

A Generative AI-Based Construction Safety Assistant Using Retrieval-Augmented Generation

Xiaoyu Hou¹, Aakash Walavalkar² and Bo Xiao³

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

²Master Student, Dept. of Health Informatics, Michigan Technological Univ., Houghton, MI. 49931, U.S.A. Email: awalaval@mtu.edu

³Assistant Professor, Dept. of Civil, Environmental, and Geospatial Engineering, Michigan Technological Univ., Houghton, MI. 49931, U.S.A. Email: boxiao@mtu.edu (Corresponding Author)

Abstract –

Safety remains a central challenge in the construction industry, where OSHA standards define most regulatory requirements. Although comprehensive, these standards are dispersed across lengthy documents, making it difficult to locate provisions relevant to specific tasks or safety concerns. This complexity is especially evident in education and training settings, where instructors must interpret regulatory text and learners often lack familiarity with its structure and terminology. Recent advances in generative Artificial Intelligence (AI) offer new opportunities for improving access to technical information, but general-purpose models lack the grounding needed for authoritative regulations. This study presents a Retrieval-Augmented Generation (RAG) platform that provides clause-linked answers based on OSHA construction standards. The platform integrates document preprocessing, semantic retrieval, and citation-aware response generation through a web interface that enables direct inspection of referenced clauses. Automated evaluation using thirty representative safety questions showed strong performance, with average scores of 0.95 for context relevance, 0.99 for factual accuracy, and 0.90 for response relevance. The study contributes a verifiable and regulation-aligned approach for delivering accessible safety guidance, offering an adaptable framework for construction safety education and training.

Keywords-

Construction safety; OSHA; RAG; ISARC

1 Introduction

In 2023, the United States construction industry recorded 1,075 fatal work injuries, which represented

nearly one fifth of all occupational deaths nationwide [1]. The consistently high number of fatalities highlights the persistent challenges associated with managing safety in an inherently hazardous sector [2]. Regulatory requirements issued under the Occupational Safety and Health Administration (OSHA) [3] construction standards therefore serve a central role in guiding hazard prevention across activities such as fall protection, scaffolding, electrical installation, excavation, equipment operation, and hazard communication. The standards offer comprehensive coverage of the risks commonly present on construction sites, yet the regulations are distributed across numerous subparts and contain extensive cross-referenced provisions [4]. The scale and organization of this regulatory corpus often make it difficult for workers, engineers, and students to identify the specific clauses relevant to a particular task or condition. Locating applicable requirements within lengthy documents requires substantial manual effort, which can hinder timely decision making and reduce the effectiveness of safety management practices.

Currently, users depend on keyword searches and personal familiarity with the structure of OSHA's construction regulations to locate relevant provisions, a process that often yields partial retrieval and inconsistent interpretation among individuals with limited regulatory experience [5]. The difficulty of identifying applicable clauses is evident across workforce training, onboarding programs, continuing professional development, and university instruction, where learners must connect specific safety questions to authoritative requirements without a clear mechanism for navigating the structure of the standards. Existing digital resources, including OSHA's eTool libraries and CPWR training materials [6, 7], improve access to regulatory information but remain primarily document centered and require substantial manual navigation. The absence of tools that can align user queries with the semantic content of regulatory text limits the role of OSHA standards as a foundation for

accurate, transferable, and instructionally meaningful safety knowledge. These limitations create a pressing need for a system capable of guiding users from practical safety questions to the corresponding OSHA clauses while providing transparent explanations that support both learning and decision making.

Advances in generative AI have created new opportunities for addressing the challenges associated with linking safety questions to relevant regulatory content. Recent developments in large language models have demonstrated strong capabilities in complex question answering, context interpretation, and the generation of instructional explanations, suggesting potential value for safety education and training [8]. Retrieval-Augmented Generation (RAG) has become particularly promising for applications that require the integration of authoritative sources with model-generated responses [9], since retrieved text can anchor explanations in verifiable regulatory clauses. Several studies have applied retrieval-based generative approaches to tasks such as building code interpretation, engineering documentation support, and domain-specific instructional assistance [10]. However, research focused on OSHA construction safety remains limited. Existing efforts rarely incorporate clause-level grounding in OSHA's construction regulations or provide mechanisms that align user queries with specific provisions in a manner suitable for education and training. The limited attention to OSHA-centered generative systems indicates the need for a solution capable of connecting user questions to regulatory clauses while delivering explanations that support accurate, transparent, and instructionally meaningful safety learning.

