A Retrieval-Augmented Generation Platform for Construction Education: Development and User Study

Xiaoyu Hou[1], Anusha Kannan[2], Shane Mueller[3] and Bo Xiao[4]

[1]PhD Student, Dept. of Civil, Environmental, and Geospatial Engineering, Michigan Technological Univ., Houghton, MI. 49931, U.S.A. Email: xiaoyuho@mtu.edu

[2]Master Student, Dept. of Psychology and Human Factors, Michigan Technological Univ., Houghton, MI. 49931, U.S.A. Email: anushaka@mtu.edu

[3]Professor, Dept. of Psychology and Human Factors, Michigan Technological Univ., Houghton, MI. 49931 Email: shanem@mtu.edu

[4]Assistant Professor, Dept. of Civil, Environmental, and Geospatial Engineering, Michigan Technological Univ., Houghton, MI. 49931, U.S.A. Email: boxiao@mtu.edu (Corresponding Author)

Abstract –

The continued digitalization of construction practice, including the expanded use of BIM, automation, and data-centric workflows, has increased the demands placed on construction management education. As curricula adapt to these developments, both instructors and students face growing challenges in organizing and interpreting course content during independent study. Traditional slide-based review materials provide limited support for individualized learning and offer few mechanisms to help students reconstruct conceptual relationships. Although general-purpose generative AI tools can provide on-demand explanations, they are not aligned with course-specific terminology and rarely offer verifiable links to instructional materials, which limits their suitability for guided review. Guided by constructivist perspectives on explanation and information integration, this study develops and evaluates a course-aligned review platform that uses Retrieval-Augmented Generation (RAG) to deliver contextualized explanations grounded in lecture slides. A formative user study was conducted in which participants viewed a common instructional video and completed a review session using either the AI-driven platform or conventional slides, followed by a user-experience questionnaire and a post-review assessment. The AI review condition received higher ratings on key UEQ dimensions, including Perspicuity, Efficiency, Dependability, and Attractiveness, and showed a modest advantage on explanatory performance. The findings indicate that domain-bounded, citation-aware generative AI can enhance transparency and support more interpretable review practices in construction education.

Keywords –

Construction education; Generative AI; BIM; RAG

1 Introduction

The construction sector continues to advance through the growing adoption of Building Information Modeling (BIM), immersive visualization technologies, automation tools, and data-driven management practices [1]. These developments are reshaping how projects are analyzed and coordinated, and they broaden the technical competencies expected of future professionals. In educational settings, this evolving knowledge is typically introduced through structured instructional materials, such as lecture slides, specifications, and supporting documents that represent key aspects of project information [2, 3]. As construction processes become more computationally mediated and interdisciplinary, the associated knowledge base expands in scope and complexity, placing heightened expectations on higher-education programs responsible for cultivating digital project literacy.

The increasing integration of digital tools into construction practice has also influenced the structure of construction management curricula. Prior studies report that course content in this domain often brings together project information, technical procedures, and software-related representations, which students must connect to understand construction workflows [4, 5]. In higher education, lecture slides remain a primary medium for organizing and delivering such materials, yet research on slideshow-based instruction has noted that PowerPoint-style presentations often segment content and provide limited support for showing relationships across concepts, procedures, and project context [6]. Studies of BIM-enabled construction education further indicate that students have trouble during self-directed review, particularly when course materials are closely tied to specialized terminology and representation conventions [7]. These findings point to a need for learning support that is grounded in course materials and better supports students in tracing conceptual connections across technical content during independent study.

Generative AI technologies have attracted increasing attention for their ability to generate explanations, summarize instructional materials, and respond to learner queries. Studies suggest that such tools can support feedback and conceptual understanding in educational settings [8]. However, their application in construction management education remains limited. Existing models are typically trained on open-domain corpora and lack alignment with terminology, course context, and instructional expectations of construction coursework [9]. They also seldom provide traceable citations to course materials, making it difficult to verify the source and accuracy of generated explanations [10]. In addition, most existing RAG-based tools operate at the document level, generating responses from individual files in isolation. In construction management education, course materials are organized across lectures and topics, and effective review requires following how concepts and procedures develop across this structure. This creates a need for course-aligned learning support that supports connected understanding across instructional materials.

