

Learning with AI, Not Just Using AI: A Human-in-the-Loop Curriculum for Critical Literacy

Jiwon Jang*
Graduate School of Data Science
Seoul National University
Seoul, Republic of Korea
jjw6424@snu.ac.kr

Jaeseong Ju*
Graduate School of Data Science
Seoul National University
Seoul, Republic of Korea
cg3731@snu.ac.kr

Dokyung Lee*
Graduate School of Data Science
Seoul National University
Seoul, Republic of Korea
didroldk@snu.ac.kr

Jaeun Seo*
Graduate School of Data Science
Seoul National University
Seoul, Republic of Korea
jaeunseo@snu.ac.kr

Hyunwoo Park†
Graduate School of Data Science
Seoul National University
Seoul, Republic of Korea
hyunwoopark@snu.ac.kr

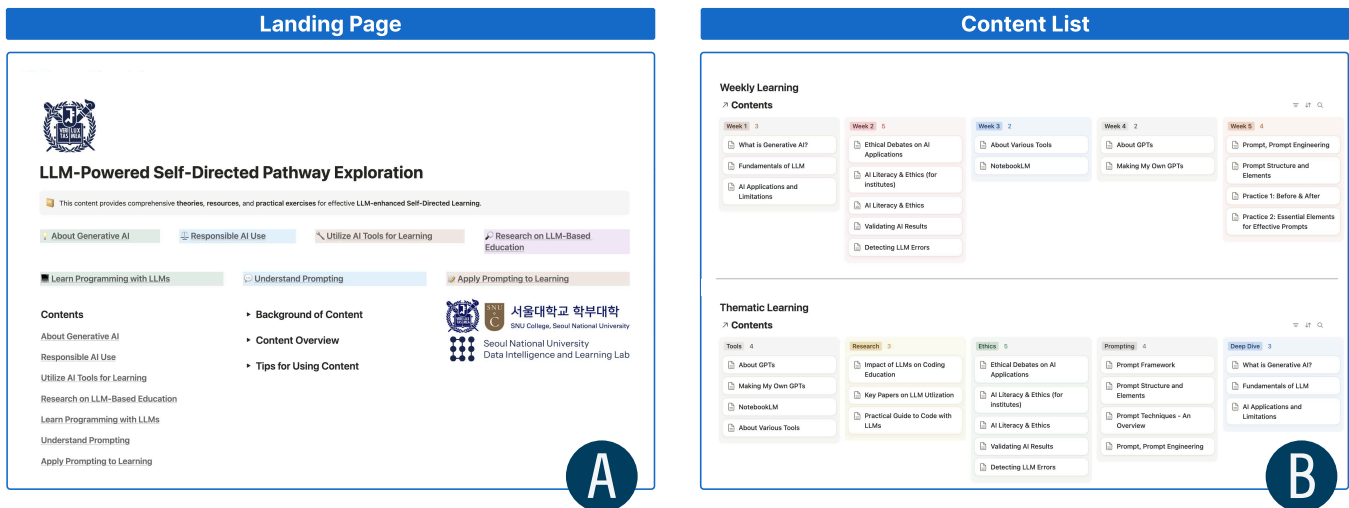


Figure 1: Overview of the LLM-Based Educational Platform: (A) Landing page introducing the pedagogical framework. (B) Content organization showing weekly and thematic learning modules designed to support AI-mediated learning. The full content (in Korean) is available at <https://buly.kr/ChqMXP5>.

Abstract

The widespread accessibility of Large Language Models (LLMs) has disrupted traditional pedagogical norms, forcing a critical re-evaluation of authorship, literacy, and academic integrity. This work presents the design and reflection of a university-level curriculum that aims not merely to teach operational AI tool usage, but to

cultivate critical, responsible, and research-informed engagement with generative AI. We describe a structured instructional framework that integrates conceptual foundations of probabilistic generation, responsible academic practices, tool differentiation, research literacy, and programming-prompting workflows, emphasizing human-in-the-loop reasoning and process transparency through mechanisms such as audit trails and structured reflection. Through this implementation, we identify three key tensions shaping AI-mediated education: tool obsolescence in rapidly evolving ecosystems, the need for process-centered assessment, and disparities in AI interaction quality across disciplinary environments. We argue that sustainable AI pedagogy requires balancing tool-centered engagement with concept-centered literacy, enabling transferable competencies that persist beyond any single platform.

