The Society For
Transdiciplinary Studies

JTS
The Journal of Transdisciplinary Studies
p-ISSN 2586-5439
e-ISSN 2713-7457

Original Article

Designing an AI-Driven Safety Education System for Foreign Construction Workers

Myunghee Jung[1], Younghee Noh[2*], Jong-Hwa Jang[3], Bo-Ram Hyun[4], Sang Hyeon Ju[5*]

[1] Professor, Department of Software Engineering, Anyang University, Republic of Korea
[2] Professor, Department of Library and Information Science, Konkuk University, Republic of Korea (irs4u@kku.ac.kr)
[3] Professor, Department of Dental Hygiene, College of Health Science, Dankook University, Republic of Korea (jhj@dankook.ac.kr)
[4] Research Professor, Department of Performing Arts, College of Health and Human Development, Korean National Sport University, Republic of Korea (boram0770@hanmail.net)
[5] Professor, Department of Public Administration, Jeonbuk National University, Republic of Korea (ju1414@jbnu.ac.kr)
*Corresponding Author : irs4u@kku.ac.kr; ju1414@jbnu.ac.kr

ABSTRACT

Objectives: As the Korean construction industry increasingly relies on foreign workers to address labor shortages, limitations of conventional safety education, characterized by one-directional instruction and insufficient linguistic adaptability, have emerged as a critical risk factor. These approaches fail to accommodate language barriers and heterogeneous learning needs, thereby increasing accident risks and hindering workplace integration. This study proposes an intelligent safety education framework and an AI-driven training architecture for foreign construction workers in Korea, integrating perspectives from construction safety, educational technology, and artificial intelligence. **Methods:** The proposed framework conceptualizes safety education as a progressive process comprising four stages: digitized, adaptive, immersive, and predictive. Based on this framework, the architecture leverages Large Language Models (LLMs) and multimodal data analysis to automatically generate multilingual, context-aware training content and to deliver personalized learning pathways according to individual proficiency levels and risk profiles. **Results:** An implementation case demonstrates the technical feasibility, scalability, and regulatory compatibility of the proposed system with Korean industrial safety standards. **Conclusions:** The findings suggest that the proposed approach can enhance safety education in multicultural construction environments and contribute to the development of a sustainable, technology-driven safety culture that supports both safety competency development and social inclusion.

Keywords: Context-aware content, Foreign construction worker, Generative AI-based training, Multilingual accessibility, Safety education

Received 29 Nov 2025
Revised 15 Dec 2025
Accepted 23 Dec 2025

Citation: Jung M, Noh Y, Jang JH, Hyun BR, Ju SH. Designing an AI-Driven Safety Education System for Foreign Construction Workers. The Journal of Transdisciplinary Studies. 2025; 9(3): 105-119. https://doi.org/10.22685/jts.2025.9.3.105

1. Introduction

In the construction industry, effective safety training is a fundamental requirement to prevent severe injuries and fatalities caused by human error and dynamic site conditions. However, traditional safety-education approaches employed in this sector, including classroom lectures, printed manuals, and static video demonstrations, provide limited interactivity and lack the personalized learning pathways necessary to engage workers effectively. Moreover, these methods offer little capacity for assessing learners' comprehension in real time, further reducing their instructional efficacy [1].

In Korea, construction safety programs have historically relied on one-directional, classroom-based instruction delivered primarily in Korean. Such programs do not sufficiently address linguistic diversity or adapt to the heterogeneous learning needs of the workforce. At the same time, a declining domestic working-age population has increased dependence on foreign labor. Foreign construction workers now constitute an essential human resource for maintaining productivity in labor-intensive and accident-prone environments [1,2].

Despite their indispensable contributions, many foreign workers continue to face systemic barriers that undermine both safety and workplace integration. These challenges include language and cultural gaps, limited access to formal and comprehensible safety training, and inconsistent communication of site-specific hazards [2]. Such barriers pose direct risks to individual workers and, more broadly, threaten the overall stability, safety performance, and competitiveness of the construction industry.

Considering these challenges, there is a growing need for an evolutionary redesign of safety education systems—one capable of dynamically responding to heterogeneous learner characteristics and rapidly changing worksite environments. While recent breakthroughs in artificial intelligence and immersive media technologies offer opportunities to reimagine safety training, a unified framework specifically tailored for the linguistic and contextual needs of foreign workers remains underexplored.

This study focuses on establishing a technical and pedagogical roadmap for transforming construction safety education into an adaptive, data-driven ecosystem. The paper proposes an evolutionary framework for safety education and a Generative AI–driven system architecture designed specifically for foreign construction workers in Korea. The framework conceptualizes safety education as a continuous evolution across four stages—digitized, adaptive, immersive, and predictive. Within this framework, the proposed architecture leverages Generative AI to provide multilingual support and context-aware safety guidance, bridging the gap between static instructions and real-world practice.

Furthermore, the study presents an implementation case demonstrating that the proposed system aligns with Korean industrial safety standards and possesses practical feasibility. By integrating AI-based personalized learning and multilingual generative technologies, this research aims to strengthen the safety competencies and social integration of foreign construction workers, thereby advancing the sustainable development of the Korean construction industry.

2. Background and related work

2.1. Limitations of current construction safety education

Current safety training in construction continues to rely on passive, regulation-driven approaches such as classroom lectures, printed manuals, and one-time orientation sessions. These methods are largely checklist-oriented and fail to adequately incorporate learners' behavioral understanding, resulting in low engagement and limited transfer of knowledge to real on-site practice. The predominance of monolingual and non-interactive materials further restricts access for foreign workers who often lack sufficient Korean language proficiency or prior technical training [3].

Although digital transformation efforts such as computer-based instruction, mobile learning, virtual reality (VR), and augmented reality (AR) have demonstrated improvements in knowledge retention and risk perception by simulating hazardous scenarios [5, 6], these systems remain largely static and instructor-authored. They are generally incapable of adaptively modifying training content based on learner performance or real-time worksite conditions. Moreover, high development costs and the requirement for specialized equipment limit scalability, particularly for small and medium-sized construction firms. Collectively, these limitations underscore the pressing need for construction safety education to evolve into a more adaptive, data-driven, and responsive learning ecosystem.

2.2. Trends in AI and generative AI for education and safety applications

Artificial intelligence has made continuous advancements in both educational and occupational safety domains, evolving from early rule-based systems and predictive analytics into highly adaptive, multimodal generative models. In traditional educational technology, AI has primarily supported learner analytics, automated assessment, and intelligent tutoring through structured machine-learning pipelines [7].

In construction and industrial safety, early applications of AI centered on computer-vision techniques for detecting personal protective equipment (PPE), recognizing unsafe behaviors, and predicting accident risks. Building on these capabilities, advanced machine learning techniques now process and analyze multimodal datasets, including sensor streams, incident and near-miss reports, and longitudinal learner performance logs, to identify latent risk patterns and recommend targeted, evidence-based interventions. Although these systems improved monitoring and regulatory compliance, they remained limited in their ability to personalize instruction, support multilingual workers, or adapt dynamically to situational contexts.

Recent developments in generative AI—driven by Large Language Models (LLMs) and Vision–Language Models (VLMs) have significantly expanded these capabilities. LLMs and generative AI systems autonomously produce multilingual instructional materials, role-based scenario simulations, and context-aware guidance. This automation significantly reduces the cost and labor associated with developing localized and culturally appropriate safety training resources [8]. In the field of education, this technology now enables personalized and dynamic learning pathways by producing micro-lessons, quizzes, and formative feedback aligned with each learner's proficiency and cognitive profile [9]. This shift marks a transition from static content delivery to AI-mediated instruction capable of responding to individual learning needs in real time.