The objective of this study is to develop and evaluate a course-aligned learning platform that employs Retrieval-Augmented Generation (RAG) to support students' interpretation of construction management materials across multiple courses. The platform provides explanations grounded in lecture slides and supports learning within and across course contexts. Its effectiveness was investigated through a formative user study in which participants viewed a common instructional video and subsequently completed a structured review session using either the AI-driven platform or conventional slides. Usability was assessed using a standardized instrument, and learning outcomes were evaluated through multiple-choice and open-ended measures. The findings contribute to ongoing efforts to integrate generative AI into construction education by supporting course-level review and helping students follow relationships across instructional materials.

2 Literature Review

2.1 Generative AI in Intelligent Tutoring

Generative AI has expanded the capabilities of intelligent tutoring systems by enabling automated content synthesis, real-time question answering, and personalized explanation. Unlike earlier AI applications focused primarily on retrieval or classification, contemporary generative models integrate multimodal reasoning and synthesis to address complex instructional needs [11]. In writing-intensive and design-related learning environments, such systems help articulate abstract ideas, refine conceptual reasoning, and support iterative development through prompt-based feedback [12, 13]. Text-to-image models further enhance visual thinking by transforming descriptive prompts into conceptual renderings, complementing early-stage ideation processes [14].

A key advance of modern intelligent tutoring is the combination of vector-based retrieval and generative synthesis, enabling context-specific explanations grounded in technical documents or uploaded instructional materials [15, 16]. This integration, commonly implemented through Retrieval-Augmented Generation (RAG), constrains generative outputs to relevant sources while preserving the flexibility of natural-language explanation. This dual mechanism supports learners who struggle with domain-specific terminology or abstract construction concepts by bridging textual complexity with accessible explanation [17]. The architecture typically involves document ingestion, embedding-based retrieval, and generative response modeling, which together facilitate scalable, document-grounded tutoring workflows applicable to construction education scenarios.

Generative AI has also supported computational learning tasks by translating natural-language instructions into executable scripts, assisting with debugging, and guiding algorithmic reasoning across tools such as Python and parametric modeling platforms [18]. These capabilities lower the barrier for procedural tasks and accelerate iterative testing and refinement within computational environments [19].

However, most implementations remain preliminary and depend on general-purpose tools that lack alignment with course-specific materials. This suggests a need for domain-tailored, document-grounded tutoring systems capable of supporting the conceptual demands of construction education.

2.2 Evaluation of AI Learning Support

Evaluations of AI-driven learning support systems commonly focus on two complementary dimensions: usability and learning effectiveness. Prior work across higher education settings shows that AI-based learning assistance influences how students engage with course materials and how they perform on related assessments. Large-scale empirical studies indicate that learners frequently adopt AI tools in professional and technical coursework, reporting generally positive experiences with their usability and instructional value [20]. However, these findings also reveal variability in the accuracy and contextual alignment of AI-generated responses, underscoring the need to examine both user perceptions and task performance when integrating AI into domain-specific instruction [21].

Usability evaluations are often conducted to determine whether AI systems can be incorporated into learning activities. Research on AI-supported feedback environments demonstrates that well-designed systems can enhance the quality of learner output while achieving favorable usability ratings [22]. Such evaluations typically include measures of perceived ease of use, clarity of information, and integration with existing workflows, reflecting broader practices in educational technology research [23]. At the same time, studies note that users may express reservations about trust, transparency, or long-term instructional implications, suggesting that usability encompasses both functional and affective components [24].

Performance-based evaluations examine whether AI-supported learning leads to measurable improvements in task outcomes. Research on AI-enabled assessment and automated evaluation systems reports a wide range of performance metrics, including accuracy, precision, recall, and other task-specific indicators [25]. Although these systems differ from learning support tools in purpose, they reflect a consistent emphasis on empirical validation and structured performance measurement. Similar approaches have been applied in studies assessing AI tutors or writing assistants, where learning outcomes are measured through conceptual tests, analytical tasks, or open-ended explanations [26].

