

Agentic AI for Autonomous Construction Issue Detection: A Multimodal Framework Using Model Context Protocol

Simge Girgin, Ph.D.*¹; Fopefoluwa Bademosi, Ph.D.²; Parastoo Varshosaz³; and Damon Lau*⁴

¹ Autodesk Research, 1 Market St #400, San Francisco, CA 94105; ORCID: <https://orcid.org/0000-0001-9746-1388>; e-mail: simge.girgin@autodesk.com, simge@stanford.edu

² Autodesk Research, 23 Drydock Ave, Boston, MA 02210; ORCID: <https://orcid.org/0000-0002-6046-9359>; e-mail: fope.bademosi@autodesk.com

³ WZMH Architects; Giraffe Apps Inc., 95 St. Clair Avenue West, Suite 1500, Toronto, ON M4V 1N6, Canada; ORCID: <https://orcid.org/0009-0004-8850-4429>; e-mail: pvarshosaz@wzmh.com

⁴ Autodesk Research, 19 Morris Ave, Building 128, Brooklyn, NY 11205; ORCID: <https://orcid.org/0000-0001-5173-0378>; email: damon.lau@autodesk.com

* Corresponding Author

ABSTRACT

Quality management remains a persistent challenge in the construction industry, with rework accounting for significant cost growth and schedule delays. While digital inspection technologies have improved data capture, defect identification still relies on manual, error-prone processes. This paper introduces an agentic AI framework, implemented through the Model Context Protocol (MCP), for autonomous construction quality management. The framework integrates retrieval-augmented reasoning over specifications and drawings, multimodal site imagery, and structured issue generation through orchestrated large language model (LLM) agents. A case study on an office renovation project demonstrates the feasibility of the approach. The system accurately retrieved construction information from specifications, achieving an average cosine similarity of 0.83 on queries compared to manual responses and successfully identified 9 of 10 documented quality issues, including design misalignments and field non-conformances. These findings provide empirical evidence that agentic AI can extend beyond descriptive reporting toward actionable quality management. The contributions are twofold: (i) a modular MCP-based architecture for orchestrating construction AI agents, and (ii) validation through deployment in a live project environment. While evaluation is limited in scope, results highlight both the promise and challenges of integrating agentic AI into daily workflows, paving the way for larger-scale studies and future digital twin integration.

INTRODUCTION

A significant challenge in the construction industry is its costly and error-prone approach to quality management, which is regarded as one of the most critical pain points in the sector. Rework can account for up to 52% of total project cost growth (Ayalp & Arslan, 2025) and leads to frequent schedule overruns. Nearly 98% of major construction projects experience significant delays, averaging 20 months, and cost increases of approximately 80% (Changali et al., 2021). To improve quality control and reduce rework, many firms have begun digitizing inspections using 360° cameras, drones, and autonomous robots (Viact.ai, 2024; PlanRadar, 2024; OpenSpace, 2024).

However, identifying defects requires manual oversight, a time-consuming process prone to inconsistency.

These limitations highlight the need for automation that extends beyond data capture to decision-making. Recent advances in agentic artificial intelligence (AI) offer a new paradigm. Agentic AI refers to a system of specialized AI agents that collaborate to complete complex tasks without continuous human oversight (Sapkota et al., 2025). Early implementations in adjacent industries have demonstrated that agentic AI systems have achieved significant operational improvements. For instance, in semiconductor manufacturing, multi-agent reinforcement learning reduced machine delays and energy use by over 60% (Zhang et al., 2025), while in automotive assembly, agent-based coordination systems cut response times to part supply disruptions by similar margins (Montoya-Zapata et al., 2024). Despite this momentum, construction has yet to realize the full potential of agentic AI.