This study investigates how a retrieval-augmented generative AI framework can support accurate and transparent access to OSHA construction safety regulations through clause-grounded retrieval and explanation. To address this problem, a retrieval-augmented framework was developed to ground generative AI responses in authoritative OSHA construction regulations. The framework integrates document retrieval and controlled response generation to provide clause-level traceability and transparent evidence for each answer. A web-based platform was implemented to operationalize the framework and to present both the generated explanation and the corresponding regulatory text in an accessible interface. The platform was evaluated using automated metrics, including relevance, factual alignment, and citation correctness, to examine the extent to which system responses remain consistent with regulatory content. The results demonstrate that the proposed approach can effectively connect user questions with verifiable OSHA clauses while maintaining transparency in the response generation process.

2 Literature Review

This section reviews the state-of-the-art literature related to generative AI in construction education, which includes technologies in construction education, generative AI in intelligent tutoring, and generative AI in domain-specific education.

2.1 Construction Safety Management

Construction safety management has been widely examined due to the persistently high incidence of injuries and fatalities in the sector. Recent studies emphasize that structured safety management systems and systematic risk assessment play a critical role in reducing incident likelihood. Sadeghi et al. demonstrated that ensemble risk-assessment models can improve the identification of high-risk situations on construction projects [11], while Zhang and Mohandes developed a holistic occupational health and safety framework that links risk control measures to project-level decision making [12]. More recent work has shown that targeted management interventions and safety management measures can have measurable effects on workers' safety behaviors and incident outcomes [13, 14].

OSHA construction standards constitute the primary regulatory foundation governing hazard control in the United States and are frequently cited as the formal basis for site-level safety practices. However, prior research has indicated that regulatory requirements are difficult to explain and apply in practice. Semantic analyses of construction safety knowledge highlight that relevant information is highly fragmented and embedded in complex technical language, which complicates efforts to translate formal requirements into actionable guidance [15, 16]. Studies on construction safety knowledge systems similarly note that safety-related concepts are dispersed across documents and topics, and that workers often rely on local practices and informal guidance rather than directly consulting formal regulations [15]. Although digital resources, such as OSHA e-tools and online training portals, have broadened access to regulatory information, existing platforms still require substantial manual navigation and provide limited semantic support for locating clause-level requirements within lengthy regulatory documents [13].

Challenges related to regulatory comprehension also affect safety education and workforce training. Survey-based studies report that trainees and frontline workers frequently struggle to interpret technical terminology and to connect safety requirements with realistic site scenarios, which constrains the transfer of formal knowledge to field practice [17]. Eye-tracking and experimental training studies further show that limited safety knowledge undermines hazard-identification performance, even when workers have access to visual

examples and instructional materials [18]. The challenges highlight a continuing need for approaches that improve the accessibility and contextual alignment of safety regulations within both education and workforce training.

2.2 Generative AI Applications in Construction

Recent research has begun to explore how generative artificial intelligence can be incorporated into construction management, reflecting the broader digital transformation underway in the sector. Ghimire, Kim, and Acharya (2023) offered an early synthesis of the opportunities and challenges associated with generative AI [19], noting the potential to automate documentation, summarize technical information, and support managerial decision-making. Their analysis also emphasized that the rapid emergence of large language models has created new expectations for access to timely and interpretable information, while simultaneously raising concerns regarding reliability, data governance, and the need for domain-sensitive safeguards in safety-critical applications.

Industry-focused investigations provide complementary evidence regarding the growing relevance of generative AI. Nyqvist et al. (2024) examined how generative AI tools can be aligned with common construction management tasks and reported substantial perceived value in communication support [20], meeting administration, situational awareness, data analysis, and risk monitoring. Their evaluation, conducted through focus group discussions and practitioner surveys, suggests that generative AI can strengthen several routine managerial functions, although its effective use requires structured workflows and mechanisms that prevent excessive dependence on uncontrolled model behavior.