However, evaluations grounded in discipline-specific materials and conducted under domain-authentic conditions remain limited. This gap is particularly notable in technical fields such as construction education, where the complexity of instructional artifacts and specific learning objectives may shape both the perceived usefulness of AI support and its impact on learning.

3 System Design and Architecture

This study presents a RAG platform that delivers explanations aligned with course documents to support learning in construction education. The system converts course materials into an interactive environment in which students ask natural-language questions and receive responses anchored to specific instructional sources. As illustrated in Figure 1, the architecture combines document preprocessing, semantic indexing, retrieval, prompt assembly, and citation-aware answer generation into a unified workflow that enables verifiable, material-aligned feedback during learning activities.

3.1 Core System Architecture

The platform is built on a RAG designed to provide citation-linked explanations aligned with instructional materials. As shown in figure 1, the system operates through three coordinated layers that together support document-grounded interaction. The first layer processes course artifacts and prepares them for retrieval. The platform accommodates multiple file formats, including Word documents, PowerPoint slides, and PDF files. Uploaded documents are parsed, cleaned, and segmented into semantically coherent units that retain metadata indicating their position within the original source.

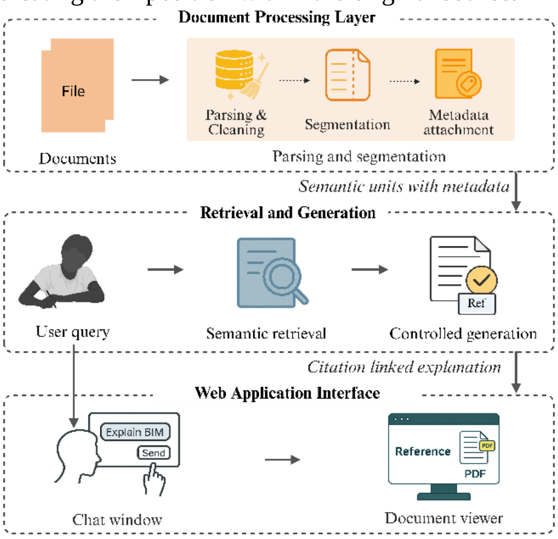

Figure 1. System architecture of the RAG-based instructional support platform

The second layer performs semantic retrieval and controlled generation. User queries are converted into vector representations and matched against the indexed content to identify document fragments that reflect the underlying intent of the question. These fragments are assembled into a structured input for the language model, which is instructed to generate explanations solely based on retrieved material. This design constrains the model to remain within course boundaries and ensures that generated responses are grounded in verifiable instructional sources.

The third layer presents explanations through a web-based interface that emphasizes transparency and interpretability. Retrieved fragments are incorporated into the output together with their metadata so that users can inspect the corresponding portions of the original documents while reviewing the explanation. This integrated architectural design supports document-grounded reasoning and facilitates a learning experience in which students can navigate course materials with greater clarity and contextual awareness.

3.2 System Workflow

The operation of the platform follows a structured workflow that links user queries to document-grounded explanations in a transparent manner. The workflow begins with query submission through the web interface.

When a student enters a question, the system converts the input into a semantic representation using the same embedding model applied to the instructional corpus. This alignment allows the system to compare the query with the indexed materials based on conceptual rather than lexical similarity.

The query is matched against the repository of instructional fragments to identify content that reflects the intent of the question. Semantic proximity is evaluated using cosine similarity, expressed as

$$\text{Sim}\ (q, d_i) = \frac{q \cdot d_i}{\|q\|\|d_i\|} \qquad (1)$$

where q denotes the query embedding and d_i represents the embedding of a document fragment. Fragments with higher similarity values are considered more relevant and are therefore selected for subsequent processing. This retrieval step ensures that the model operates on content closely aligned with the learner's information needs while preserving the structure of instructional materials.