*Authors contributed equally to this research.

†Corresponding author

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
Conference'17, Washington, DC, USA

© 2026 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 978-x-xxxx-xxxx-x/YYYY/MM
<https://doi.org/10.1145/nnnnnnn.nnnnnnn>

Keywords

Generative AI in Education, AI Literacy, Human–AI Collaboration, Process-Centered Assessment

ACM Reference Format:

Jiwon Jang, Jaeseong Ju, Dokyung Lee, Jaeun Seo, and Hyunwoo Park. 2026. Learning with AI, Not Just Using AI: A Human-in-the-Loop Curriculum for Critical Literacy. In . ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/nnnnnnnn.nnnnnnn>

1 Introduction

The rapid integration of generative AI into higher education has fundamentally altered academic work, offering efficiency while challenging traditional notions of literacy and intellectual ownership [5, 12]. While AI tools reduce friction in tasks like drafting and debugging, their rapid evolution risks rendering tool-specific instruction immediately irrelevant [25]. Furthermore, an “AI divide” persists: computational disciplines often embrace optimization, whereas humanities fields may perceive conflicts with creative practices [6, 32]. This landscape complicates the distinction between learning *with* AI and learning *replaced* by AI, making traditional artifacts unreliable indicators of understanding [18, 26].

To address these systemic shifts, this study proposes a resilient pedagogical framework that transcends technological flux. We investigate how to design curricula that balance practical tool exposure with enduring cognitive skills, fostering universal AI literacy across technical and non-technical disciplines [14, 19]. Central to our approach is shifting the educational focus from tool proficiency to “process visibility,” where reasoning, validation, and revision become the primary objects of assessment [21, 24].

We present the design and reflection of an interdisciplinary curriculum structured around four pillars: (1) conceptual foundations of probabilistic systems, (2) responsible verification practices, (3) research-informed tool differentiation, and (4) human-centered workflows. By synthesizing these elements, we offer a validated model for AI-mediated learning that ensures conceptual durability [17, 19]. Ultimately, this work navigates key tensions in assessment and disciplinary divergence, aiming to shift the discourse from how to use tools” toward how to learn meaningfully in the presence of AI.”

2 Background & Related Works

2.1 Generative AI in Higher Education and the Assessment Paradox

Initially adopted as auxiliary tools for information retrieval or grammar correction, generative AI systems have evolved into versatile conversational agents capable of producing essays, programming code, research summaries, and creative artifacts on demand [5, 12]. This expansion of capability has generated both optimism and concern within educational communities. On the one hand, AI offers opportunities for personalized feedback, accessibility support, and reduced cognitive barriers for novice learners [10]. Conversely, it introduces risks related to academic integrity, overreliance, epistemic trust, and the erosion of foundational reasoning skills [7, 26].

A growing perspective suggests that introducing AI tools at the undergraduate entry stage is particularly critical. Rather than mere

technical skill acquisition, early exposure functions as the development of cross-disciplinary cognitive infrastructure. It enables students to contextualize these systems before disciplinary habits become rigidly specialized, allowing them to approach AI not as a domain-specific instrument but as a flexible cognitive partner [17]. Importantly, the pedagogical value lies less in mastering interfaces and more in cultivating transferable competencies that remain relevant despite rapid technological change [14].

However, this pedagogical imperative creates a fundamental paradox in assessment. Traditional evaluation systems often rely on the quality of final artifacts as proxies for understanding. Yet, when AI systems can generate polished outputs with minimal human effort, artifact-centered assessment becomes fragile as an indicator of genuine learning [18]. Consequently, academic integrity is shifting from a focus on detecting unauthorized assistance to redefining legitimate evidence of learning. Institutions are increasingly prioritizing process-centered assessment, emphasizing transparency, verification practices, and reflective documentation, rather than relying on the final output alone.