In the safety domain, generative models and simulation tools can create rich training scenarios and "what-if" analyses that were previously impractical to produce manually. AI-driven virtual instructors can engage learners in interactive safety drills, while VR environments can be populated with AI-generated hazard scenarios. When integrated with Extended Reality (XR) platforms and edge-AI technologies, these AI-driven tools enable the creation of real-time, site-specific learning environments that seamlessly bridge the gap between simulated training and real-world practice. Together, these innovations constitute a new paradigm for intelligent safety education that supports continuous, adaptive, and highly personalized learning for diverse worker populations.

2.3. *Gaps and research motivation*

Although prior studies have demonstrated the effectiveness of VR and AI in enhancing safety training [3], a critical gap remains in establishing a unified framework that integrates dynamic content generation, adaptive feedback mechanisms, and real-time monitoring. Existing approaches rarely incorporate multilingual personalization or employ generative AI for scenario synthesis [10]. In addition, modern safety-training systems increasingly require real-time data analytics sourced from IoT sensors and wearable devices to ensure that instructional content reflects actual on-site conditions [11].

These limitations underscore the necessity for a holistic, data-driven, and generative framework capable of evolving in parallel with technological advancements and the growing diversity of the workforce. Future research must therefore (1) formulate an evolutionary roadmap for modernizing safety education, (2) develop a modular AI architecture that supports generative, adaptive, and multimodal instruction, and (3) validate its practical utility through implementation scenarios tailored to Korea's multicultural construction workforce.

3. System design for AI-driven safety education in construction

This study proposes a generative AI–driven safety training architecture designed to deliver adaptive, multilingual, and reliable instructional content to construction workers—including foreign laborers who often operate under significant linguistic and technological constraints. The technical features presented in the preceding sections establish the fundamental functional requirements that the architecture must satisfy and articulate how these capabilities are organized and operationalized within a multilayered, AI-enabled infrastructure.

Furthermore, the study elaborates on the generative mechanisms that support dynamic scenario synthesis, multilingual content transformation, and personalized instructional adaptation. These mechanisms enable the system to generate context-aware hazard scenarios in real time, simplify complex safety instructions, translate training materials across multiple languages, and tailor educational content to each worker's proficiency and risk profile.

Taken together, these components constitute a coherent conceptual and technical foundation for advancing next-generation safety training systems powered by generative AI.

3.1. *Intelligent safety education system*

The intelligent safety education system is built upon four fundamental technical capabilities: multilingual instruction, personalized learning pathways, modular deployment, and integrated safety guardrails as shown in Fig. 1. These features represent the essential functional expectations for an AI-based educational platform. The diverse and multilingual nature of contemporary construction workforces necessitates real-time language adaptation [12], while the high variability in worker skill levels requires tailored content delivery. Furthermore, deployment flexibility is vital given the heterogeneity of on-site infrastructure, and safety guardrails are indispensable for ensuring regulatory compliance and preventing harmful AI outputs. Collectively, these capabilities define what the system must be able to accomplish in order to function as an inclusive, adaptive, and reliable safety education solution.

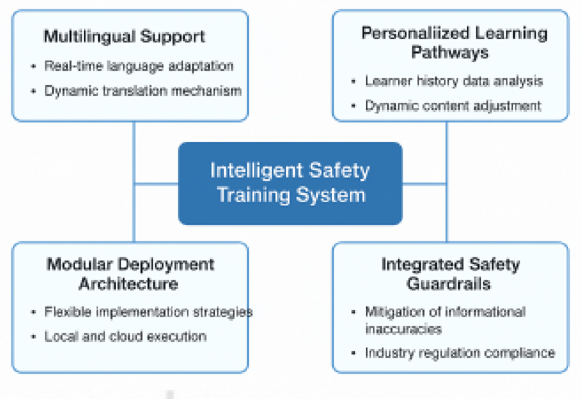

Fig. 1. Technical features of the intelligent safety education system

3.1.1. Multilingual instruction capability

To address the extensive linguistic diversity characteristic of modern construction sites, the system integrates multilingual training capabilities powered by advanced translation application programming interfaces (APIs). This functionality enables real-time language adaptation and ensures seamless delivery of instructional content across multiple languages. Such capabilities are particularly critical for accommodating foreign workers and individuals with varying levels of literacy, thereby ensuring equitable access to essential safety information regardless of linguistic background. Moreover, the system incorporates dynamic translation mechanisms designed to preserve semantic accuracy while adapting instructional materials to the language proficiency of individual users [13, 14].

3.1.2. Personalized learning pathways

The personalization module leverages comprehensive learner history data and detailed interaction logs to generate individualized educational experiences. Through continuous analysis of workers' performance trends and behavioral patterns, the system dynamically adjusts the training content, calibrates difficulty levels, and customizes feedback mechanisms to align with each user's needs. This adaptive approach optimizes knowledge retention and facilitates the practical application of safety protocols by accounting for each worker's unique risk profile, learning pace, and comprehension capacity [15]. The personalization algorithm iteratively refines its recommendations based on accumulated user data, thereby supporting progressive improvements in learning outcomes.

3.1.3. Modular deployment architecture

Given the heterogeneous and often unstable infrastructure environments typical of construction sites, the system employs a modular deployment architecture that supports flexible and context-sensitive implementation strategies. This design accommodates local execution modes—including on-device and edge computing—while also supporting cloud-based deployment options. Such architectural flexibility ensures operational continuity and system resilience even under conditions of limited or intermittent network connectivity. This dual-mode capability allows construction organizations to select and optimize deployment strategies according to their specific technological infrastructure and operational constraints.

3.1.4. Integrated safety guardrails

compliance, comprehensive safety guardrails are embedded throughout the platform. These safeguards are engineered to reduce informational inaccuracies, prevent hallucinations in AI-generated outputs, and enforce adherence to industry-specific safety regulations. The verification framework upholds the integrity and credibility of training materials by implementing multi-layered content validation protocols and controlled output-generation mechanisms. These safeguards are essential for establishing user trust and meeting stringent occupational health and safety standards [16].

In addition to content reliability, data privacy constitutes a critical component of the system architecture. To protect sensitive worker data such as proficiency levels and nationality, the system enforces strict data privacy protocols in compliance with Korea's Personal Information Protection Act (PIPA). All personally identifiable information (PII) is pseudo-anonymized and encrypted both at rest (AES-256) and in transit (TLS 1.3). Access controls are implemented based on the principle of least privilege, ensuring that only authorized safety managers can view individual worker profiles.

3.2. Generative AI-driven architecture

To operationalize the core technical capabilities outlined in Section 3.1—multilingual support, personalized learning pathways, modular deployment, and integrated safety guardrails—this study proposes a four-layer generative AI-based architecture that seamlessly connects on-site learning contexts with cloud-based generative intelligence and supervisory governance.

At the base of the framework, the On-Site Interaction Layer functions as the primary interface for foreign construction workers. Through mobile devices, kiosks, and lightweight headsets, this layer delivers multilingual microlearning content, VR-based hazard simulations, and voice-guided instructional modules. It is responsible for directly reflecting the worker's immediate learning needs and the situational demands of the construction environment.

Above this, the AI Safety Education Engine serves as the core computational hub of the system. Leveraging large language models (LLMs) and recommendation algorithms, the engine dynamically generates simplified explanations, multilingual translations, and task-specific instructional modules tailored to individual workers. A cloud-edge hybrid design further ensures operational continuity by maintaining essential training delivery even in environments with limited or unstable network connectivity [17].