Applications at the project and design levels further illustrate the expanding role of generative AI across construction activities. Recent studies have shown that generative models can accelerate conceptual design, produce alternative layouts, and enhance sustainability assessments when integrated with BIM and digital twin environments, particularly by enabling rapid scenario exploration and reducing material waste [20, 22]. In project management settings, the use of large language models such as ChatGPT has demonstrated potential for automating routine planning descriptions, refining resource allocation, and supporting team coordination [23]. Emerging discussions also highlight generative AI as a candidate technology for advancing safety and sustainability outcomes, particularly by improving planning accuracy and facilitating data-driven decisions [24].

Although these developments indicate accelerating

interest, current generative AI applications in construction largely operate as general-purpose assistants. Most systems do not incorporate domain-specific constraints, authoritative regulatory sources, or mechanisms that ensure traceable and verifiable outputs. Prior studies repeatedly point to the need for generative AI systems that can operate within clearly defined domain boundaries and provide transparent reasoning aligned with industry standards. The absence of such capabilities highlights a continuing gap, particularly for safety-critical contexts where the interpretation of formal regulatory requirements must be both accurate and verifiable.

3 System Design

3.1 Overview

The Construction Safety Assistant is implemented as a retrieval-augmented platform that supports OSHA-grounded question answering for construction safety education and training. The system consists of three primary components: OSHA document preprocessing, semantic retrieval, and controlled answer generation, all integrated into a web-based interface. The workflow begins with the preparation of OSHA construction safety documents, which are parsed, segmented, and transformed into an indexed knowledge base to enable efficient semantic search. When a user submits a query through the interface, the system encodes the input and matches it against the indexed corpus. Retrieved OSHA excerpts are then incorporated into a structured prompt for the generative model to ensure grounded and citation-constrained responses. The platform outputs an OSHA-referenced answer along with citation links pointing to the corresponding regulatory documents. Figure 1 presents the system architecture and the end-to-end flow from query submission to retrieval, generation, and citation display.




Figure 1. Overview of the construction safety assistant

3.2 OSHA Document Processing

The Construction Safety Assistant is built on OSHA construction safety regulations obtained in PDF format. Each document is converted into machine-readable text through an automated extraction procedure that preserves the regulatory reading order while removing layout artifacts inherent to PDF formatting. Non-informative elements such as headers, footers, page numbers, and residual formatting symbols are removed to ensure that only substantive regulatory content is retained.

After extraction, the regulatory text is segmented into fixed-length fragments to support efficient semantic retrieval. Each fragment is paired with metadata identifying the source OSHA document and its corresponding page index within the original file. These metadata elements enable transparent citation during answer presentation and allow users to inspect the regulatory basis of each response.

The processed fragments are then transformed into vector representations using a pretrained embedding model and stored in an indexed corpus that supports top-k similarity search. This indexing structure allows user queries to be mapped to relevant OSHA content with high precision, forming the foundation for verifiable and document-grounded question answering.

3.3 Retrieval-Augmented Answer Generation

The answer generation process integrates semantic retrieval with controlled natural-language synthesis to produce grounded explanations drawn from OSHA construction safety regulations. When a user submits a question, the system converts the input into a dense semantic representation that supports comparison with the indexed regulatory corpus. Let q denote the user query. The embedding function $f(\cdot)$ maps the query into a vector space as shown in Equation (1):

$$e_q = f(q) \quad (1)$$

This representation allows the system to evaluate conceptual similarity between the user query and the embedded OSHA fragments.

Each OSHA text fragment is represented by a vector e_i , and the semantic similarity between the query and each fragment is computed using cosine similarity. The similarity score is given in Equation (2):

$$\text{Sim}(e_q, e_i) = \frac{e_q \cdot e_i}{\|e_q\| \|e_i\|} \quad (2)$$

The retrieved fragments are ranked according to this similarity measure. The system selects the top-k fragments that exhibit the strongest alignment with the query, as expressed in Equation (3):

$$R_k(q) = \text{TopK}_{e_i \in E} \text{Sim}(e_q, e_i) \quad (3)$$

where E denotes the full embedded corpus. Each retrieved fragment includes metadata identifying the source OSHA document and page location, which is

preserved throughout the subsequent generation process.