2.2 From Critical Literacy to Human-in-the-Loop Learning

The integration of AI into education necessitates a shift from operational proficiency toward broader forms of AI literacy. AI literacy extends traditional notions of digital and information literacy by emphasizing not only the ability to use tools but also the capacity to understand their limitations, biases, and probabilistic nature [19]. Rather than equating literacy with efficiency, recent scholarship frames it as the cultivation of critical awareness, skepticism, and epistemic responsibility. In AI-mediated environments, this means recognizing that generated content is neither authoritative nor neutral, but contingent upon training data and prompt structure [25].

This critical foundation enables the *Human-in-the-Loop* (HITL) learning model, which conceptualizes AI not as an autonomous replacement but as an augmentative partner. This paradigm emphasizes mixed-initiative interaction, where both human and machine contribute to problem solving while maintaining human oversight [3]. Concepts such as cognitive offloading and co-creative interaction illustrate how AI can extend human capabilities without displacing core reasoning processes [21]. In educational contexts, this approach reframes generative AI from an answer-producing mechanism into a dialogic medium that supports inquiry, iteration, and metacognitive reflection, ensuring that students actively guide and critique the AI’s output [21].

2.3 Context of the Study: AI-Integrated Curriculum Initiative

This study stems from the “Advanced Technology Educational Content Development” initiative at the Undergraduate College of Seoul National University. The program seeks to innovate pedagogy by utilizing emerging technologies, such as AI, AR/VR, and the Metaverse, to create immersive and personalized learning environments. It invites proposals that integrate these tools into classroom instruction to drive digital transformation in undergraduate education.

In response to this call, the authors proposed and executed a project specifically focused on “Learning how to learn with LLMs.” Our objective was to move beyond the novelty of the technology to explore how Large Language Models could function as effective cognitive partners for self-regulated learning. Over a six-month period, we conducted a comprehensive development process that included extensive literature reviews, empirical evaluation of generative tools, student interviews, and iterative content design. The resulting curriculum and guidelines were synthesized into a web-based educational resource, the specific structure and components of which are detailed in the following section.

3 Implementation

3.1 Foundations and Responsible Use

3.1.1 Generative AI Foundations. We introduced generative AI as a class of probabilistic systems that can produce plausible outputs across text, code, and multimodal content, while remaining fundamentally non-authoritative. Instead of teaching a single product workflow, this module aimed to build transferable conceptual understanding that remains applicable as tools evolve [4]. This emphasis is particularly important for students in undeclared-major programs, who will encounter generative AI in diverse disciplinary settings (e.g., humanities writing, design ideation, quantitative analysis, and programming). Accordingly, we focused on the practical implications of how generative models behave. We explained how instruction following emerges as conditional generation, why outputs can appear confident despite being incorrect, and what typical failure modes look like in academic work, including hallucinated facts, fabricated citations, and hidden assumptions. This framing helps students treat model responses not as final answers, but as drafts that require scrutiny, contextualization, and revision.

3.1.2 Responsible Use as an Academic Practice. We framed responsible use not as abstract ethics but as a set of actionable academic practices that align AI assistance with learning goals while mitigating risks related to integrity, privacy, and harm [8]. The module emphasized concrete norms by asking students to verify claims through triangulation, document prompts and outputs, avoid disclosing sensitive information, and check for bias and overgeneralization. Crucially, we articulated boundaries between permissible support (e.g., brainstorming, debugging) and impermissible substitution (e.g., submitting AI-generated work as one’s own). To operationalize these norms, students engaged in reflective exercises comparing outputs under controlled prompts and produced lightweight audit trails recording what was asked, returned, verified, and used.