After retrieval, the selected fragments are assembled into a structured prompt that guides the generative model to produce an explanation that remains within the boundaries of the supplied material. The prompt incorporates both the learner's question and the retrieved content, and the model is instructed to rely exclusively on this information when forming its response. In this study, the prompt takes a general form such as: "You are a helpful assistant. User Query: {user_query}. References: {reference_block}. Provide an answer to the query using only the information contained in the references." In this template, {user_query} denotes the learner's submitted question, and {reference_block} represents the retrieved instructional fragments together with their metadata, including the document name and page indicator. During execution, these placeholders are replaced with the actual query and the corresponding retrieved segments, ensuring that the generated explanation remains grounded in the instructional materials. This controlled prompting strategy reduces unsupported inferences and helps maintain conceptual and terminological alignment with course content.

The system then generates an explanation accompanied by citations that reference the underlying document segments. These citations provide students with a direct pathway to the source material, enabling them to verify the origin of each component of the explanation. The final output is presented through the web interface, where students can review both the synthesized response and the associated source excerpts in a single environment. This workflow supports targeted exploration of course content and encourages learners to connect explanations with the instructional context from which they are derived.

3.3 Web Application Interface

The platform is accessed through a web-based interface that supports interaction with the RAG system. The application provides two main functions designed to facilitate question submission and source-linked explanation review.

The first function is a chat-style interface that allows users to submit queries related to course materials in natural language. Generated explanations are displayed within the same window, creating a continuous interaction flow that supports iterative question refinement. The interface is designed to present responses clearly and to maintain conversational context without introducing unnecessary visual complexity. The second function supports citation-based navigation of instructional materials. When an explanation includes a citation, selecting it opens the corresponding portion of the source document in a parallel viewing pane. The citation display includes the document name and page indicator, enabling users to verify the origin of the retrieved content while examining the generated explanation. Presenting the explanation and its supporting material side by side enhances interpretability and reinforces the connection between the system's output and the instructional documents.

3.4 Implementation Summary

The implementation of the platform involves three primary components: preparation of the instructional corpus, backend retrieval and generation services, and a browser-based client interface. All instructional materials were consolidated into PDF format and processed to support semantic retrieval. The documents were parsed page by page, and the extracted text was segmented into fixed-length units. Each unit was stored together with metadata identifying the document name and page index. The processed corpus was uploaded to Azure Blob Storage, and a Shared Access Signature was assigned to each segment to enable secure citation links within system outputs.

The backend was developed in Python using the Flask framework and integrates several Azure services to support retrieval-augmented generation. Azure Cognitive Search indexes the instructional corpus and performs vector-based semantic retrieval. User queries are encoded through an Azure OpenAI deployment of the text-embedding-3-large model, and the most relevant document segments are retrieved from the index. These segments, together with their associated metadata, are assembled into a structured prompt that instructs a deployed instance of GPT-4o to generate an explanation grounded in the retrieved material. If no segment meets the retrieval threshold, a fallback prompt is issued to prevent unsupported claims and avoid fabricated citations. The entire backend is hosted on Azure App Service to ensure stable and scalable performance.

The client interface is implemented using HTML, CSS, and JavaScript and communicates with the backend through RESTful endpoints. The interface presents a query window and a dynamic answer display, along with a citation viewer that renders the referenced document snippet beside the generated response. Figure 2 illustrates the website interface.

