

EMBEDDING AI LITERACY WITHIN DISCIPLINARY CURRICULUM THROUGH STRUCTURED ASSIGNMENT DESIGN

Ronald J. Glotzbach

Purdue University / Associate Professor (USA)

Abstract

The rapid expansion of generative and agentic artificial intelligence (AI) in higher education is prompting renewed discussions about curriculum integration, academic integrity, and responsible technology adoption. While institutions debate policy and regulation, a central pedagogical question remains: how can AI literacy be embedded meaningfully within disciplinary coursework rather than treated as a standalone skill or external add-on? This paper explores an instructional design approach that integrates AI directly into structured course assignments while preserving disciplinary rigor. The discussion is grounded in the redesign of a specific undergraduate course assignment, used as a concrete example of discipline-embedded AI integration. Rather than positioning AI as a shortcut or replacement for student effort, the approach reframes AI as a tool that must be evaluated, constrained, and verified within existing learning objectives. In the redesigned assignment, students first demonstrate foundational competency independently before engaging generative AI to extend, analyze, and refine their work. Students document their prompts, test AI-generated outputs, identify inaccuracies, and justify revisions, ensuring that conceptual understanding remains central. The structured integration emphasizes that AI literacy involves understanding system limitations, critically evaluating outputs, and recognizing AI's influence on disciplinary thinking. By embedding AI use within established coursework, curriculum redesign becomes incremental rather than disruptive, allowing faculty to integrate AI without replacing core assignments. The paper will outline the design principles guiding this integration, including scaffolding, productive struggle, and transparent documentation, and will discuss how the model can be adapted across disciplines to support responsible AI adoption in higher education. As generative AI continues to influence academic work, instructional design plays a central role in shaping student engagement with these systems. Embedding AI literacy within disciplinary curriculum offers a sustainable pathway forward that balances innovation with rigor and prepares students to engage critically and ethically with emerging technologies.

Keywords: *AI literacy, curriculum integration, instructional design, generative AI, higher education.*

1. Introduction

Generative and agentic artificial intelligence (AI) systems are beginning to change the day-to-day reality of coursework in higher education. In technical courses especially, students now have easy access to tools capable of generating programming code, written explanations, database queries, and analytical responses within seconds. In many cases, the output looks convincing enough that students may assume it is correct before ever testing or evaluating it. This is starting to create some interesting challenges for pedagogy. Currently, many assignments in higher education are structured with the explicit goal of granting students insight into their level of understanding of course material. But as generative AI becomes more sophisticated, assignments like these are increasingly likely to reveal more about the efficacy of a prompt than the depth of knowledge that a student has acquired. To date, most of the conversation in higher education about the role of AI has to do with questions of academic integrity, authorship, and institutional policy. In an increasing number of contexts, students use a variety of tools to collect and analyse data outside of formal schooling and indeed outside of teacher control. The more immediate question for teachers and instructors is then how these tools can be used within the classroom so that students learn to evaluate and work through AI-generated outputs rather than simply relying on them uncritically (Holmes et al., 2019; Zawacki-Richter et al., 2019).

Of increasing interest as more educational tools incorporate AI into their functionalities, is the idea of AI literacy. Though there is still a lack of consensus when it comes to fully defining this term, it is generally considered to encompass knowledge of how AI systems operate, awareness of AI systems'

limitations, and a critical perspective when interacting with AI generated content (Ng et al., 2021). Discussions of AI literacy in academia tend to be separate from discipline specific coursework and typically encompass topics such as workshops, institutional policies and statements, and discussions regarding the proper use of AI within society. Little is made of what AI literacy looks like within the context of specific disciplines. Although students may learn about the flaws of AI and how it can hallucinate, or produce incorrect information, they still struggle to recognize flawed SQL queries, weak relational structures, or misleading outputs inside a database assignment.

One thing that became increasingly noticeable in technical coursework was how quickly students turned to AI systems for complete solutions once those systems became widely available. In SQL assignments, students could often generate entire blocks of code with only a short prompt. Sometimes the generated SQL worked correctly. Sometimes it looked correct but failed once executed in PhpMyAdmin. In other cases, the AI-generated tables ignored relational constraints entirely or produced JOIN statements that did not function as intended. Students who already understood relational databases could usually identify and fix those issues. Students with weaker conceptual understanding often accepted the generated output at face value. That distinction became important because it suggested that the central instructional issue was not simply whether AI was being used, but whether students retained enough disciplinary understanding to evaluate what the AI system produced (Luckin, 2025; Popenici & Kerr, 2017).