The third layer, the Management and Monitoring Layer, provides supervisors and safety managers with real-time dashboards, risk analytics, and content-approval tools. This layer reinforces transparency and accountability across all stages of content generation, distribution, and verification, specifically through a 'Pre-Deployment Approval' workflow. In this Human-in-the-loop (HITL) process, safety managers review AI-generated scenarios and translated content via the dashboard before they are pushed to worker devices. The system automatically flags low-confidence outputs (e.g., AI confidence score < 60%) for mandatory human verification, ensuring that no hallucinatory or non-compliant instructions reach the workers. This rigorous validation process ensures that training materials remain compliant with regulatory and organizational safety standards.

Finally, the Continuous Feedback Loop facilitates iterative improvement by analyzing learner performance, identifying recurring hazards or misconceptions, and regenerating updated training content. As regenerated content is reviewed and validated by field experts or instructional staff prior to redeployment, the system exhibits self-adaptive characteristics that allow it to improve continuously over time [18].

By operating these four layers as an integrated system, the proposed architecture establishes a personalized, culturally and linguistically responsive, and continuously evolving ecosystem for intelligent safety education. Figure 2 summarizes the overall flow of the proposed AI-based safety training system. Worker-facing devices collect multimodal inputs—including images, speech, and behavioral signals—which are processed through on-site edge intelligence for low-latency inference. The synchronized data is then transmitted to the cloud, where large-scale generative processing, adaptive learning, policy retrieval, and advanced analytics are performed. The application and management layer subsequently distributes personalized safety training content, enforces governance mechanisms, and provides supervisors with real-time compliance and training insights.

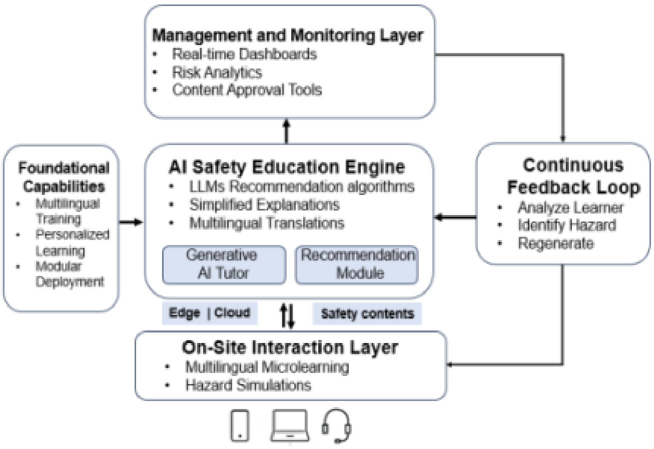

Fig. 2. AI-driven safety education system overview architecture

Table 1. Functional features and roles of each system layer

Layer	Who / What	Main Functions	Generated / Collected Outputs
1. On-Site Interaction Layer (Workers & Devices)	• Foreign construction workers • Smartphones / tablets • Simple VR headsets • Kiosks at site entrance (briefings) • Optional PPE-check cameras	• Display short, visual safety lessons • Run VR simulations for high-risk tasks • Provide voice-based multilingual guidance • Conduct short quizzes and comprehension checks	• Quiz scores • Frequently missed concepts • VR simulation performance • Basic engagement metrics (time spent, completion records, etc)
2. AI Safety Education Engine (Cloud + Edge Hybrid)	**• Core AI Modules** - Generative AI Tutor (LLM) - Recommendation Engine, - Safety Content Library **• Edge / Cloud Setup** - Edge devices at site - Cloud backend for heavy computation	**• Generative AI Tutor** - Simplifies safety text - Multilingual translation - Auto-generates quizzes & explanations **• Recommendation Module** - Assigns next-best lesson based on worker profile **• Edge Device Functions** - Offline fallback - Caching critical content **• Cloud Functions** - Large LLM/VLM execution - Data storage & analytics - Content updates to edge	• Personalized learning paths • AI-generated multilingual content • Risk-tailored scenarios • Updated content packs for edge devices
3. Management & Monitoring Layer (Supervisors & Safety Managers)	• Supervisors • Safety managers • Training officers	**• Training Dashboard** - Training completion tracking - Identification of high-risk groups **• Safety Insight View** - Detection of common unsafe behaviors - dentification of topics requiring more training **• Content Approval Panel** - Review/approve AI-generated lessons - Check compliance with Korean regulations	• Trainee progress reports • Risk and vulnerability heatmaps • Approved and validated content • Compliance documentation
4. Continuous Feedback Loop (Learn & Improve Cycle)	• Workers (learning) • AI engine (analysis) • Instructors(approving updates)	• Capture learner performance & site risk patterns • Analyze error frequency, misconceptions, hazard trends • Regenerate improved content • Instructor verification & deployment	• Iteratively improved content library • Updated scenarios • Revised multilingual modules • Adaptive curriculum aligned with new risks and regulations

The architecture links on-site interactions with cloud-based generative processing, enabling multilingual, personalized, modular, and safety-assured training. Edge devices collect multimodal inputs, the cloud provides adaptive and generative intelligence, and the management layer delivers tailored content and governance. Together, they form a unified system that supports reliable and continuously improving safety education.

3.3. *Generative scenario creation and multilingual/ personalized safety content*

Building on the architectural framework introduced in Section 3.2, this section details the generative mechanisms that enable the system to produce customized, multilingual, and context-aware safety training content. The generative layer integrates large language models (LLMs), vision–language models (VLMs), and structured safety knowledge bases to transform sensor-derived data and contextual site models into actionable instructional materials that align with worker needs and site-specific hazard conditions [19].

The generative process encompasses three primary functions: (1) scenario generation constructs realistic hazard simulations based on real-time site conditions; (2) multilingual content generation ensures accessibility across diverse linguistic backgrounds; and (3) personalized instructional adaptation tailors delivery methods and difficulty levels to individual worker profiles and learning progress.

These components do not function as isolated modules; rather, they operate as tightly integrated elements within the Sense → Understand → Generate → Validate → Deliver → Learn pipeline (Fig. 3). This integrated workflow enables the system to maintain the situational appropriateness, linguistic accessibility, and dynamic responsiveness required to address the evolving safety risks present in Korea's multicultural construction environments.

3.3.1. *Generative scenario creation*

The system employs construction-specific hazard ontologies, worker-context embeddings, and environmental telemetry to automatically generate realistic and pedagogically structured training scenarios. These scenarios may include VR-based fall-prevention simulations, short video demonstrations of common hazards, or step-by-step textual walkthroughs of lock-out/tag-out (LOTO) procedures. The generative model conditions its outputs on key contextual variables—such as task type, worker role, site layout, environmental risk factors, and recent incident history—to ensure situational relevance and instructional appropriateness [20]. This approach aligns with established findings showing that generative scenario variation enhances learning transfer and improves hazard recognition in industrial training environments.