After the relevant fragments are identified, the system constructs a structured prompt that combines the original user query with the retrieved regulatory text. This prompt formulation establishes the context and boundaries within which the generative model operates, limiting the model's responses to material drawn from the retrieved fragments. When the retrieval stage fails to return fragments that meet the predefined relevance threshold, the system activates a fallback mechanism that guides the model to request clarification or to provide a neutral response rather than infer unsupported information. Figure 2 illustrates the structured prompt construction process.




Figure 2. Structured prompt construction process

Using the retrieved content and structured prompt, the generative model synthesizes an answer that restates and interprets the OSHA provisions. Citation metadata is embedded directly into the response and linked to the original OSHA documents, allowing users to inspect the regulatory basis of each statement through the interface. This grounding mechanism supports transparent and verifiable use of OSHA standards in safety education and training.

3.4 Web Interface

Following the retrieval and generation pipeline, the Construction Safety Assistant presents the grounded responses through a web-based interface designed to support transparent review of OSHA regulatory material. The interface provides two coordinated components: a conversational panel for user interaction and a document viewer for inspecting cited OSHA content. In the conversational panel, users submit safety-related questions in natural language. The system's responses appear in the same panel and include embedded citation markers that identify the OSHA document and page from which the supporting content was retrieved. These markers are integrated directly into the response text to

maintain readability while indicating the regulatory source of each referenced statement. Selecting a citation opens the corresponding OSHA PDF page in a side panel, allowing users to review the original regulatory wording without leaving the ongoing interaction.

4 Implementation and Evaluation

4.1 Implementation

The platform is composed of a Python backend and a browser-based frontend. OSHA regulations are stored in PDF format and processed using a text extraction pipeline that removes headers, footers, page numbers, and formatting artifacts. The cleaned text is segmented into fragments of approximately 300 characters, each paired with metadata identifying the source PDF filename and page index.

The processed fragments and metadata are uploaded to Azure Blob Storage, where a Shared Access Signature (SAS) URL is generated for each fragment to provide direct access to its corresponding location within the original PDF. Text embeddings are created using the text-embedding-ada-003 model, and both the embeddings and their metadata are indexed through Azure Cognitive Search to enable vector-based similarity retrieval.

At runtime, the backend receives a user query, computes its embedding, and issues a vector search request to Cognitive Search. The retrieved OSHA fragments and their SAS URLs are assembled into a structured prompt that is sent to a GPT-4o model for grounded answer generation. Citation links are inserted into the output using the stored metadata so that each referenced clause can be opened directly.

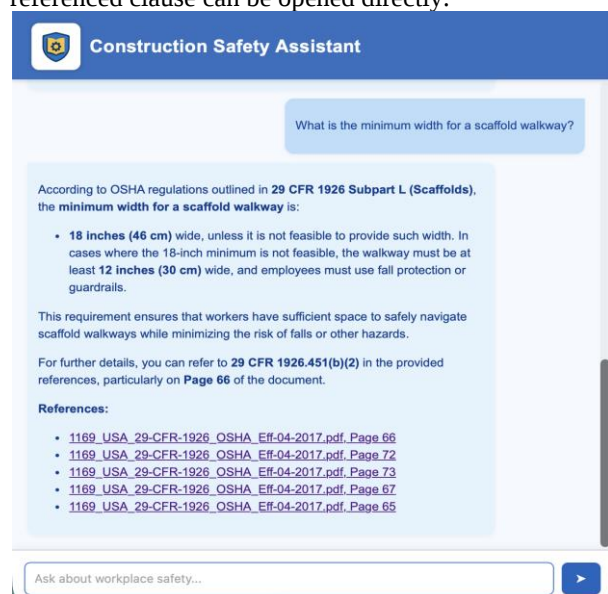


Figure 3. Web interface of the platform

The backend exposes REST endpoints for query submission and citation retrieval. The application runs in a cloud environment that supports stateless request-response communication with the frontend, which is implemented using standard web technologies. As illustrated in Figure 3, the web interface allows users to submit queries, view generated responses, and interact with citation links to access the corresponding OSHA PDF pages in an embedded viewer, completing the end-to-end interaction between retrieval, generation, and document-level inspection. When users select a citation in the generated response, the front end loads the associated SAS URL and displays the referenced OSHA PDF page.