The instructional approach described in this paper emerged from redesign work within an undergraduate web development course focused on relational databases and Structured Query Language (SQL). Earlier versions of the assignment required students to manually create tables and write SQL statements independently. Once AI tools became available, instead of having students use them to replace the foundational tasks, an AI activity was created that builds on those foundational steps. Students complete tasks, document prompts that were used, execute the resulting SQL statements, identify why some failed and recursively refine the resulting SQL until optimal results are achieved based on relational database principles. While verification was supposed to happen after the assignment was submitted, it was actually one of the most valuable parts of the assignment. Students quickly discovered that plausible-looking AI-generated SQL often still contained structural or logical problems. Similar instructional concerns involving AI-supported learning and structured assignment integration have also been discussed in recent work examining AI literacy within technical coursework (Glotsbach, 2026).

While there is great value in fostering AI literacy, it is generally more meaningful when incorporated into coursework and discipline-specific instruction, rather than as an add-on skill. Using the redesigned SQL assignment as a concrete example, the discussion examines how structured assignment design can integrate AI into existing coursework while preserving disciplinary rigor. Particular attention is given to scaffolding, verification, prompt documentation, and productive struggle as instructional design principles that help keep conceptual understanding central to the assignment. As generative AI continues to influence higher education, the role of instructional design may become less about restricting AI use and more about structuring how students interact with these systems during learning.

2. AI literacy as an instructional design problem

One of the difficulties instructors currently face is a gap between how AI literacy is conceptualized and how it can be taught to students in higher education. While many workshops and conferences about AI literacy focus on concepts and policy, students are rarely given disciplinary-based assignments that make these concepts relevant. Awareness of issues such as ethics, bias, hallucinations, and responsible use are important for students to possess, but unless and until they are given the opportunity to practice critical thinking about these issues in the context of meaningful work, they will have little idea how to apply what they have learned. This disconnect suggests that AI literacy cannot be fully developed through isolated instruction alone. It must be situated within disciplinary practice where students are required to make decisions, evaluate outputs, and justify revisions using domain-specific knowledge (Ng et al., 2021; Zawacki-Richter et al., 2019).

This issue became increasingly visible in technical coursework where students had access to AI systems capable of producing complete programming solutions with minimal prompting. In database assignments involving SQL, students could quickly obtain functional CREATE, INSERT, SELECT, UPDATE, and DELETE statements without necessarily understanding why those statements worked or failed. In some cases, the generated code appeared correct on the surface but violated basic relational principles or produced unintended outcomes once executed in the database environment. Students who lacked foundational understanding often struggled to recognize those problems. The concern was therefore not simply whether students were using AI, but whether they retained enough disciplinary understanding to evaluate what the AI system produced. Similar concerns have emerged in broader discussions of

AI-supported learning environments, particularly regarding how AI systems may influence student autonomy, reasoning, and decision-making processes (Popenici & Kerr, 2017).

From an instructional design perspective, this creates a problem that extends beyond academic integrity. Most academic assignments are designed with the assumption that, in order to complete them, students will go through a cognitive process involving the information and skills taught in class. However, conversational generative AI models change this equation, as it is now possible to complete academic assignments that test skills and knowledge by simply asking a computer to produce the work. What is seen, for example a completed assignment, may not reflect what has been learned. And this is particularly troubling in technical disciplines, where simple errors in logic, syntax, or structure can have unapparent yet serious consequences. As AI grows in complexity, the way that AI is taught must evolve as well; moving beyond a simple focus on results and towards an emphasis on the underlying reasoning code that produces those results (Holmes et al., 2019; Luckin, 2025).

The redesign described in this paper emerged from the recognition that simply prohibiting AI use would not solve the larger instructional issue. This assignment redesign tackles the more fundamental pedagogical problem. Students are already using these tools outside of the classroom, experimenting with how to incorporate them into their work. Instead of ignoring or banning what students are doing with AI in their free time, this redesign brings the activity into the open and teaches students to use the technology intentionally while challenging them to utilize it appropriately and reflect on the distinctions between AI generated content and human content. Similar arguments have emerged in recent scholarship suggesting that AI integration requires movement beyond purely prohibition-based approaches toward models that emphasize ethical and intentional engagement with AI systems (Eaton, 2023). After executing their queries, students had to evaluate the results for correctness. They also had to critically evaluate and correct any incorrect SQL and justify their decisions. Students demonstrated both AI literacy and disciplinary reasoning.

This strategy for incorporating AI into the SQL assignment started from well-established educational practices for scaffolding and productive struggle. Students first complete a set of foundational tasks learning key relational database concepts. Then, students review, test and modify generated SQL statements, identifying errors and tuning the language used in the prompts to produce more complete and accurate AI code. Many of the most productive learning opportunities occur when the generated code fails to execute correctly. Students must then rely on their understanding of database concepts and query structure to identify the problem and correct it. Productive struggle remains part of the learning process even when AI is involved (Kapur, 2008; Merrill, 2002).