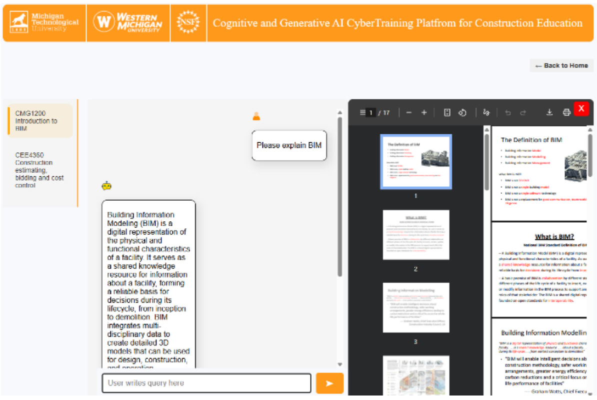

Figure 2. Web interface of the proposed platform

4 Experiment and Results

4.1 Experiment Design

The experiment was implemented as an online study delivered through the Qualtrics platform [27]. Participants were undergraduate students at Michigan Technological University, and only individuals without prior experience in construction, BIM, or related domains were recruited to control baseline knowledge. The study received Institutional Review Board approval, and informed consent was obtained electronically at the beginning of the session. A total of thirty students completed the study and were randomly assigned in equal numbers to two review conditions. One group reviewed the instructional content using the course-aligned retrieval-augmented platform developed in this study, whereas the comparison group examined the same material using a conventional slide deck.

All participants proceeded through a fixed sequence of activities within Qualtrics. The session began with an introduction and consent procedure, followed by a ten-minute recorded lecture derived from the instructional slides. After viewing the lecture, students entered the review phase corresponding to their assigned condition. The duration of this phase was held constantly to ensure comparable exposure time across groups. Upon completing the review activity, participants proceeded to two assessment components designed to capture perceptions of the assigned review tool and understanding of the lecture content.

User experience was evaluated using the short version of the User Experience Questionnaire (UEQ) [28], which provides standardized scales for pragmatic and hedonic interaction qualities. Learning performance was assessed through a post-review test that included four multiple-choice questions and two open-ended explanation prompts aligned with the lecture material. Each multiple-choice item contributed one point, and open-ended responses were scored with a five-point analytic rubric, yielding a maximum possible score of fourteen. Two trained raters independently scored the open-ended responses and resolved discrepancies through discussion to ensure consistent application of the rubric.

4.2 Experiment Results

User-experience ratings and post-review test scores were analyzed to compare the two review conditions. User-experience outcomes are reported according to the structure of the User Experience Questionnaire (UEQ), whose conceptual model distinguishes six dimensions of interaction quality. Attractiveness captures the overall impression of the system; Perspicuity reflects how easily users can become familiar with and understand the interface; Efficiency addresses the degree to which tasks can be completed without unnecessary effort; Dependability concerns users' sense of control and the predictability of system behavior; Stimulation captures the degree to which the interaction is perceived as motivating or exciting; and Novelty reflects the perceived innovativeness of the system. Each dimension is measured on a –3 to +3 semantic differential scale, with positive values indicating more favorable evaluations. These dimensions together form a comprehensive assessment of pragmatic and hedonic qualities and are therefore used to structure the reporting of user-experience results in this study.

Table 1 presents the descriptive statistics for the six UEQ scales across the two review conditions, including mean values and standard deviations. Mean ratings were positive in both groups, with the AI review condition showing higher values on pragmatic-oriented scales (Perspicuity, Efficiency, Dependability) as well as Attractiveness, while differences in Stimulation and Novelty were comparatively limited.

Statistical comparisons were conducted using independent-samples t-tests. The results indicate that Perspicuity showed a statistically significant difference between the two conditions ($p < 0.01$), whereas differences in other dimensions did not reach statistical significance. These results suggest that the AI platform improves clarity and ease of understanding, while differences in other aspects remain moderate.

Table 1. UEQ comparison with statistical analysis

Dimensions	PPT group Mean	PPT group SD	AI group Mean	AI group SD	p
Attractiveness	0.17	1.12	0.56	1.08	0.31
Perspicuity	0.20	1.25	1.50	1.10	0.004
Efficiency	0.53	1.18	1.05	1.02	0.16
Dependability	0.85	1.05	0.90	0.98	0.89
Stimulation	-0.32	1.30	0.37	1.22	0.12
Novelty	-0.32	1.28	-0.17	1.15	0.70

Learning performance results are summarized in Table 2. Multiple-choice accuracy (maximum four points) was similar across conditions, suggesting comparable recognition-level understanding of the lecture material. For the open-ended questions (maximum ten points), consensus rubric scores showed a slight advantage for the AI review condition; however, the overall difference in total test scores (maximum fourteen points) was small. Inter-rater agreement was high prior to reconciliation, confirming the consistency of the scoring process.