This study presents an agentic AI framework using Model Context Protocol (MCP) (Anthropic, 2024), which integrates multiple large language model (LLM) agents capable of reasoning across multimodal construction data, including site imagery, drawings, specifications, and reports. Deployed on a commercial office renovation case, the system autonomously reviewed documentation, analyzed site images, and identified quality issues due to design misalignments and field non-conformances. The framework coordinated agent workflows, applied vision transformers for defect detection, and used retrieval-augmented generation (RAG) to ground reasoning in project-specific contexts. Together, these components achieved a cosine similarity of 0.83 on information retrieval queries and were able to detect 8 out of 10 RFIs. These results provide insights for project-based deployment of AI agents and pave the way for further evaluation using a larger dataset.

BACKGROUND AND RELATED RESEARCH

The growing maturity of agentic AI frameworks, such as MCP, LangGraph, CrewAI, and AutoGen, enables multi-step reasoning, persistent memory, and role-based collaboration—capabilities well-aligned with the complex and dynamic nature of construction workflows (LangChain, 2024; Assouline, 2024; Wu et al., 2023). While general AI research highlights the shift from reactive assistants to proactive, collaborative agents (Wang et al., 2024), construction AI remains largely focused on single-purpose applications, with few examples of multi-agent systems that address integrated, end-to-end workflows (Pan and Zhang, 2023).

AI Applications for Construction Data Retrieval. Recent advancements in AI for construction management have focused on improving access to complex project information. Wu et al. (2025) proposed a construction-tailored RAG pipeline that retrieves and reasons over complex documents via hierarchical knowledge modeling and preference learning, thereby improving communication efficiency. However, it primarily operates as a retrieval chatbot rather than an autonomous issue detector. Similarly, Pal et al. (2025) introduced a schedule management interface that utilizes knowledge graphs and natural language processing (NLP) for real-time, natural-language queries, thereby enhancing engagement with schedule data while maintaining a retrieval-centric approach. However, this tool focuses on information retrieval without cross-referencing site conditions or integrating into automated issue detection workflows.

AI for Visual Monitoring and Construction Documentation. Advancements in computer vision have also contributed to the development of AI-driven construction monitoring. Yang et al. (2024) developed a vision transformer (ViT)-based visual question answering system designed to extract information from construction site images, such as workforce levels, activities, and

progress. This system demonstrates the potential of multimodal AI for generating structured construction logs through real-time site monitoring. However, the system remains focused on descriptive reporting and lacks integration with field quality assessment. Similarly, Yung et al (2024) developed a ViT-based image captioning model that reviews field imagery and reports daily progress and work activity log. Although the model offers a way to retrieve field images precisely, the detection of issues within the retrieved imagery was not within the scope of analysis.

Gaps and Opportunities. Existing approaches emphasize retrieval and descriptive reporting, but none directly integrate AI agents into construction workflows for autonomous quality management. This study addresses that gap by introducing an end-to-end framework that combines retrieval, multimodal analysis, and agentic orchestration to monitor site conditions, cross-reference documentation, and generate actionable quality records.

FRAMEWORK DEVELOPMENT AND CASE STUDY ANALYSIS

The methodology combines a proposed agentic AI framework with a case study evaluation. The framework defines how construction quality issues are detected through retrieval-augmented reasoning, multimodal inspection, and structured reporting under the MCP. The case study demonstrates its application on a live office renovation project, providing real-world validation of the approach.

Agentic AI Framework for Construction Quality Agent. The proposed framework integrates three core capabilities: (i) retrieval-augmented reasoning over project documentation, (ii) image-based inspection of construction progress, and (iii) automated generation of structured quality issues. These capabilities are orchestrated within the MCP, providing a modular architecture where tools are dynamically invoked in response to evolving site conditions and user queries (Figure 1).




Figure 1: MCP architecture for the construction quality agent.

Retrieval-augmented reasoning. Construction specifications, drawings, and related documentation are hierarchically indexed using vector embeddings. At inference, the agent first identifies relevant documents and then narrows the retrieval to specific pages or sheets using similarity thresholds and drawing relationship links. Drawings are rendered as raster images, summarized via optical character recognition (OCR), and linked to metadata such as titles and sheet numbers, enabling context-aware access to project requirements.