Viewing AI literacy as an instructional design problem rather than simply a technology problem changes how curriculum integration is approached. The emphasis shifts from teaching students about AI in the abstract to designing assignments that require meaningful interaction with AI systems inside disciplinary practice. In this framework, prompting, verification, revision, and reflection become embedded learning activities rather than supplemental skills. The assignment itself becomes the mechanism through which students develop AI literacy while continuing to build disciplinary competence.

3. Structured assignment integration

The assignment used in this redesign is part of an undergraduate web development course in which students learn foundational relational database concepts and Structured Query Language (SQL). Prior to the redesign, the assignment focused entirely on manual database construction and manipulation. Students created relational tables, inserted records, and wrote SELECT, UPDATE, and DELETE statements directly in PhpMyAdmin. The assignment was intentionally procedural. Students needed to understand how primary keys functioned, how relationships between tables were established, and how improperly scoped queries could affect stored data. Those foundational concepts remain central to the course and were not removed during the redesign process.

To incorporate AI-assisted work in this assignment, a third component was introduced to the current workflow; however, this component did not supersede the existing workflow. The traditional workflow of creating tables, filling them with data and manually performing a series of SQL commands to extract specific information to answer academic generated questions remained the same as the first two components of the assignment and were to be completed prior to the introduction of the additional component and the utilization of generative AI tools. The newly introduced component consequently expanded on the traditional workflow rather than substituting it. The problem of using AI to look up answers before demonstrating basic competence with a particular skill was avoided by moving the use of AI to after students had shown some grasp of a skill. Using AI this way allowed them to continue developing and refining their skills (Merrill, 2002).

In the third component of the assignment, students prompt an AI system to design an additional relational database table that meaningfully connects to the schema they previously created. The assignment requires students to specify constraints directly in their prompts, including primary keys, appropriate data types, and at least one foreign key relationship when relational connections are present. Students then prompt the AI system to generate SQL statements for creating the table and manipulating data through INSERT, UPDATE, DELETE, and SELECT operations. At least one relational JOIN statement must also be included when appropriate. Rather than treating the AI-generated SQL as a completed solution, students are required to execute each statement within PhpMyAdmin and verify that the queries function correctly within the actual database environment.

The verification step ended up changing the assignment more than expected. Students had to verify not only their own work, but also the SQL code generated by an artificial intelligence program for a variety of tables. Although the code looked good to the students, upon copying and pasting the code into PhpMyAdmin, a number of errors were encountered including relational errors, syntax errors and missing constraints. In some cases, the tables generated by the AI for tables that were meant to have a primary key, failed to declare a primary key. Other code had field definitions that were defined incorrectly within the overall schema. Students learned that simply copying and pasting from the AI program is not good enough in these assignments. The students are forced to understand the AI generated code, process it internally, find errors and correct them, all the while drawing upon prior knowledge of relational databases. This is the most instructionally valuable part of the assignment.

A second key aspect of this assignment is the documentation of the actual prompts written to the AI model, the resulting SQL code generated by the AI model, and the resulting final corrected and executable SQL code provides valuable insight into both the AI's process and the results of using it, as well as how prompting affects technical communication outcomes in terms of quality and organization. Well-structured prompts resulted in high quality, organized SQL code that adhered to relational database design principles.

The redesign also preserves disciplinary rigor by maintaining traditional evaluation standards. Students are not graded on whether they used AI effectively in a general sense. Instead, grading remains grounded in SQL competency and relational database accuracy. Queries must execute correctly, constraints must function properly, and relational integrity must be maintained throughout the schema. AI use does not exempt students from understanding the consequences of poorly scoped UPDATE statements or improperly defined keys. The database environment itself becomes a form of verification because MySQL either accepts or rejects the generated code based on structural validity. As a result, AI functions less as an authority and more as a drafting partner whose outputs require continual evaluation.

Embedding AI within the assignment in this structured way changes the instructional role of the technology. Rather than operating outside the learning process, AI becomes integrated into disciplinary practice under clearly defined constraints. Students engage with AI-generated code while still remaining responsible for verification, revision, and conceptual understanding. The assignment therefore treats AI literacy not as a separate competency, but as something developed through direct interaction with AI systems inside authentic technical coursework.

4. Implications for curriculum design

Rather than simply engaging in policy discussions, curriculum redesign can introduce students to AI in ways that afford deeper and nuanced engagement with the technology through activities specific to their major discipline. In this paper, we describe how a course activity was redesigned to address these goals and how engagement with AI within a design activity afforded students' use of AI in relation to their major discipline. Specifically, students in the redesigned activity engaged in design, testing, revising, and evaluating as they utilized AI tools. This integration helps preserve the connection between AI literacy and disciplinary thinking.