4.2 Automated Evaluation Framework

The performance of the platform was assessed using an automated evaluation framework configured with UpTrain [25], which enables systematic comparison between system-generated answers and authoritative OSHA regulatory text. UpTrain provides normalized metric scores between 0 and 1 for each evaluation criterion, supporting large-scale, clause-level assessment without manual scoring. A benchmark set of thirty representative safety questions was constructed to cover major regulatory domains such as fall protection, excavation and trenching, scaffolding, electrical hazards, confined spaces, and personal protective equipment. Each question was paired with a ground-truth citation from the corresponding OSHA construction standard, forming the reference set used in the evaluation. Table 1 summarizes the benchmark questions and their associated regulatory topics.

Table 1 Questions used for automated evaluation

No.	Question	Type
1	What is the minimum width for a scaffold walkway?	
2	What does OSHA define as a “competent person” in excavation work?	
3	What is the permissible exposure limit (PEL) for noise in construction?	
4	What is the maximum allowable slope for a Type C soil trench?	
5	What class of hard hat offers protection from electrical hazards?	What
6	What are the signage requirements near a trenching operation?	
7	What are the minimum illumination levels required in construction tunnels?	
8	What are the minimum spacing requirements for cross braces on	

	scaffolds?	
9	What is the definition of “confined space” in construction?	
10	When is fall protection required for workers on low-slope roofs?	
11	When should a safety harness be inspected?	
12	When should protective systems be installed in excavations?	When
13	How often must trenches be inspected by a competent person?	
14	How often must cranes be inspected?	
15	How should workers be protected from cave-ins during trenching?	
16	How does OSHA define “confined space” in construction?	
17	How do scaffolds and aerial lifts differ in terms of fall protection?	How
18	How should hard hats be worn in construction zones?	
19	How often must safety harnesses be inspected?	
20	Why is eye protection mandatory during welding operations?	
21	Why must hard hats be worn in construction zones?	
22	Why are GFCIs required on construction sites for power tools?	Why
23	Why is ventilation important in welding operations inside a confined space?	
24	Is it acceptable to stack bricks more than 7 feet high without restraint?	
25	Is training required before operating a scissor lift?	Is /
26	Is daily site inspection required for all excavation work regardless of depth?	Are
27	Are reflective vests mandatory for night work near roadways?	
28	Can a scaffold be used if it is partially disassembled?	
29	Can a ladder be used horizontally as a platform?	Can
30	Can a trench shield be used in unstable soil conditions?	

Three quantitative metrics were used to characterize system performance. Context relevance evaluates whether the retrieved OSHA passages correspond to the substantive focus of the user query. Factual accuracy assesses the correctness of statements in the generated answer relative to the regulatory source. Response relevance measures how well the answer addresses the intent of the question in a clear and semantically appropriate manner. In addition to these metrics, a

Reference Match check was included to verify that the citation links embedded in the output point to the expected OSHA clauses, ensuring traceability and source transparency.

4.3 Evaluation Results

The automated evaluation produced consistent evidence of strong grounding performance across the selected metrics. The system achieved an average context relevance score of 0.95, indicating that retrieved OSHA passages typically aligned with the substantive focus of the questions. Factual accuracy reached 0.99, reflecting near-perfect agreement between generated statements and the corresponding regulatory clauses. The average response relevance score of 0.90 suggests that the generated answers generally addressed the user intent with appropriate specificity. All citation links passed the Reference Match check without errors. Table 2 summarizes the metric scores for all thirty benchmark questions.

Table 2 Automated evaluation results

No.	Context relevance	Factual accuracy	Response relevance
1	1	1	1
2	1	1	1
3	0.67	1	1
4	1	1	1
5	1	1	1
6	0.67	1	1
7	1	1	1
8	1	1	1
9	1	1	1
10	0.67	1	0.67
11	1	1	1
12	0.67	1	1
13	1	1	1
14	1	1	1
15	0.67	1	1
16	1	1	0.67
17	1	1	1
18	1	1	0.67
19	1	1	0.67
20	1	1	0.67
21	1	1	0.67
22	1	1	0.67
23	1	0.67	0.67
24	1	1	1
25	1	1	1
26	1	1	1
27	1	1	0.67
28	1	1	1
29	1	1	1
30	1	1	1
Ave. score	0.95	0.99	0.90

5 Discussion

The evaluation results highlight several characteristics of the system's behavior across different question types and regulatory topics. High factual accuracy reflects the constraint imposed by retrieval-augmented generation, which restricts responses to OSHA source text and minimizes reliance on model priors. Variations in context relevance are largely associated with questions whose answers span multiple regulatory clauses, particularly in areas such as confined spaces and ventilation where relevant information is distributed across several subparagraphs. Lower response relevance in some how and why questions stem from the broader interpretive nature of these items, which may not be fully represented within a single retrieved fragment.

