, readings, and assignments. It ensures that student interactions remain aligned with instructor-defined content and learning objectives.

The second layer is a *longitudinal learning companion*. This component maintains a persistent representation of student work across courses and over time. It enables students to revisit prior learning, identify patterns in their academic development, and connect knowledge across domains. This layer also supports career-oriented tasks by mapping academic experiences to external requirements such as job descriptions or graduate applications. Together, these layers establish a continuous learning structure that links short-term instructional support with long-term knowledge development.

4.3 Core Components

The framework is built on three primary components: knowledge grounding, pedagogical guidance, and longitudinal memory.

Knowledge grounding (C) is implemented through retrieval mechanisms that allow the system to access course-specific materials at query time. This ensures that responses are consistent with instructor-provided content rather than relying solely on pre-trained knowledge.

Pedagogical strategy (P) governs how responses are generated. Instead of directly providing answers, the system incorporates instructional techniques such as scaffolding, guided reasoning, and Socratic questioning. These strategies aim to promote deeper understanding rather than surface-level task completion.

Longitudinal memory (M) captures prior interactions and accumulated student work. This component enables the system to maintain continuity across sessions, supporting progressive learning and the integration of knowledge over time.

4.4 Agentic Interaction

A key feature of the proposed framework is its agentic capability. The system is designed to perform multi-step reasoning and adapt its

responses based on both context and interaction history. Rather than generating a single static response, the system can guide students through a sequence of steps, adjusting its behavior based on intermediate feedback. This interaction model supports iterative learning, where students engage in dialogue that evolves over time [9]. The incorporation of pedagogical strategies within this process ensures that reasoning is structured and aligned with instructional goals.

4.5 System Pipeline

The overall interaction process can be described as a sequence of operations. Given a student query x_t , the system first retrieves relevant course materials from C . It then applies pedagogical strategies P to determine the form and structure of the response. Finally, it incorporates historical context from M_{t-1} to ensure continuity. This results in a response y_t that reflects not only the immediate query but also the broader instructional and learning context. The memory component is then updated to include the current interaction:

$$M_t = M_{t-1} \cup \{x_t, y_t\}$$

This iterative process enables the system to function as a continuous learning agent, supporting both immediate academic tasks and long-term knowledge development.

The overall architecture of the proposed Instructor-Grounded Agentic AI framework, including its grounding, agentic reasoning, and longitudinal learning components, is illustrated in Figure 1.




Figure 1. Architecture of Instructor-Grounded Agentic AI, integrating knowledge grounding, agentic reasoning, and longitudinal memory for continuous learning and career alignment.

5. Research Design and Methodology

To evaluate the proposed Instructor-Grounded Agentic AI framework, this work outlines a mixed-methods research design intended to assess the impact of instructor grounding, pedagogical guidance, and longitudinal memory on learning and career alignment. The methodology is presented as a structured evaluation plan that can be applied in future empirical studies.

5.1 Study Design

The proposed evaluation follows a two-phase design. The first phase consists of a retrospective analysis of system interactions, intended to examine patterns of student engagement and reasoning behavior. The second phase is designed as a quasi-experimental study to compare learning outcomes between students using the proposed system and those relying on general-purpose AI tools. This design enables both observational and comparative analysis, supporting the study of interaction dynamics as well as measurable learning outcomes.

5.2 Data Sources

The proposed evaluation framework considers multiple data sources to capture different aspects of the learning process.

These include interaction logs containing timestamped student queries and system responses, assessment data such as short-answer and applied learning tasks, student feedback collected through surveys and reflective responses, and instructor input obtained through semi-structured interviews. Together, these sources are intended to provide both quantitative and qualitative perspectives on system behaviour, student engagement, and learning outcomes.

5.3 Evaluation Metrics

The proposed evaluation focuses on three primary dimensions: learning depth, interaction continuity, and career alignment.

- Learning depth: Can be assessed using a revised version of Bloom's taxonomy [11], which categorizes responses based on levels of cognitive complexity from recall to higher-order reasoning; interaction transcripts can be coded to evaluate the extent to which the system supports analytical and reflective thinking.
- Interaction continuity: Can be measured through session-based metrics such as interaction length, return frequency, and progression across sessions, capturing sustained engagement over time.
- Career alignment: Can be evaluated by analyzing how effectively students relate academic experiences to professional contexts, including structured tasks mapping coursework to job descriptions or academic applications, as well as expert evaluation of response quality.

5.4 Pedagogical Evaluation

The framework proposes evaluating the effectiveness of embedded pedagogical strategies, including scaffolding, guided reasoning, and Socratic questioning. Prior work in intelligent tutoring systems suggests that such strategies support deeper learning by encouraging active engagement and metacognitive reflection [1], [12]. Interactions can be examined to determine whether the system promotes reasoning processes rather than direct answer

generation, providing insight into the instructional quality of responses.

5.5 Analytical Approach

The proposed evaluation framework includes both quantitative and qualitative analytical components. Quantitative data can be analyzed using comparative statistical methods, such as analysis of variance (ANOVA), to examine differences in learning outcomes between experimental conditions. Interaction-based metrics, including session length and progression patterns, can be used to identify trends in student engagement and learning continuity over time.

Qualitative data, including student feedback and instructor input, can be analyzed using thematic coding. This approach enables the identification of recurring patterns related to system usability, perceived learning benefits, and alignment with instructional objectives. Together, these analytical methods are intended to provide a comprehensive understanding of both measurable outcomes and user experience.