As students incorporate AI into their assignments, it is likely that the manner in which they utilize the technology will be shaped by the design of the assignment rather than by the features of the technology itself. In order to ensure that students utilizing AI remain engaged in disciplinary concepts and processes, it is crucial for instructors to design assignments that demand documentation of prompts, verification of reliability, repeated regeneration of content, and explanation of processes. While typical discussions about tools center on their rapid evolution, a more fruitful conversation might focus on how those tools are used, and the design principles (scaffolding, transparency, and productive struggle) that guide their use. (Kapur, 2008; Merrill, 2002).

The redesign also offers a practical model for faculty adoption. Many instructors remain hesitant to integrate AI into coursework because they perceive the process as requiring major curriculum redesign or entirely new course structures. This assignment demonstrates a more incremental approach. The foundational assignment remained largely intact while AI interaction was added through a structured extension activity. Existing learning objectives were preserved, and traditional disciplinary expectations continued to guide evaluation. Similar instructional approaches involving structured AI integration within technical coursework have also been explored in related work focused on AI literacy and assignment design (Glotzbach, 2026). For faculty members seeking manageable entry points into AI integration, modifying existing assignments may be more sustainable than large-scale curricular replacement.

Finally, this approach reinforces the idea that AI literacy is inseparable from disciplinary judgment. In the SQL assignment, students do not simply learn how to prompt an AI system. They learn how to evaluate relational structures, identify flawed outputs, and reconcile generated code with actual database behavior. The disciplinary concepts remain central throughout the interaction. This distinction is important because AI literacy risks becoming superficial when disconnected from authentic disciplinary practice. Embedding AI literacy into curriculum through structured assignment design helps ensure that students continue to develop the reasoning, verification, and evaluative skills necessary for meaningful learning in AI-supported environments.

5. Conclusion

Generative and agentic AI systems are increasingly influencing how students approach academic work, creating new challenges for instructional design in higher education. This paper argued that AI literacy should not be treated as a separate or supplemental competency, but instead embedded directly within disciplinary coursework through structured assignment design. Using a redesigned SQL assignment as a concrete example, the discussion illustrated how AI can be integrated into existing coursework while preserving disciplinary rigor and meaningful engagement with core concepts.

The assignment redesign positioned AI as a drafting and extension tool rather than a replacement for learning. Through sequencing, verification, prompt documentation, and revision, students remained responsible for evaluating and correcting AI-generated outputs using disciplinary knowledge. In this model, instructional design determines whether AI interaction supports conceptual understanding or substitutes for it. As AI systems continue to evolve, the challenge for educators will not simply be whether these tools are allowed in coursework, but how learning environments are designed around their use. Embedding AI literacy within authentic disciplinary practice offers one practical pathway forward for integrating emerging technologies while preserving the intellectual goals of higher education.

References

- Eaton, S. E. (2023). Postplagiarism: Transdisciplinary ethics and integrity in the age of artificial intelligence and neurotechnology. *International Journal for Educational Integrity*, 19(1), 23.
- Glotzbach, R. J. (2026). Embedding AI Literacy within Technical Coursework: An Instructional Design Approach. *Proceedings of the 37th international conference of the Society for Information Technology and Teacher Education (SITE 2026)* (pp. 1203-1207). Philadelphia, PA.
- Holmes, W., Bialik, M., & Fadel, C. (2019). Artificial intelligence in education. In M. Hildebrandt, J. W. B. G. Peters, & B. van den Hoven (Eds.), *Data ethics: Building trust* (pp. 619–653). Globethics.
- Kapur, M. (2008). Productive failure. *Cognition and Instruction*, 26(3), 379–424.
- Luckin, R. (2025). Nurturing human intelligence in the age of AI: rethinking education for the future. *Development and Learning in Organizations: An International Journal*, 39(1), 1-4.
- Merrill, M. D. (2002). First principles of instruction. *Educational Technology Research and Development*, 50(3), 43–59.
- Ng, D. T. K., Leung, J. K. L., Chu, S. K. W., & Qiao, M. S. (2021). Conceptualizing AI literacy: An exploratory review. *Computers and Education: Artificial Intelligence*, 2, 100041.
- Popenici, S. A. D., & Kerr, S. (2017). Exploring the impact of artificial intelligence on teaching and learning in higher education. *Research and Practice in Technology Enhanced Learning*, 12(1), 1–13.
- Zawacki-Richter, O., Marín, V. I., Bond, M., & Gouverneur, F. (2019). Systematic review of research on artificial intelligence applications in higher education. *International Journal of Educational Technology in Higher Education*, 16(1), 1–27.