Table 2. Post-review learning performance

Questions		PPT group	AI platform group
Multiple-choice	1	0.93	1.00
	2	0.87	0.93
	3	0.53	0.47
	4	1.00	0.93
Open-ended question	5	3.40	3.47
	6	3.00	3.73
Avg. score		9.73	10.53

5 Discussion

5.1 Interpretation of User-Experience

Across the UEQ dimensions, the observed differences are most evident in the pragmatic aspects of interaction, particularly those related to clarity and task-oriented use. This pattern suggests that the platform supports a more direct and structured review process, allowing students to engage with course content in a focused manner. By enabling question-driven interaction and linking responses to specific instructional materials, the system reduces the need for manual navigation across slides. This may facilitate more efficient access to relevant information and support the interpretation of key concepts during review.

In contrast, the limited variation observed in Stimulation and Novelty indicates that affective responses were relatively stable across conditions. This outcome is consistent with short-duration interventions and the widespread familiarity of students with AI-based tools, both of which can constrain perceived differences in engagement or novelty.

With respect to learning performance, the small differences between conditions suggest that both review formats were effective in supporting basic understanding. The slightly higher performance observed in open-ended responses under the AI condition may reflect differences in how students organize and express their understanding, rather than substantial differences in content acquisition.

5.2 Interpretation of Learning Performance

Scores on both the multiple-choice and open-ended items show a modest advantage for students who reviewed with the AI-driven platform. The difference in multiple-choice performance was limited, suggesting that both review conditions supported short-term recall of the lecture content. A clearer difference appeared in the open-ended responses, where students in the AI condition more often produced answers that were fuller in scope and better organized.

This contrast reflects the different demands of the two assessment formats. Multiple-choice items capture recognition or recall, whereas open-ended questions require students to organize and express their understanding. The stronger performance of the AI group therefore appears to relate to how students structured their responses rather than to differences in factual recall. One possible explanation is that the platform enabled more targeted review through question-driven interaction, allowing students to clarify key concepts during the review process. In the slide-based condition, students relied on self-directed navigation to locate and assemble relevant information within the same time constraint.

The learning results do not suggest a large separation in basic content retention between the two conditions. The more noticeable difference lies in explanatory performance, where the AI-supported review condition appears to have better supported the construction of clearer and more complete written responses.

5.3 Limitations

Several limitations should be acknowledged regarding both the platform and the experimental design. At the system level, the current implementation is restricted to lecture-slide materials and does not yet incorporate other artifacts central to construction education, such as drawings, models, schedules, or specifications. As a result, the platform does not fully represent the multimodal information environment in which construction-related learning typically occurs. In addition, the retrieval and explanation mechanisms have not been evaluated in scenarios that require multi-step reasoning or integration across multiple document types, which limits conclusions about the platform's applicability to more advanced instructional tasks.

At the empirical level, the study involved modest sample size and a single short review session, constraining the detection of longer-term or cumulative learning effects. Participants were students without prior exposure to construction or BIM-related topics; while this controlled for differences in initial knowledge, it may not reflect how domain-aware learners interact with course-aligned AI tools. The assessment also focused on basic conceptual understanding and brief explanatory responses, leaving open the question of how the platform supports higher-order competencies. Finally, all interactions occurred within the structured environment of an online survey instrument, which differs from classroom or project-based settings where review behaviors may be more varied. These considerations should be considered when interpreting the findings.

5.4 Future Research

Future work may advance the platform along several directions. At the system level, extending retrieval and explanation capabilities beyond lecture slides to encompass drawings, BIM models, schedules, and specifications would more fully reflect the multimodal information landscape of construction education. Incorporating mechanisms for cross-document reasoning and multi-step explanation would also enable the platform to support instructional tasks that require synthesizing procedural and representational knowledge across artifacts.