3.3.2. *Multilingual and comprehension-adaptive content generation*

With a workforce comprising diverse linguistic groups such as Vietnamese, Thai, Chinese, and Uzbek speakers, the Korean construction industry requires training that functions independently of language proficiency. The proposed Multilingual and Comprehension-Adaptive Content

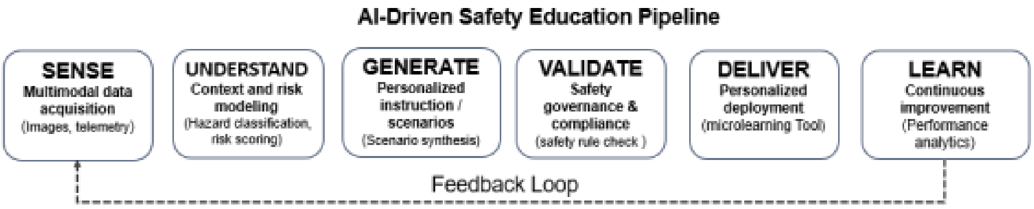

Fig. 3. Six-stage pipeline from data acquisition to continuous learning with feedback integration

Table 2. Generative AI: functionality and products per architectural layer

Architecture Layer	Generative AI Role	Example Output
Sense	Input for contextual adaptation	Images, hazards, voice queries
Understand	Determine learner needs and risk context	Language level, task type
Generate	Main generative actions	scenarios, translated modules, microlearning
Validate	Ensure safety, factuality, compliance	KOSHA-aligned content
Deliver	Adaptive distribution	VR headsets, mobile apps, kiosks
Learn	Model updating & data-driven iteration	Risk prediction updates

Generation module leverages Cognitive Load Theory and Inclusive Learning principles to mitigate the extraneous cognitive load of second-language processing, ultimately facilitating the internalization of critical safety behaviors.

During the Generate stage, large language models (LLMs) and vision–language models (VLMs) convert technical safety manuals, work procedures, and incident reports into accessible microlearning materials [3]. The system automatically rewrites simplified textual explanations, visual hazard cues, and VR/AR narration in each worker's native language. To ensure instructional accuracy, the generative pipeline incorporates controlled vocabularies and safety-specific terminology maps, preventing misinterpretation of regulatory or procedural terminology [21].

In the Deliver stage, the system employs Voice-First Conversational Safety Tutors that provide hands-free explanations, guided walkthroughs, and interactive Q&A support. Accent-tolerant automatic speech recognition (ASR) models accommodate pronunciation patterns that diverge from standard Korean or English, while edge-based speech processing ensures operability in low-connectivity construction environments.

3.3.3. Personalized instructional adaptation

Personalization is embedded into the Generate, Deliver, and Learn stages of the architecture. Personalized Instructional Adaptation refines the system's generative outputs to ensure that each worker receives an instructional sequence tailored to their proficiency level, performance history, and risk exposure. To enable individualized learning pathways, generative outputs are adapted through a structured sequence of transformations defined by the function G, which is composed of three hierarchical sub-functions:

$$G = f_3 \circ f_2 \circ f_1$$

where

f_1 is the Proficiency-Based Filtering function,
f_2 is the Performance-Based Adjustment function, and
f_3 is the Risk-Based Prioritization function.

These operators work together to sequentially: 1) filter instructional content based on the learner's proficiency level and training history; 2) update and recalibrate the learner's effective proficiency based on recent performance signals (e.g., quiz scores, VR simulation behavior); and 3) reorder and prioritize the remaining content according to the worker's personalized risk profile.

G operates as a layered personalization mechanism that adaptively tailors the instructional sequence to each learner's needs. Thus, G is best understood not as a monolithic function but as a layered transformation pipeline, consistent with the way generative AI safety systems operate in practice. This formulation is grounded in Adaptive Learning Theory and aligns with evidence showing that contextualized, learner-specific instruction improves retention and situational hazard awareness in high-risk industries [22].

Following the generation of personalized content, the system selects the next instructional step using a mastery-based decision criterion:

$$\text{Next Modulation}(i) = \begin{cases} ReinforcementModule(i) & \text{if } S_i < \theta \\ AdvancedModule(i+1), & \text{if } S_i \geq \theta \end{cases}$$

where S_i is the learner's skill score and θ is a predefined mastery threshold. This design draws from Mastery learning theory, ensuring that hazardous or high-complexity tasks such as fall-prevention simulations are introduced only after sufficient baseline understanding is demonstrated [21].

During the Learn stage, learner interactions (e.g., micro-assessment performance, VR task behaviors, and correction logs) continually refine the worker model using the profile-update operator f:

$$S_{t+1} = S_t + f(\triangle E_t, \triangle R_t)$$

where ΔE_t represents changes in engagement and ΔR_t represents improvements in risk-related behavior. The operator f is defined as a lightweight additive model:

$$f(\triangle E, \triangle R) = \alpha \triangle E + \beta \triangle R$$

with α and β denoting the sensitivity weights assigned to engagement and risk-behavior improvement, respectively. Rather than implementing full reinforcement learning or complex Bayesian updating, this formulation provides a computationally efficient and interpretable mechanism for incrementally refining learner states in environments with limited connectivity and constrained hardware.

Taken together, the personalization function G and the update operator f are integrated into the overall architectural pipeline: Sense/Understand supply learner and contextual inputs; Generate applies G to produce individualized content; Deliver controls the adaptive release of training modules; and Learn invokes f to adjust the learner model for subsequent iterations.

4. Research results: implementation and validation of AI-driven safety education

4.1. *Overview of the implemented system*

This section presents the research results derived from a fully functional Minimum Viable Product (MVP). Built upon the architectural design proposed in Section 3, this system successfully operationalizes the four-layer generative AI-driven architecture. Consequently, it demonstrates the feasibility of delivering multilingual, personalized safety education, with a specific focus on foreign laborers facing linguistic constraints.

The implemented system comprises 27 core project files, including 7 JavaScript modules, 11 EJS (Embedded JavaScript) template files, and 2 CSS stylesheets, collectively forming a comprehensive web-based platform. The system architecture adheres strictly to the theoretical framework outlined in the research design, providing empirical validation of the conceptual model through practical implementation. The system successfully instantiates all four layers of the proposed architecture shown in Table 3.

4.2. *Functional Capabilities and Feature Implementation*

4.2.1. *Multilingual instruction capability*

The system demonstrates robust multilingual functionality through its custom translation utility module. The implementation achieves comprehensive language coverage across the platform. First, the interface incorporates 4 multilingual UI components and supports three safety lessons in both Korean and their translated counterparts. Real-time language switching is enabled based on worker profiles, ensuring seamless user adaptation during operation. Additionally, all safety category labels are localized into six languages, and proficiency indicators (beginner, intermediate, advanced) are consistently translated, enabling global usability across diverse worker populations.

The translation mechanism is structured into three coordinated layers—UI element translation, safety category translation, and difficulty-level translation—ensuring semantic consistency across interface components, instructional categories, and user-facing learning indicators. This modular architecture allows the system to maintain a unified localization strategy while supporting future extension to additional languages.

Language detection and application are performed through profile-based selection, automatically rendering all interface elements and educational materials in the worker's preferred language. This approach demonstrates that real-time multilingual adaptation is both technically feasible and computationally lightweight, thereby validating the design principles established in Section 3.1. The successful implementation confirms that generative AI-driven safety education systems can effectively accommodate linguistically diverse workforces without imposing significant operational overhead.