3.2 Learning with AI Tools & Research-Informed Practice

3.2.1 Learning with AI Tools. By experiencing specialized tools designed for different purposes (research assistance, content creation, workflow automation), students learn to distinguish between tools rather than treating GenAI as a monolithic entity. Learning to create custom GPTs shifts students from passive consumers to informed designers, fostering critical awareness of how these tools are constructed and what affordances they provide. This hands-on

exposure is particularly crucial for students in undeclared major programs who will encounter GenAI across vastly different disciplinary contexts—from literature analysis to data science, from design to engineering. Before committing to specific majors, students need broad exposure to understand which AI tools align with particular disciplinary practices and learning goals. Rather than accepting or rejecting AI wholesale, they develop the capacity to make nuanced, context-specific judgments about when and how to integrate these tools into their learning processes across diverse academic domains.

3.2.2 LLM-Based Learning Research. Fixed guidelines become outdated as technology evolves rapidly. Equipping students with research literacy skills allows them to critically engage with emerging evidence about how LLMs affect learning processes, building a transferable approach that remains relevant despite technological change. By analyzing empirical studies alongside usage guidelines, particularly in programming education where LLM integration is already widespread, students learn to distinguish between anecdotal claims and evidence-based conclusions about AI’s educational impact. This research-informed stance is essential for self-determined technology use. Students learn to identify both opportunities (personalized feedback, error explanation) and risks (overreliance, reduced problem-solving practice) documented in research, enabling strategic adoption and resistance. Rather than following institutional mandates or tech advocacy, they can ground their decisions about when AI supports learning goals versus when it introduces counterproductive shortcuts in empirical evidence. This critical engagement models an approach students can apply to future tools and contexts they will encounter throughout their academic and professional lives.

3.3 Bridging Fundamentals with AI: Programming & Prompting

3.3.1 Democratizing Computational Thinking through Structured AI Usage. We presented programming as a universal tool for problem-solving and automation rather than solely a Computer Science discipline. Generative AI has significantly lowered entry barriers through natural language programming. However, we argue that effective utilization still requires a fundamental understanding of computational mechanisms. To bridge the gap between intent and execution, we structured the curriculum around a systematic 11-stage programming workflow, grounded in traditional software construction principles [15]. This workflow ranges from problem analysis and algorithm design to variable planning, implementation, and documentation. In each stage, we integrated specific LLM interactions. For example, students used LLMs to draft pseudocode during the design phase or to interpret error messages during debugging. This structure ensures that LLMs augment the thinking process instead of allowing students to bypass the foundational understanding of code construction.

3.3.2 Prompt Engineering as a Framework for Active Inquiry. We framed prompting not as a collection of optimization hacks, but as a methodology to understand the capabilities and limitations of LLMs.

Drawing from recent survey paper [23], we curated prompting techniques specifically for self-directed learning contexts. These techniques included Chain-of-Thought (CoT) [30] and Persona adoption [31]. For instance, students applied 'Persona' to generate customized quizzes and 'CoT' to decompose complex study topics. A core pedagogical emphasis was that high-quality AI output depends on high-quality human input. We facilitated a "Self-Learning Lab" where students applied these strategies to personal learning goals, such as acquiring a new language or web development skills. This process demonstrated that effective GenAI usage requires users to define problems and structure inquiries deeply. Consequently, it reinforces the "human-in-the-loop" as the primary orchestrator of the learning process.

4 Discussion

4.1 Tool obsolescence and pedagogical sustainability

A central tension in implementing AI curricula is the risk of tool obsolescence within rapidly evolving ecosystems. While platforms like GPTs and NotebookLM provide authentic experience, the rapid shift in their capabilities exposes the fragility of tool-specific knowledge. Consequently, our framework emphasizes transferable mental models [14]—such as probabilistic generation and verification practices—over fixed procedural workflows. However, tool-centered learning remains a vital entry point that grounds abstract concepts in recognizable academic tasks [19]. Pedagogical sustainability is thus achieved by treating specific tools as interchangeable exemplars while maintaining invariant competencies, like research literacy and human-in-the-loop reasoning, as the stable core [29]. By oscillating between concrete engagement and conceptual abstraction, we enable students to adapt their understanding as new tools inevitably emerge.