From an educational perspective, longer-term classroom deployments are needed to examine how course-aligned AI support influences learning when integrated into recurring study routines, project work, or module-level instruction. Studies involving learners with prior exposure to construction or BIM would provide additional insight into how domain knowledge shapes the use and effectiveness of AI-driven review tools. Further work could also investigate the platform's role in supporting higher-order competencies, such as coordinating design intent, interpreting model-based relationships, or evaluating alternative construction strategies.

Methodologically, future studies may benefit from incorporating process-level data, including interaction logs, query patterns, and temporal measures of review behavior. Such analyses would allow finer-grained characterization of how students engage with AI-generated explanations and how these behaviors relate to performance outcomes. Comparative evaluations across different configurations of generative AI, including variations in retrieval granularity or explanation constraints, may also help clarify which design features most effectively promote reliable, transparent, and pedagogically aligned learning support.

6 Conclusion

This study explored how generative AI, supported by a Retrieval-Augmented Generation framework, can be integrated into construction management education to improve the accessibility, relevance, and traceability of learning resources. By combining domain-bounded retrieval with transparent citation mechanisms, the proposed approach offers a structured pathway through which students can engage with course materials in alignment with instructor-approved content while supporting the active construction of knowledge, consistent with constructivist perspectives, by helping learners relate retrieved information to their developing conceptual understanding.

The analysis of opportunities and challenges indicates that the effective use of generative AI in educational settings depends on appropriate model constraints, reliable data sources, and thoughtful pedagogical integration. When embedded within a controlled and well-designed instructional environment, generative AI has the potential to enrich learners' engagement with construction management concepts and to provide more interpretable pathways for reviewing complex materials. Future research may extend this work by examining the framework with broader learner populations, incorporating additional instructional artifacts such as drawings and models, and evaluating its influence on higher-order learning outcomes across diverse curricular contexts.

References

[1] G. Mahajan and P. Narkhede. Integrating BIM with digital technology trends in the construction industry: Implementation insights for 2023. Library of Progress – Library Science, *Information Technology & Computer*, 44(3), 2024.

[2] X. Hou, C. Li, and Q. Fang. Computer vision-based safety risk computing and visualization on construction sites. *Automation in Construction*, 156:105129, 2023.

[3] N. Rane. Integrating leading-edge artificial intelligence (AI), internet of things (IoT), and big data technologies for smart and sustainable AEC industry: Challenges and future directions, September 24, 2023. *Engineering and Construction (AEC) Industry: Challenges and Future Directions*.

[4] T. Olowa, E. Witt, and I. Lill. Building information modelling (BIM)–enabled construction education: Teaching project cash flow concepts. *International Journal of Construction Management*, 23(9):1494–1505, 2023.

[5] R. Elias, R. R. A. Issa, and W. Wu. Progress on building information modeling education and talent

acquisition. *International Journal of Construction Education and Research*, 19(4):363–382, 2023.

[6] C. M. Ferreira, A. I. Santos, and S. Serpa. Electronic slideshow presentations in the higher education teaching and learning process. *Journal of Education and Training Studies*, 6(2):120–124, 2018.

[7] J. Jin, K.-E. Hwang, and I. Kim. A study on the constructivism learning method for BIM/IPD collaboration education. *Applied Sciences*, 10(15):5169, 2020.

[8] A. Irons and S. Elkington. Enhancing learning through formative assessment and feedback, volume 1. *Routledge*, 2021.

[9] A. Jamal. Automatic construction of ontology with public-domain datasets for personalized tutoring with ECA, Master's thesis, *University of Windsor*, 2024.

[10] R. Kharipova, I. Khaydarov, S. Akramova, D. Lutfullaeva, S. Saidov, A. Erkinov, and N. Erkinova. The role of artificial intelligence technologies in evaluating the veracity of scientific research. *Journal of Internet Services and Information Security*, 14(4):554–568, 2024.

[11] R. C. Sharma and A. Bozkurt. Transforming education with generative AI: Prompt engineering and synthetic content creation. IGI Global, 2024.