Table 3. Roles and implementation methods per layer

Layer	Core Functions & Roles	Application Technology & Implementation Methods
Layer 1: On-Site Interaction Layer	• Access and interaction with safety education content • Device-agnostic learning environment • Delivery of three main safety modules (ex. falls, electrical, equipment)	• Responsive web application for mobile, tablet, and PC • Multilingual UI (4 languages) • Composed of structured web pages (authentication → learning → tracking) • safety education content operation
Layer 2: AI Safety Education Engine	• Generates user-customized training paths • Manages learning data and interaction logs • Generates personalized adaptive learning scenarios	• Backend: Node.js (v14+), Express.js (v4.18.2) • MVP level DB (future transition to PostgreSQL or MongoDB) • Personalized recommendation algorithm based on proficiency • Adaptive learning path creation using worker history + behavior logs
Layer 3: Management & Monitoring Layer	• Data visualization for safety managers • Worker risk analysis and classification (high/medium/low) • Supports real-time monitoring of worker performance	• Visualization: Chart.js (v4.4.0) • Manager dashboard • Real-time statistical processing and risk analysis algorithm • Worker profile analysis+risk-level classification
Layer 4: Continuous Feedback Loop	• Collects learning performance and progress indicators • Dynamically updates training recommendations • Refines user proficiency models over time	• Data collection for quizzes & completion metrics • Proficiency-based learning progress logic • Updating of user proficiency based on cumulative performance • Dynamic recommendation adjustment aligned with learning progression

4.2.2. Personalized learning pathways

The personalized learning system implements the theoretical model described in Section 3.3. Each worker profile stores key attributes—including language preference, job role, proficiency level, completed lessons, quiz history, and recent activity—which together form the basis for individualized recommendations. The recommendation module follows a three-tier strategy. First, proficiency-based filtering selects lessons that match the user's current skill level while excluding completed modules, generating the top three recommendations. Second, dynamic proficiency adjustment recalculates learner proficiency after every quiz using average scores, promoting users to Intermediate at ≥70% and Advanced at ≥85%. Finally, risk-based prioritization identifies high-risk learners (average score <60%) and adjusts the curriculum by emphasizing foundational content and modifying progression speed. This implementation enables continuous, data-driven personalization of safety training pathways.

4.2.3. Modular deployment architecture

The system's modular architecture ensures high deployment flexibility, guaranteeing reliable operation across heterogeneous construction environments. To meet infrastructure adaptability requirements, the platform supports four deployment configurations: (1) local on-site server, (2) cloud deployment (AWS, Azure, Google Cloud), (3) hybrid edge–cloud operation with offline synchronization, and (4) Docker-based containerized deployment for scalable orchestration. System parameters and security credentials are managed via .env files, facilitating seamless transitions between modes and ensuring consistent performance under the varying technological constraints of construction sites.

4.2.4. Integrated safety guardrails

The system incorporates a multi-layered safety guardrail framework that aligns with the theoretical requirements outlined in Section 3. Content validation is enforced through pre-deployment reviews against KOSHA standards, expert verification of quiz answers, and native-speaker checks for critical safety terminology. At the system level, safeguards include input sanitization, session validation, hash-based content integrity verification, and error-boundary mechanisms for graceful failure handling. In addition, an automated risk assessment algorithm classifies workers into high, medium, or low risk based on their recent performance, using the last five quiz attempts and updating the risk level after each completion. Together, these components ensure regulatory compliance, system reliability, and the safe delivery of training content.

4.3. User interface and experience implementation

This section details the UI/UX implementation by presenting screen examples from two core user journeys. The 'Worker Journey' (Scenario 1) demonstrates the learner's path from the initial 'Worker Login' to their personalized dashboard, where they access recommended lessons and review their statistical profile. Concurrently, the 'Manager Monitoring' (Scenario 2) illustrates the administrative workflow, following the 'Manager Login' to the main dashboard for an overview of worker statistics and then to the 'Analytics' page for in-depth data visualization.

Figure 4 presents representative result screens generated by the prototype implementation of the proposed AI-driven safety education system. The figure illustrates three functional components: (1) LLM/VLM-based scenario generation, (2) multilingual instructional content delivery, and (3) user-specific interfaces for workers and supervisors, including the management dashboard and content-verification workflow.

Figure 4(a)–(e) present the worker and manager interfaces of the implemented system. Workers access streamlined microlearning modules, voice-guided instructions, and scenario-based training exercises, enabling quick and context-aware learning. Managers, in contrast, interact with an integrated dashboard that visualizes risk patterns, training progress, hazard distributions, and worker-level risk indices. The dashboard also provides a Content Approval function corresponding to the Validate Layer in the system architecture. It allows supervisors to review, modify, and approve AI-generated materials prior to deployment, ensuring regulatory compliance and preventing the release of unsafe or inaccurate outputs.

Figures 4(f) and 4(g) illustrate example of the multilingual support module. After a scenario is generated, users can select their preferred language (e.g., Vietnamese, Chinese, or English), upon which the LLM produces simplified and culturally accessible instructional text. This feature enables workers with limited Korean proficiency to receive clear, comprehensible safety guidance. The generative model then synthesizes the contextual information into a pedagogically structured scenario such as hazard descriptions, required protective measures, and task-specific corrective actions. These examples demonstrate the system's ability to transform real-world visual inputs into actionable and multilingual safety training materials.

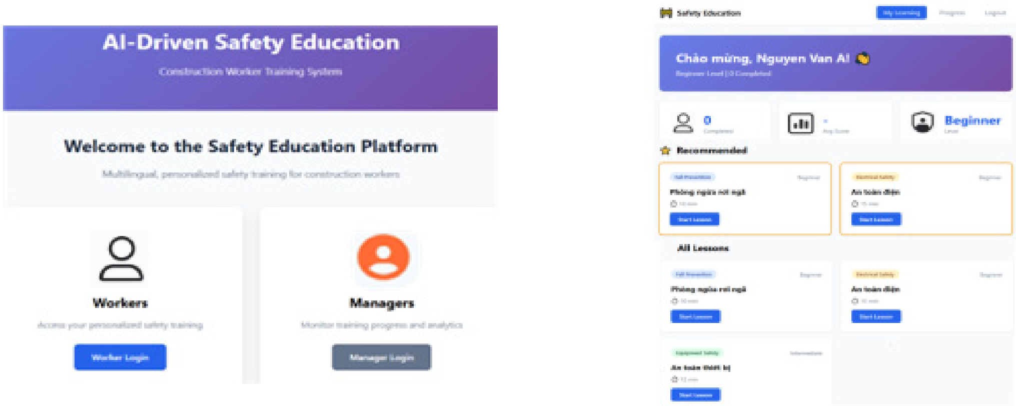

(a) Worker/Manager login interface (b) Worker learning profile

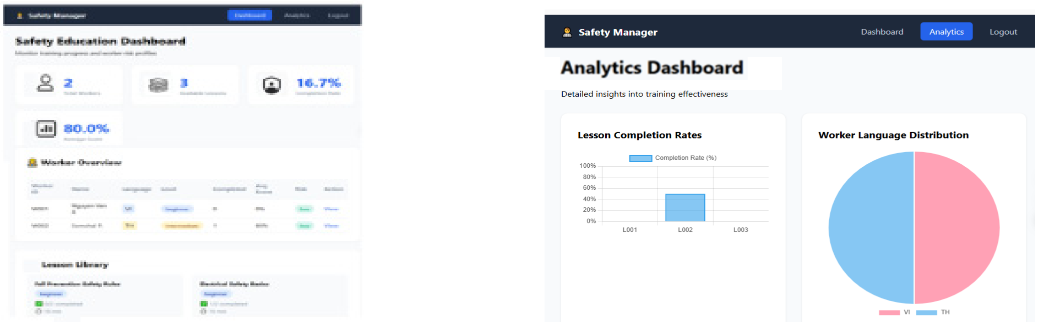

(c) Manager Dashboard

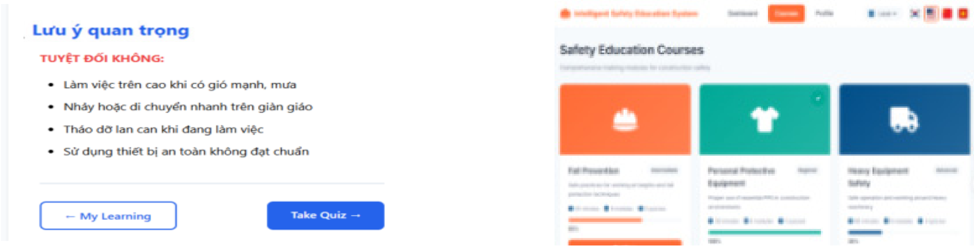
Lưu ý quan trọng

TUYỆT ĐỐI KHÔNG:

- Làm việc trên cao khi có gió mạnh, mưa
- Nhảy hoặc di chuyển nhanh trên giàn giáo
- Tháo dỡ lan can khi đang làm việc
- Sử dụng thiết bị an toàn không đạt chuẩn

(d) Learning recommendation (e) Learing courses selection

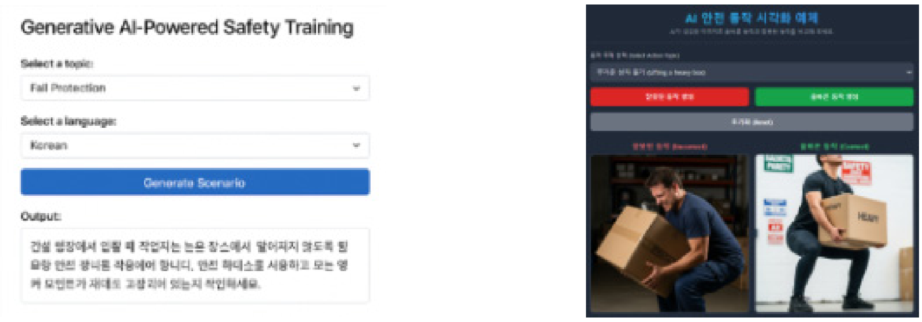

(f) language-aware quiz generator (g) action scenario creation

Fig. 4. Screen examples of practical user interface (UI) and user experience (UX) implementation

4.4. System validation and testing results

This section summarizes the results of functional testing, usability evaluation, and performance benchmarking conducted on the prototype safety education platform. Overall, the system demonstrated high operational reliability, strong user experience quality, and sufficient performance for real-time use in construction-site environments.

4.4.1. Functional testing outcomes

Comprehensive functional tests were performed across authentication, content delivery, quiz operation, recommendation logic, and analytics. All critical functionalities operated as intended as shown in Table 4.. Authentication correctly handled valid/invalid credentials, multilingual lesson content rendered reliably across devices, quizzes produced accurate scores, and the recommendation system adjusted proficiency and lesson suggestions dynamically. Manager-facing analytics—including statistics calculation, chart rendering, and risk-level classification—updated correctly in real-time.

Table 4. Key functional testing results

Test area	Core outcome	Result
Authentication	Correct login handling, session management, logout flow	**Pass**
Content Delivery	Lessons, images, formatting, and multilingual interfaces render correctly	**Pass**
Quiz System	Question display, scoring accuracy, progress updates	**Pass**
Recommendation Engine	Proficiency filtering, exclusion of completed lessons, dynamic level updates	**Pass**
Analytics Dashboard	Metrics calculation, chart rendering, real-time updates	**Pass**

4.4.2. Experimental evaluation of the personalization function

To empirically validate the behavior of the proposed personalization mechanism, we implemented a prototype version of the function

$$G = f_3 \circ f_2 \circ f_1$$

where f_1 \performs proficiency-based filtering, f_2 adjusts module ordering based on recent performance, and f_3 prioritizes hazard-relevant content.

A test dataset consisting of eight safety education modules was created to evaluate the system under controlled conditions (Fig. 5). Each module contains three attributes—difficulty (1–3), risk level (1–3), and hazard type—representing simplified but realistic variations in construction-site safety content. The dataset allowed direct observation of how each stage of G influenced the resulting recommendations. f_1.

The Python implementation applies the three transformation functions sequentially:

f_1 (Proficiency Filter) excludes modules that exceed the worker's capability, retaining only content within one difficulty level above the worker's current proficiency. f_2 (Performance Adjustment) reorders modules depending on recent performance scores: low-scoring workers receive easier modules first and high-scoring workers receive more advanced and higher-risk modules. f_3 (Risk Prioritization) elevates modules related to hazards the worker is most frequently exposed to (e.g., fall-related content for high-risk workers). The overall pipeline therefore produces a ranking that is simultaneously ability-appropriate, performance-sensitive, and hazard-specific. The annotated implementation example code is visualized in Fig. 6.

Running the personalization pipeline produced differentiated learning sequences for each worker profile. Table 5

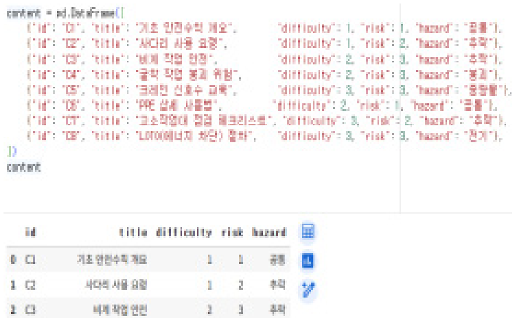

```
content = pd.DataFrame([
    {"id": "C1", "title": "기초 안전수칙 개요",          "difficulty": 1, "risk": 1, "hazard": "공통"},
    {"id": "C2", "title": "사다리 사용 요령",            "difficulty": 1, "risk": 2, "hazard": "추락"},
    {"id": "C3", "title": "비계 작업 안전",              "difficulty": 2, "risk": 3, "hazard": "추락"},
    {"id": "C4", "title": "굴착 작업 붕괴 위험",         "difficulty": 2, "risk": 3, "hazard": "붕괴"},
    {"id": "C5", "title": "크레인 신호수 교육",          "difficulty": 3, "risk": 3, "hazard": "중장비"},
    {"id": "C6", "title": "PPE 실제 사용법",            "difficulty": 2, "risk": 1, "hazard": "공통"},
    {"id": "C7", "title": "고소작업대 점검 체크리스트", "difficulty": 3, "risk": 2, "hazard": "추락"},
    {"id": "C8", "title": "LOTO(에너지 차단) 절차",     "difficulty": 3, "risk": 3, "hazard": "전기"},
])
content
```

```
   id               title  difficulty  risk  hazard
0  C1   기초 안전수칙 개요           1     1    공통
1  C2     사다리 사용 요령           1     2    추락
2  C3       비계 작업 안전           2     3    추락
```

Fig. 5. Eight test modules with difficulty (difficulty 1~3), risk (risk 1~3, and hazard type)

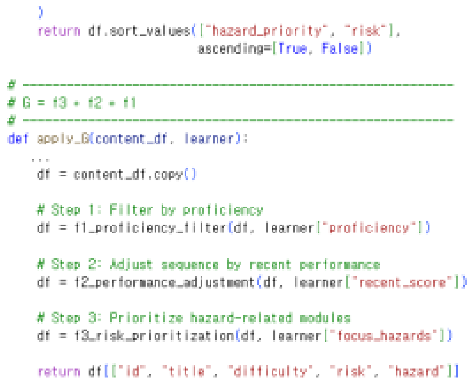

```
    )
    return df.sort_values(["hazard_priority", "risk"],
                          ascending=[True, False])

# ------------------------------------------------------------------
# G = f3 ∘ f2 ∘ f1
# ------------------------------------------------------------------
def apply_G(content_df, learner):
   ...
    df = content_df.copy()

    # Step 1: Filter by proficiency
    df = f1_proficiency_filter(df, learner["proficiency"])

    # Step 2: Adjust sequence by recent performance
    df = f2_performance_adjustment(df, learner["recent_score"])

    # Step 3: Prioritize hazard-related modules
    df = f3_risk_prioritization(df, learner["focus_hazards"])

    return df[["id", "title", "difficulty", "risk", "hazard"]]
```

Fig. 6. Python prototype code example

Table 5. Summary of personalized learning sequences generated for three representative worker profiles.