4.2 Recalibrating Assessment for Learning with Generative AI

The proliferation of generative AI has significantly lowered the cost of producing high-quality artifacts, making it increasingly difficult to infer student understanding from final outputs alone [13]. When assessment remains fixated on these products, it incentivizes the use of AI as a performance shortcut rather than a learning support [9]. The critical challenge is therefore to redefine valid evidence of learning by shifting evaluation focus from the final answer to the process of inquiry.

In AI-mediated environments, meaningful indicators of learning lie in how students frame problems, interrogate sources, and revise outputs. We propose using "lightweight process records"—such as revision traces or verification logs—as pedagogical instruments to make these invisible cognitive steps accountable [11]. This approach prioritizes human judgment, evaluating whether learners treat AI outputs as provisional drafts requiring critical refinement.

Operationally, this necessitates rubrics that explicitly reward verification, error correction, and uncertainty management alongside correctness [20]. To ensure equity, assessment criteria must prioritize epistemic rigor over mere fluency with specific tools [1]. By formally recognizing and rewarding the labor of verification

and justification, assessment can transform generative AI from a mechanism for score optimization into a catalyst for deeper, critical learning.

4.3 Navigating the AI Divide: From Usage Quality to Disciplinary Ecosystems

Our observations reveal a significant divergence in interaction quality, where power users "accelerate growth while others exhibit counterproductive behaviors like blindly accepting unverified outputs [2]. This misuse often stems from a disparity in accumulated trial-and-error experiences, preventing students from developing the literacy of skepticism" needed to navigate hallucinations [29].

We argue that this experience gap is closely tied to disciplinary ecosystems, as a student's major shapes their intrinsic motivation to automate. For instance, while Computer Science prioritizes optimization, Humanities students may view such tools as secondary to creative processes, rendering the utility of AI a subjective perception shaped by the academic environment [16].

Therefore, fostering true self-determination [22] requires abandoning a monolithic curriculum that risks being inefficient for experts while unapproachable for novices. Instead, we propose a flexible framework [27] that establishes a baseline of critical literacy while branching into context-aware applications that respect diverse disciplinary needs [28]. The ultimate goal is to scaffold this learning to bridge the experience gap "without imposing a singular, engineering-centric view of AI proficiency.

5 Conclusion

We explored how higher education can move beyond binary narratives of adoption versus resistance and instead foster critical engagement with generative AI. Through the design and reflection of a university-level curriculum, we argued that the value of AI integration lies not in mastering specific tools but in cultivating transferable competencies such as verification practices, research literacy, and human-in-the-loop reasoning.

This work offers a reflective framework rather than a finalized model. Future efforts may extend this approach through empirical evaluation across diverse institutional contexts and collaborative curriculum design. Ultimately, the goal is not to normalize AI use uncritically, but to support learners in exercising informed judgment, knowing when to adopt, when to question, and how to engage with AI as part of their evolving academic and professional practices.

References

- [1] Riordan Alfredo, Vanessa Echeverria, Yueqiao Jin, Lixiang Yan, Zachari Swiecki, Dragan Gašević, and Roberto Martinez-Maldonado. 2024. Human-centred learning analytics and AI in education: A systematic literature review. *Computers and Education: Artificial Intelligence* 6 (2024), 100215. doi:10.1016/j.caeai.2024.100215
- [2] Laura Kristen Allen and Panayiota Kendeou. 2024. ED-AI Lit: An interdisciplinary framework for AI literacy in education. *Policy Insights from the Behavioral and Brain Sciences* 11, 1 (2024), 3–10.
- [3] Saleema Amershi, Maya Cakmak, William Bradley Knox, and Todd Kulesza. 2014. Power to the people: The role of humans in interactive machine learning. *AI magazine* 35, 4 (2014), 105–120.
- [4] Ravinithesh Annapureddy, Alessandro Fornaroli, and Daniel Gatica-Perez. 2025. Generative AI Literacy: Twelve Defining Competencies. *Digit. Gov. Res. Pract.* 6, 1, Article 13 (Feb. 2025), 21 pages. doi:10.1145/3685680
- [5] David Baidoo-Anu and Leticia Owusu Ansah. 2023. Education in the era of generative artificial intelligence (AI): Understanding the potential benefits of ChatGPT in promoting teaching and learning. *Journal of AI* 7, 1 (2023), 52–62.