[12] M.-L. Tsai, C. W. Ong, and C.-L. Chen. Exploring the use of large language models (LLMs) in chemical engineering education: Building core course problem models with Chat-GPT. *Education for Chemical Engineers*, 44:71–95, 2023.

[13] A. Giretti, D. Durmus, M. Lemma, M. Vaccarini, M. Zambelli, A. Guidi, and F. Ripa di Meana. Knowledge engagement in art and design education: About the role of AI in creativity education. *In Artificial Intelligence for Supporting Human Cognition and Exploratory Learning in the Digital Age*, pages 3–24. Springer, 2024.

[14] Z. M. Albaghajati, D. M. Bettaieb, and R. B. Malek. Exploring text-to-image application in architectural design: Insights and implications. *Architecture*, Structures and Construction, 3(4):475–497, 2023.

[15] M. O. Chimezie. Leveraging retrieval-augmented generation in large language models for effective learning: A data structures & algorithms learning assistant. *Ashesi University*, 2024.

[16] G. D. Fatouros. Design, optimization, development and explanation of artificial intelligence algorithms, capable of running on distributed big data systems, based on multidimensional and complex datasets. 2024.

[17] L. Wong. Gaita: A RAG system for personalized computer science education. 2024.

[18] L. Sass and R. Oxman. Materializing design: The implications of rapid prototyping in digital design. *Design Studies*, 27(3):325–355, 2006.

[19] A. Nag, J. Boricha, and A. Sarkar. Architectural education: AI integrative challenges and prospects. *In Effective Instructional Design Informed by AI*, pages 383–408, 2025.

[20] Sun, Y., Yang, H., Yu, H. K., & Suen, R. (2025). Boon or Bane? Evaluating AI-driven learning assistance in higher education professional coursework. *Education and Information Technologies*, 1-34.

[21] N. Jiang, W. Zhou, S. Hasanzadeh, and V. G. Duffy. ExpertGen: A comparative analysis of user performance, cognitive workload, and trust in domain-tailored generative AI. *In Proceedings of the International Conference on Human-Computer Interaction*, pages 178–190, 2025.

[22] Pishtari, G., Sarmiento-Márquez, E. M., Rodríguez-Triana, M. J., Wagner, M., & Ley, T. (2023, August). Evaluating the impact and usability of an AI-Driven feedback system for learning design. *In European Conference on Technology Enhanced Learning* (pp. 324-338). Cham: *Springer Nature Switzerland*.

[23] M. R. Boland, A. Rusanov, Y. So, C. Lopez-Jimenez, L. Busacca, R. C. Steinman, and C. Weng. From expert-derived user needs to user-perceived ease of use and usefulness: A two-phase mixed-methods evaluation framework. *Journal of Biomedical Informatics*, 52:141–150, 2014.

[24] M. Al-Kfairy, A. Alomari, M. Al-Bashayreh, O. Alfandi, and M. Tubishat. Unveiling the metaverse: A survey of user perceptions and the impact of usability, social influence and interoperability. *Heliyon*, 10(10), 2024.

[25] Chinnasamy, P., Pushpalatha, M., Nirmala, V., Ayyasamy, R. K., & Priya, D. (2024, December). AI-Driven Assessment and Evaluation: Enhancing Sustainability in Educational Practices. *In 2024 3rd International Conference on Automation, Computing and Renewable Systems (ICACRS) (pp. 1319-1325). IEEE.*

[26] V. J. Owan, K. B. Abang, D. O. Idika, E. O. Etta, and B. A. Bassey. Exploring the potential of artificial intelligence tools in educational measurement and assessment. *Eurasia Journal of Mathematics, Science and Technology Education*, 19(8):2307, 2023.

[27] Qualtrics. Understand every customer. Act when it counts. On-line: https://www.qualtrics.com/, Accessed: 12/07/2025.

[28] M. Schrepp. User Experience Questionnaire Handbook. On-line: https://www.ueq-online.org/Material/Handbook.pdf, Accessed: 02/09/2025.