Worker type	Rank	Content ID	Title	Difficulty	Risk	Hazard type
Low-proficiency worker	1	C3	Scaffold work safety	2	3	fall
Low-proficiency worker	2	C2	Safe ladder use	1	2	fall
Low-proficiency worker	3	C1	Basic safety rules overview	1	1	general
High-proficiency worker	1	C8	LOTO (energy isolation) procedure	3	3	electric
High-proficiency worker	2	C3	Scaffold work safety	2	3	fall
High-proficiency worker	3	C7	Aerial lift inspection checklist	3	2	fall
Fall-focused high-risk worker	1	C3	Scaffold work safety	2	3	fall
Fall-focused high-risk worker	2	C2	Safe ladder use	1	2	fall
Fall-focused high-risk worker	3	C7	Aerial lift inspection checklist	3	2	fall

summarizes the Top-3 recommended modules per worker after applying *G*.

The results confirm that the layered architecture of *G* successfully adapts the curriculum to learner specificities:

Low-Proficiency Adaptation: For the low-skilled worker, the system prioritized easier, fundamental content. As seen in Table 5, the sequence (C3, C2, C1) includes "Basic Safety Rules" and excludes high-difficulty tasks (Level 3), focusing instead on building a safety baseline.

High-Proficiency Adaptation: In contrast, the high-proficiency worker received a sequence dominated by high-difficulty, high-risk modules. The top recommendation, "LOTO Procedures" (C8), is a Level 3 difficulty task, ensuring the training remains challenging and relevant rather than repetitive.

Risk-Based Prioritization: For the worker identified with fall risks, the f3 function successfully reordered the content to prioritize hazard-specific modules. The resulting sequence (C3, C2, C7) exclusively targets fall prevention scenarios related to scaffolding, ladders, and aerial Lifts, demonstrating the system's ability to align training with immediate situational hazards.

The experiment verifies that the personalization mechanism functions as designed, dynamically adjusting content difficulty, ordering, and hazard relevance for heterogeneous worker profiles. This empirical evidence supports the theoretical formulation presented in Section 3.3 and demonstrates the feasibility of integrating rule-based personalization within a generative AI–driven safety education system.

4.4.3. Performance benchmarking

Performance testing under local deployment conditions confirmed that the system meets real-time responsiveness requirements. Average page load time was approximately 180 ms, and API endpoints responded within an average of 20 ms. Chart.js visualizations rendered smoothly with sub-100 ms latency. The server maintained a low memory footprint (~50 MB), even at 100 concurrent sessions, and peak memory usage during stress testing remained within acceptable limits.

Key Performance Highlights:

- Average Page Load Time: 180 ms
- Average API Response Time: 20 ms
- Chart Rendering Time (avg): 71 ms
- Server Memory Footprint: ~50 MB (peak 127 MB under stress)

These results indicate the platform can reliably support interactive learning workflows, multilingual content rendering, and real-time analytics without performance degradation in typical construction-site operational environments.

4.4.4. Browser compatibility testing

Cross-browser testing confirmed full compatibility across all major desktop and mobile browsers, including Chrome, Firefox, Edge, Safari, iOS Safari, and Android Chrome. No functional issues were observed, with only a minor CSS gradient rendering difference noted in Safari. All touch interactions and responsive layouts were successfully verified on mobile platforms.

4.5. Current system limitations

The MVP system has several limitations stemming from its in-memory architecture and early-stage feature set. The database is non-persistent and operates only on a single server, constraining scalability. In terms of functionality, the system does not yet support VR/AR hazard simulations, voice-guided instructions, offline edge-computing capabilities, or video-based training modules, as all content is pre-generated. These limitations define the key areas targeted for future enhancements.

From a pedagogical perspective, the current validation is limited to system usability and functional correctness rather than learning outcomes. This study has not yet longitudinally verified whether AI-generated multilingual content leads to long-term knowledge retention or tangible behavioral changes among foreign workers. Furthermore, while the AI tutor provides linguistic accessibility, it currently lacks the empathetic interaction and nuanced situational judgment provided by experienced human safety instructors, which may limit its effectiveness in fostering a deep safety culture beyond mere rule compliance.

5. Discussion and future work

This study proposed and validated a generative AI-driven safety education architecture specifically designed to address the linguistic diversity, heterogeneous skill distribution, and dynamic risk conditions of contemporary construction sites. Through the implementation of a prototype system and empirical evaluation, several important observations, practical implications, and remaining challenges have emerged. This chapter discusses these findings and outlines key directions for future research.

5.1. Summary of key findings

This study set out to address the critical vulnerability of linguistic rigidity and lack of personalization in construction safety education by proposing a Generative AI-driven architecture. The development and validation of the Minimum Viable Product (MVP) demonstrated that the proposed four-layer architecture is technically feasible and effective in bridging the gap between static training protocols and the diverse needs of a multicultural workforce. The implementation results confirmed that the system could successfully operationalize multilingual content generation, translating safety microlearning modules into four languages while maintaining semantic consistency across safety categories. Furthermore, the empirical evaluation of the personalization function (G) validated the effectiveness of the three-tier adaptation mechanism. As evidenced in the test scenarios, the system successfully differentiated learning pathways; low-proficiency workers were guided toward fundamental safety rules, while high-proficiency workers were challenged with complex tasks like LOTO procedures. Finally, performance benchmarking confirmed that the system meets real-time responsiveness requirements, indicating that the architecture is computationally efficient enough for deployment in resource-constrained construction environments.

5.2. Practical implications

The findings of this study offer several practical implications for field practitioners and safety managers aiming to modernize their training ecosystems. First, the system significantly enhances operational efficiency through automation. Traditionally, developing localized training materials for a multi-national workforce is labor-intensive. This study demonstrates that LLMs can automate the production of multilingual instructional materials, reducing the cost and labor associated with developing culturally appropriate resources. Second, the proposed modular deployment architecture addresses the infrastructural reality of construction sites. By supporting both local edge and cloud execution, the system ensures operational continuity even in environments with limited or intermittent network connectivity. Third, the integration of personalization and multilingual support fosters a technology-driven safety culture that directly promotes social inclusion. By adapting content to the worker’s native language and proficiency level, the system not only mitigates accident risks but also supports the integration of foreign workers into the broader safety management framework.

5.3. Limitations and future work

Despite the demonstrated feasibility, several limitations must be acknowledged to guide future implementation efforts. First, the current MVP relies on an in-memory architecture and a single-server deployment model, which restricts its ability to scale for concurrent users and lacks data persistence. Scaling the system will require scalable databases and container orchestration. Second, while the system generates text and image-based content, it does not yet support the real-time generation of VR/AR simulations or voice-guided instructions due to current latency and cost constraints. Third, the reliance on generative AI introduces potential risks regarding informational accuracy (hallucinations). Although this study implemented guardrails and KOSHA-aligned validation, ensuring absolute reliability in safety-critical contexts remains a challenge. Future iterations must incorporate human-in-the-loop verification processes and reinforcement learning-based refinement to ensure sustained regulatory compliance and content integrity.

6. Conclusions

This study establishes a generative AI-driven safety education framework designed to address the linguistic and structural limitations inherent in traditional training for foreign construction workers. By bridging field-level interactions with cloud-scale generative reasoning, the proposed architecture enables the delivery of dynamic, multilingual, and personalized safety content that evolves alongside real-time site conditions.