- [6] Cecilia Ka Yuk Chan and Wenjie Hu. 2023. Students' voices on generative AI: Perceptions, benefits, and challenges in higher education. *International Journal of Educational Technology in Higher Education* 20, 1 (2023), 43.
- [7] Debby RE Cotton, Peter A Cotton, and J Reuben Shipway. 2024. Chatting and cheating: Ensuring academic integrity in the era of ChatGPT. *Innovations in education and teaching international* 61, 2 (2024), 228–239.
- [8] Christian Djeflal. 2025. Reflexive Prompt Engineering: A Framework for Responsible Prompt Engineering and AI Interaction Design. In *Proceedings of the 2025 ACM Conference on Fairness, Accountability, and Transparency (FAccT '25)*. Association for Computing Machinery, New York, NY, USA, 1757–1768. doi:10.1145/3715275.3732118
- [9] Michail Giannakos, Roger Azevedo, Peter Brusilovsky, Mutlu Cukurova, Yannis Dimitriadis, Davinia Hernandez-Leo, Sanna Järvelä, Manolis Mavriakis, and Bart Rienties. 2025. The promise and challenges of generative AI in education. *Behaviour & Information Technology* 44, 11 (2025), 2518–2544. doi:10.1080/0144929X.2024.2394886
- [10] Wayne Holmes, Maya Bialik, and Charles Fadel. 2019. *Artificial intelligence in education promises and implications for teaching and learning*. Center for Curriculum Redesign.
- [11] Md Naimul Hoque, Tasfia Mashiat, Bhavya Ghai, Cecilia Shelton, Fanny Chevalier, Kari Kraus, and Niklas Elmqvist. 2024. The HaLLMark Effect: Supporting Provenance and Transparent Use of Large Language Models in Writing with Interactive Visualization. arXiv:2311.13057 [cs.HC] <https://arxiv.org/abs/2311.13057>
- [12] Enkelejda Kasneci, Kathrin Seßler, Stefan Küchemann, Maria Bannert, Daryna Dementieva, Frank Fischer, Urs Gasser, Georg Groh, Stephan Günemann, Eyke Hüllermeier, et al. 2023. ChatGPT for good? On opportunities and challenges of large language models for education. *Learning and individual differences* 103 (2023), 102274.
- [13] Hao-Ping (Hank) Lee, Advait Sarkar, Lev Tankelevitch, Ian Drosos, Sean Rintel, Richard Banks, and Nicholas Wilson. 2025. The Impact of Generative AI on Critical Thinking: Self-Reported Reductions in Cognitive Effort and Confidence Effects From a Survey of Knowledge Workers. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Association for Computing Machinery, New York, NY, USA, Article 1121, 22 pages. doi:10.1145/3706598.3713778
- [14] Duri Long and Brian Magerko. 2020. What is AI literacy? Competencies and design considerations. In *Proceedings of the 2020 CHI conference on human factors in computing systems*. 1–16.
- [15] S. McConnell. 2004. *Code Complete*. Pearson Education. <https://books.google.co.kr/books?id=LpVCAwAAQBAJ>
- [16] Reza Hadi Mogavi, Chao Deng, Justin Juho Kim, Pengyuan Zhou, Young D Kwon, Ahmed Hosny Saleh Metwally, Ahmed Tlili, Simone Bassanelli, Antonio Bucchiarone, Sujit Gujar, et al. 2024. ChatGPT in education: A blessing or a curse? A qualitative study exploring early adopters' utilization and perceptions. *Computers in Human Behavior: Artificial Humans* 2, 1 (2024), 100027.
- [17] Ethan R Mollick and Lilach Mollick. 2023. Using AI to implement effective teaching strategies in classrooms: Five strategies, including prompts. *The Wharton School Research Paper* (2023).
- [18] Philip Newton. 2016. Academic integrity: A quantitative study of confidence and understanding in students at the start of their higher education. *Assessment & Evaluation in Higher Education* 41, 3 (2016), 482–497.