5.6 Ethical Considerations

The proposed research design follows standard ethical guidelines for educational studies. Data collection is expected to ensure anonymization of student information, voluntary participation, and secure handling of all collected data. The system is designed to operate without using student interaction data for model training, thereby supporting privacy and data protection. These considerations are critical for ensuring responsible deployment of AI systems in educational environments.

6. Discussion and Implications

The proposed Instructor-Grounded Agentic AI framework represents a shift in how AI systems are conceptualized in educational settings. Rather than functioning as general-purpose tools, the framework positions AI as an extension of instructional design, explicitly aligned with course structure, pedagogical intent, and long-term learning objectives. This perspective is consistent

with foundational learning theories that emphasize guided interaction and the importance of structured support in cognitive development [10].

From an educational perspective, the integration of pedagogical strategies such as scaffolding and guided reasoning reflects established principles of effective instruction. Prior work suggests that learning is most effective when students are supported in progressing from basic understanding to higher-order reasoning through structured guidance [11], [12]. By embedding these strategies directly into AI interaction, the proposed framework encourages active engagement and deeper cognitive processing, rather than passive answer consumption.

The inclusion of longitudinal memory extends the role of AI from short-term assistance to continuous learning support. This aligns with the concept of connected learning, where knowledge develops through the integration of experiences across contexts and over time [13]. By maintaining a persistent representation of student work, the system enables reflection, supports knowledge transfer, and strengthens the connection between academic learning and real-world application.

From a technical standpoint, the framework highlights the importance of combining retrieval-based grounding, agentic reasoning, and memory mechanisms within a unified system. Existing approaches often treat these components independently, focusing either on retrieval accuracy, reasoning capability, or interaction modelling. The proposed integration suggests a new direction for AI system design in which context, reasoning, and continuity are jointly considered to support learning outcomes.

Finally, the framework has important implications for career readiness. By enabling students to map academic work to professional contexts, the system addresses a well-documented gap between education and workforce preparation. This positions

AI not only as a tool for academic support but also as a mechanism for translating knowledge into actionable skills, supporting students in navigating transitions beyond formal education.

7. Conclusion

This paper introduced Instructor-Grounded Agentic AI as a framework for designing educational AI systems that integrate course-specific knowledge, pedagogical strategies, and longitudinal interaction. By modeling learning as a context-conditioned and continuous process, the proposed approach

extends beyond general-purpose AI assistance toward systems that support structured learning and career alignment. The framework highlights the importance of grounding, reasoning, and memory as core components of educational AI. While the current work focuses on conceptual design and proposed evaluation, future research will focus on empirical validation and deployment across diverse educational settings. The proposed direction contributes toward a more aligned and continuous model of AI-supported learning.

References

- [1] VanLehn, K. The relative effectiveness of human tutoring, intelligent tutoring systems, and other tutoring systems. *Educational Psychologist*. 2011. 46(4): 197–221.
- [2] Kasneci, E., Sessler, K., Küchemann, S., Bannert, M., Dementieva, D., Fischer, F., Kasneci, G. ChatGPT for good? On opportunities and challenges of large language models for education. *Learning and Individual Differences*. 2023.
- [3] OpenAI. GPT-4 Technical Report. *arXiv preprint arXiv:2303.08774*, 2023.
- [4] Lewis, P., Perez, E., Piktus, A., Petroni, F., Karpukhin, V., Goyal, N., Kiela, D. Retrieval-augmented generation for knowledge-intensive NLP tasks. In: *Advances in Neural Information Processing Systems (NeurIPS)*. 2020.
- [5] Wei, J., Wang, X., Schuurmans, D., Bosma, M., Ichter, B., Xia, F., Zhou, D. Chain-of-thought prompting elicits reasoning in large language models. In: *NeurIPS*. 2022.
- [6] Yao, S., Zhao, J., Yu, D., Du, N., Shafran, I., Narasimhan, K., Cao, Y. ReAct: Synergizing reasoning and acting in language models. In: *International Conference on Learning Representations (ICLR)*. 2023.
- [7] National Association of Colleges and Employers (NACE). *Job Outlook 2023*. Bethlehem, PA: NACE.
- [8] Hart Research Associates. *Fulfilling the American dream: Liberal education and the future of work*. Association of American Colleges and Universities (AAC&U), 2018.
- [9] Park, J. S., O'Brien, J., Cai, C. J., Morris, M. R., Liang, P., Bernstein, M. S. Generative agents: Interactive simulacra of human behavior. In: *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology (UIST)*. 2023.
- [10] Vygotsky, L. S. *Mind in society: The development of higher psychological processes*. Harvard University Press. 1978.
- [11] Anderson, L. W., Krathwohl, D. R. *A taxonomy for learning, teaching, and assessing: A revision of Bloom's taxonomy*. Longman. 2001.
- [12] Graesser, A. C., D'Mello, S., Cade, W. Instruction based on tutoring. In: Mayer, R. E., Alexander, P. A. (eds.), *Handbook of Research on Learning and Instruction*. Routledge. 2012.
- [13] Ito, M., Gutierrez, K., Livingstone, S., Penuel, B., Rhodes, J., Salen, K., Watkins, S. C. *Connected learning: An agenda for research and design*. Digital Media and Learning Research Hub. 2013.