The prototype implementation and empirical validation confirmed the system's technical feasibility and practical utility. The architecture successfully demonstrated its capacity to automate the generation of context-aware scenarios and adapt instructional materials to individual proficiency levels. These findings indicate that generative AI can serve as a transformative tool for enhancing the scalability and inclusiveness of construction safety management.

However, technical feasibility does not guarantee educational efficacy. A key limitation of this study is the absence of longitudinal data linking the system’s use to sustained safety behavioral changes or accident rate reductions. Moreover, the current AI-driven approach cannot fully replicate the mentorship and psychological support provided by human instructors.

Ultimately, this research provides a foundational step toward next-generation intelligent safety systems. While challenges regarding data quality and real-world complexity remain, the proposed framework offers a robust roadmap for modernization. Future research will focus on validating the system through long-term field experiments involving actual foreign workers. This evaluation will require tracking key performance indicators (KPIs), such as near-miss reporting frequency and compliance with safety protocols, to quantitatively assess the system's impact on reducing accident rates and improving safety behaviors. Moreover, deeper personalization through reinforcement learning and integration with an IoT-based digital twin ecosystem will be essential for advancing safety education within the construction industry.

Author Contribution

Conceptualization: MH ung and SH Ju; Formal analysis: MHJung, SH Ju and YH Noh; System design: MH Jung, SH Ju, YH Noh, JH Jang; System test: JH Jang and BR Hyun; Writing-original draft: MH Jung, SH Ju, YH Noh and JH Jang; Writing-review & editing: JH Jang and BR Hyun

ORCID

Myunghee Jung: https://orcid.org/0000-0001-5729-3017
Jong-Hwa Jang: https://orcid.org/0000-0001-9044-0218
Younghee Noh: https://orcid.org/0000-0001-8358-6286
Sang-Hyun Ju: https://orcid.org/0000-0001-7975-1146

Conflicts of Interest

The authors declare no conflict of interest.

Funding

This work was supported by the Ministry of Education of the Republic of Korea and the National Research Foundation of Korea (NRF-2023S1A5A2A21083786).

References

1. Noh YH, Jang JH. A study on the perception of foreign construction workers on basic occupational safety and health training in the construction industry. Korean Comparative Government Review. 2024;28(2): 81–112. DOI : 10.18397/kcgr.2024.28.2.81
2. Lee JS, Jeong KT, Jo MJ, Park JH, Lee DH. A Study on the Safety Management of Foreign Workers in Construction Site. Architectural Image Studies. 2024; 6(4):2184-8645. https://doi.org/10.62754/ais.v6i4.402
3. Kim JM, Son KY, Yeom SK, Ahn SJ. Analyzing the risk of safety accidents: The relative risks of migrant workers in the construction industry. Sustainability. 2020;12(13):5430. https://oak.ulsan.ac.kr/handle/2021.oak/10840
4. Ahmed TM, Johanna V, Marcelo M. A Retrospective Analysis of Artificial Intelligence in Education (AIEd) Studies: Perspectives, Learning Theories, Challenges, and Emerging Opportunities. Lecture Notes in Educational Technology. Singapore: Springer; 2024:127-141. https://doi.org/10.1007/978-981-97-8638-1_9
5. Wang P, Wu P, Wang J, Chi HL. Wang X. A critical review of the use of virtual reality in construction engineering education and training. International Journal of Environmental Research and Public Health 2018; 15(6):1204. https://doi.org/10.3390/ijerph15061204
6. Li X, Yi W, Chi HL, Wang X, Chan APC. A critical review of virtual and augmented reality (VR/AR) applications in construction safety. Automation in Construction. 2018;86:150–162. https://doi.org/10.1016/j.autcon.2017.11.003
7. Kasneci E, Sessler K, Küchemann S, Bannert M, Dementieva D, Fischer F, et al. ChatGPT for good? On opportunities and challenges of large language models for education. Learning and Individual Differences. 2023;103. https://doi.org/10.1016/j.lindif.2023.102274
8. Bewersdorff A, Hartmann C, Hornberger M, Seßler K, Bannert M, Kasneci E, et al. Taking the next step with generative artificial intelligence: The transformative role of multimodal large language models in science education. Learning and Individual Differences. 2025;118. https://doi.org/10.1016/j.lindif.2024.102601
9. Khosravi H, Shibani A, Jovanovic J, Pardos ZA, Yan L. Generative AI and learning analytics. Journal of Learning Analytic 2025;12(1). https://doi.org/10.18608/jla.2025.8961
10. Pan Y, Zhang L. Roles of artificial intelligence in construction engineering and management: A critical review and future trends. Automation in Construction. 2021;122. https://doi.org/10.1016/j.autcon.2020.103517
11. Awolusi I, Marks E, Hallowell M. Wearable technology for personalized construction safety monitoring and trending: Review of applicable devices. Automation in Construction. 2018;85:96–106. https://doi.org/10.1016/j.autcon.2017.10.010
12. Pyo AJ, Lee EY. Yoon SH. Identifying industrial safety issues among foreign workers in South Korea: A BERTopic and network-based analysis of news article data under the Employment Permit System. Asia Pacific Journal of Information Systems. 2025;35(2):367

–388. https://doi.org/10.14329/apjis.2025.35.2.367

13. Naveen P, Trojovsky P. Overview and challenges of machine translation for contextually appropriate translations. iScience. 2024;27(10). https://doi.org/10.1016/j.isci.2024.110878
14. Jiao W, Wang W. Huang JT, Wang X, Shi S, Tu Z. Is ChatGPT a good translator? Yes With GPT-4 As The Engine. ArXiv. 2023; arXiv:2301.08745. https://doi.org/10.48550/arXiv.2301.08745
15. Xu S, Ni QQ, Zhang M, Li M. A personalized safety training system for construction workers. In International Conference on Smart Infrastructure and Construction 2019 (ICSIC). https://www.icevirtuallibrary.com/doi/10.1680/icsic.64669.321
16. Varshney KR. Trustworthy machine learning and artificial intelligence. XRDS: Crossroads, The ACM Magazine for Students. 2019;25(3):26–29. https://doi.org/10.1145/3313109
17. Ojghaz AS, Sadeghpour F, Bayat S. Edge Computing for real-time monitoring systems in construction sites. Proceedings of the canadian society for civil engineering annual conference. 2024;3:697. https://doi.org/10.1007/978-3-031-97697-1_24
18. Mosqueira-Rey E, Hernández-Pereira E, Alonso-Ríos D. Bobes-Bascarán J, Fernández-Leal A. Human-in-the-loop machine learning: A state of the art. Artificial Intelligence Review 2023;56:3005–3054. https://doi.org/10.1007/s10462-022-10246-w
19. Chen Q, Yin Z, Yuan B, Chen Q. Personalized safety training for construction workers: A large language model-driven multi-agent framework integrated with knowledge graph reasoning. Computers in Industry. 2026;174. https://doi.org/10.1016/j.compind.2025.104399
20. Alwashah Z, Xiao B, Liu H, Mueller AT, Shao X. Generative artificial intelligence for construction: Use cases, trends, challenges, and opportunities. Journal of Building Engineering. 2025;112(15). https://doi.org/10.1016/j.jobe.2025.113802
21. Pedro A, Pham-Hang A.T, Nguyen P.T, Pham H.C. Data-driven construction safety information sharing system based on linked data, ontologies, and knowledge graph technologies. International Journal of Environmental Research and Public Health 2022: 19(2): 794. https://doi.org/10.3390/ijerph19020794
22. Lengetti A, Kronk R, Cantrell MA. A theory analysis of mastery learning and self-regulation. Nurse Education in Practice. 2020;49:102911. https://doi.org/10.1016/j.nepr.2020.102911