- [19] Davy Tsz Kit Ng, Jac Ka Lok Leung, Samuel Kai Wah Chu, and Maggie Shen Qiao. 2021. Conceptualizing AI literacy: An exploratory review. *Computers and Education: Artificial Intelligence* 2 (2021), 100041.
- [20] Aditya Pathak, Rachit Gandhi, Vaibhav Uttam, Arnav Ramamoorthy, Pratyush Ghosh, Aaryan Raj Jindal, Shreyash Verma, Aditya Mittal, Aashna Ased, Chirag Khatri, Yashwanth Nakka, Devansh, Jagat Sesh Challa, and Dhruv Kumar. 2025. Rubric Is All You Need: Improving LLM-Based Code Evaluation With Question-Specific Rubrics. In *Proceedings of the 2025 ACM Conference on International Computing Education Research V.1 (ICER '25)*. ACM, 181–195. doi:10.1145/3702652.3744220
- [21] Ido Roll and Ruth Wylie. 2016. Evolution and revolution in artificial intelligence in education. *International journal of artificial intelligence in education* 26, 2 (2016), 582–599.
- [22] Richard M Ryan and Edward L Deci. 2024. Self-determination theory. In *Encyclopedia of quality of life and well-being research*. Springer, 6229–6235.
- [23] Pranab Sahoo, Ayush Kumar Singh, Sriparna Saha, Vinija Jain, Samrat Mondal, and Aman Chadha. 2025. A Systematic Survey of Prompt Engineering in Large Language Models: Techniques and Applications. arXiv:2402.07927 [cs.AI] <https://arxiv.org/abs/2402.07927>
- [24] Gregory Schraw and Rayne Sperling Dennison. 1994. Assessing metacognitive awareness. *Contemporary educational psychology* 19, 4 (1994), 460–475.
- [25] Neil Selwyn. 2019. *Should robots replace teachers?: AI and the future of education*. John Wiley & Sons.
- [26] Miriam Sullivan, Andrew Kelly, and Paul McLaughlan. 2023. ChatGPT in higher education: Considerations for academic integrity and student learning. *Journal of Applied Learning & Teaching* 6, 1 (2023), 31–40.
- [27] Carol Ann Tomlinson. 1999. Mapping a route toward differentiated instruction. *Educational leadership* 57 (1999), 12–17.
- [28] Aida Walqui. 2006. Scaffolding instruction for English language learners: A conceptual framework. *International journal of bilingual education and bilingualism* 9, 2 (2006), 159–180.
- [29] Yoshija Walter. 2024. Embracing the future of Artificial Intelligence in the classroom: the relevance of AI literacy, prompt engineering, and critical thinking in modern education. *International Journal of Educational Technology in Higher Education* 21, 1 (2024), 15.
- [30] Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, brian ichter, Fei Xia, Ed Chi, Quoc V Le, and Denny Zhou. 2022. Chain-of-Thought Prompting Elicits Reasoning in Large Language Models. In *Advances in Neural Information Processing Systems*, S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh (Eds.), Vol. 35. Curran Associates, Inc., 24824–24837. https://proceedings.neurips.cc/paper_files/paper/2022/file/9d5609613524ecf4f15af0f7b31abca4-Paper-Conference.pdf
- [31] Jules White, Quchen Fu, Sam Hays, Michael Sandborn, Carlos Olea, Henry Gilbert, Ashraf Elnashar, Jesse Spencer-Smith, and Douglas C. Schmidt. 2023. A Prompt Pattern Catalog to Enhance Prompt Engineering with ChatGPT. arXiv:2302.11382 [cs.SE] <https://arxiv.org/abs/2302.11382>
- [32] Olaf Zawacki-Richter, Victoria I Marin, Melissa Bond, and Franziska Gouverneur. 2019. Systematic review of research on artificial intelligence applications in higher education—where are the educators? *International journal of educational technology in higher education* 16, 1 (2019), 1–27.