The observed patterns underscore the influence of OSHA's document structure on retrieval performance. Standards that rely on explicit thresholds or enumerated requirements, such as fall protection and scaffolding, align well with fragment-level retrieval and produce consistent grounding results. Sections organized around descriptive or condition-dependent language present additional challenges, as semantically related content may reside in multiple locations. This separation affects the system's ability to assemble a complete context window for procedural or explanatory questions.

Several limitations arise from these characteristics. The retrieval process operates at the level of individual text fragments and does not aggregate related clauses unless they co-occur within the same segment. Citation links reflect the segmentation scheme rather than the full logical scope of a regulation, which may omit adjacent supporting text. Additionally, queries that rely on implicit hazard rationales or multi-step procedures may not fully align with the structure of the underlying standards, reducing the semantic match between the question and the retrieved passage. While the automated evaluation provides a consistent and scalable assessment of system performance, it does not fully capture aspects such as interpretability for practitioners or effectiveness in real-world decision-making contexts.

For future research, improvements to the retrieval process could focus on integrating multi-clause retrieval to better capture regulatory content that is distributed across several sections. Additional refinement of semantic clustering may also help group related provisions into coherent units, improving coverage for questions that require broader contextual grounding. Enhancing query processing with domain-informed expansion represents another opportunity to reduce mismatches between user phrasing and OSHA terminology. Future work will also incorporate expert-based and user-centered evaluation to complement automated assessment and to further examine the

usability and practical effectiveness of the system in construction safety contexts.

6 Conclusion

This study developed a retrieval-augmented question-answering system to support construction safety education and training through clause-linked access to OSHA standards. The system integrates document processing, semantic retrieval, and citation-aware answer generation within a web-based interface, enabling users to examine both the generated response and its regulatory source. The findings provide a direct response to the research question by demonstrating that a retrieval-augmented framework can enable accurate, transparent, and clause-level access to OSHA construction safety regulations.

Automated evaluation across thirty representative safety questions showed consistently high factual accuracy and strong performance in relevance and citation correctness, with lower performance primarily observed in questions whose regulatory basis spans multiple sections of the standards. These results indicate that clause-grounded retrieval combined with citation-aware generation can support more reliable access to OSHA regulations by aligning user queries with verifiable regulatory content and preserving transparency in the response process.

The proposed approach provides a practical mechanism for interacting with prescriptive and authoritative documents such as OSHA standards, particularly in contexts where regulatory information is distributed across multiple sections and requires structured access. Future work may focus on improving retrieval coverage and contextual aggregation for complex regulatory queries. In addition, comparative evaluation with existing tools, such as keyword-based search systems and general-purpose language models, will be conducted to further assess the relative performance and practical advantages of the proposed approach. The framework can also be extended to additional regulatory corpora and specialized application contexts.

References

- [1] U.S. Bureau of Labor Statistics. Fatal falls in the construction industry in 2023. On-line: <https://www.bls.gov/opub/ted/2025/fatal-falls-in-the-construction-industry-in-2023>.
- [2] J. E. Doodoo and H. Al-Samarraie. A systematic review of factors leading to occupational injuries and fatalities. *Journal of Public Health*, 31(1):99–113, 2023.
- [3] U.S. Department of Labor. Occupational Safety and

- Health Administration. *Occupational Safety and Health Administration*. On-line: <https://www.osha.gov/>, Accessed: 30/11/2025.
- [4] G. Chambers. Concise Guide to Workplace Safety and Health: What You Need to Know, When You Need It. *CRC Press*, 2011.
- [5] H. Kim, H. S. Lee, M. Park, B. Chung, and S. Hwang. Information retrieval framework for hazard identification in construction. *Journal of Computing in Civil Engineering*, 29(3):04014052, 2015.
- [6] U.S. Department of Labor. Occupational Safety and Health Administration. *Occupational Safety and Health Administration*. On-line: <https://www.osha.gov/>, Accessed: 30/11/2025.
- [7] CPWR – The Center for Construction Research and Training. *Training Resources*. On-line: <https://www.cprw.com/>, Accessed: 30/11/2025.
- [8] E. Kasneci, K. Seßler, S. Küchemann, M. Bannert, D. Dementieva, F. Fischer, ..., and G. Kasneci. ChatGPT for good? On opportunities and challenges of large language models for education. *Learning and Individual Differences*, 103:102274, 2023.
- [9] S. Zhao, Y. Yang, Z. Wang, Z. He, L. K. Qiu, and L. Qiu. Retrieval augmented generation (RAG) and beyond: A comprehensive survey on how to make your LLMs use external data more wisely. *arXiv preprint arXiv:2409.14924*, 2024.
- [10] M. Aqib. Development of a Question Answering System Over Building Codes using Retrieval Augmented Generation, 2025.
- [11] H. Sadeghi, S. R. Mohandes, M. R. Hosseini, S. Banihashemi, A. Mahdiyar, and A. Abdullah. Developing an ensemble predictive safety risk assessment model: Case of Malaysian construction projects. *International Journal of Environmental Research and Public Health*, 17(22):8395, 2020.
- [12] X. Zhang and S. R. Mohandes. Occupational health and safety in green building construction projects: A holistic Z-numbers-based risk management framework. *Journal of Cleaner Production*, 275:122788, 2020.
- [13] S. Tabatabaee, S. R. Mohandes, R. R. Ahmed, A. Mahdiyar, M. Arashpour, T. Zayed, and S. Ismail. Investigating the barriers to applying Internet-of-Things-based technologies to construction site safety management. *International Journal of Environmental Research and Public Health*, 19(2):868, 2022.
- [14] X. Li, L. Zhang, and Y. Wang. Causal inference of construction safety management measures towards workers' safety behaviours: A multidimensional perspective. *Safety Science*, 172:106432, 2024.
- [15] Y.-H. Huang and T.-R. Yang. Exploring on-site safety knowledge transfer in the construction industry. *Sustainability*, 11(22):6426, 2019.
- [16] J. Li, Y. Zhou, and B. Guo. Enhancing knowledge of construction safety: A semantic network analysis approach. *Buildings*, 15(17):3036, 2025.
- [17] Y. Mei, J. Huang, J. Liu, and L. Jia. A study of factors influencing construction workers' intention to share safety knowledge. *Buildings*, 14(2):440, 2024.
- [18] Q. Huang, S. Zhang, and D. Fang. Effects of construction workers' safety knowledge on hazard-identification performance via eye-movement modeling examples training. *Safety Science*, 180:106653, 2024.
- [19] P. Ghimire, K. Kim, and M. Acharya. Opportunities and challenges of generative AI in construction industry: Focusing on adoption of text-based models. *Buildings*, 14(1):220, 2024.
- [20] R. Nyqvist, A. Peltokorpi, and O. Seppänen. Integration of generative artificial intelligence across construction management. In *Proceedings of the IOP Conference Series: Earth and Environmental Science*, volume 1389(1), pages 012011, 2024. IOP Publishing, August.
- [21] K. K. Naji, M. Gunduz, A. Mohamed, and A. Alomari. Generative AI for sustainable project management in the built environment: Trends, challenges, and future directions. *Sustainability*, 17(20):9063, 2025.
- [22] V. J. Gan. BIM-based building geometric modeling and automatic generative design for sustainable offsite construction. *Journal of Construction Engineering and Management*, 148(10):04022111, 2022.
- [23] K. Cinkusz, J. A. Chudziak, and E. Niewiadomska-Szynkiewicz. Cognitive agents powered by large language models for agile software project management. *Electronics*, 14(1):87, 2024.
- [24] N. Rane. (2023). Role of ChatGPT and similar generative artificial intelligence (AI) in construction industry. Available at SSRN 4598258.
- [25] UpTrain: Open-source LLM evaluation framework. On-line: <https://github.com/uptrain-ai/uptrain>, Accessed: 05/13/2025.