Image-based inspection. Site photographs are embedded using a vision language model (VLM) and enriched with metadata such as timestamps and geolocation. These embeddings are stored in a parallel index to the document corpus. At runtime, the system retrieves both textual and visual context, integrating them into structured prompts that encode behavioral rules for traceable, evidence-based analysis.

Automated generation of structured quality issues. The construction quality agent sequentially invokes four MCP tools. The “*analyze_drawings*” tool inspects multiple construction drawings to identify the installation requirements for building components. The “*analyze_construction_image*” tool interprets photos to detect building components, activities, and materials. The “*evaluate_installation_quality*” tool compares observed conditions against retrieved documentation, identifying deviations from design intent. Finally, the “*report_quality_issue*” tool prioritizes design misalignments and field non-conformances by severity, assigns trade responsibilities, and generates structured records with supporting evidence and references, which are then submitted to the project’s management platform.

The framework operates in a modular workflow integrating retrieval pipelines, vector databases, LLMs, and MCP tools. Large-model embeddings, dynamic prompt construction, and asynchronous tool invocations enable flexible sequencing as site conditions evolve. This reduces the need for manual inspection, improving consistency and traceability to documentation, while supporting scalable quality management workflows. This modular architecture enables construction firms to deploy AI incrementally, starting with high-risk areas and expanding as confidence grows.

Case Study Evaluation. The evaluation was conducted through a nine-month case study of a 16,000-square-foot office renovation project in collaboration with WZMH Architects in Toronto, Ontario. This real-world construction environment enabled continuous on-site data collection and iterative testing of the proposed agentic AI framework. Data collection consisted of weekly site visits by the project architect, during which construction progress was documented through 360-degree videos, written reports, and formal construction documentation, including requests for information (RFIs) and change orders. Additionally, the contractor provided weekly site walkthroughs, which were captured via OpenSpace, contributing interactive 360-degree imagery to the dataset.

The collected visual data were analyzed using GPT-5 (OpenAI, 2025) to extract construction features, including built objects, construction equipment, and materials. The extracted features from the images were subsequently compared against corresponding representations in 2D and text-based documentation in the drawings and specifications. Due to the absence of accurate location tracking for interior construction sites, geotags were manually assigned to images to cross-reference photos of physical construction to their digital counterparts. The orchestration of multiple LLM agents through the MCP was necessary to process and cross-reference multimodal data and establish reasoning across the siloed data sources.

The evaluation employed 25 structured queries to assess the effectiveness of the construction quality agent across text-only and multimodal scenarios. Fifteen queries focused on text-based documentation retrieval and analysis, while 10 evaluated multimodal performance involving drawings, images, and associated metadata derived from project RFIs and change orders. These 10 RFIs were selected from a total of 40 project RFIs, representing cases of design misalignment or field non-conformance. The remaining 30 RFIs were excluded from evaluation as they related

to client-initiated design changes or unforeseeable supply chain issues outside the scope of quality assessment.

To evaluate text-based documentation retrieval and analysis, 15 ground truth answers were established by a panel of construction managers, architects, and researchers with access to the project documentation. The LLM’s response to each question was evaluated against the ground truth answers based on their semantic similarity. Cosine similarity was used as a metric of semantic similarity, assessing how closely the generated response mimics the structure and content of the ground truth answers. The OpenAI text embedding model, text-embedding-large (OpenAI, 2024), was utilized to vectorize both the ground truth answers and the LLM’s response for the calculation of cosine similarity. Next, the weighted average cosine similarity was calculated, with scores ranging from 0 (no similarity) to 1 (perfect similarity).

To evaluate the system’s performance in identifying issues using multimodal data, 10 RFIs out of 40 were selected. The remaining RFIs were excluded since they were related to design changes that were unrelated to issues. The 10 RFIs listed in Table 1 describe five misalignments across the architectural, mechanical, electrical, and plumbing drawing sets that were detected during design coordination, and five non-conformances that were detected during field installation. For each issue, the system’s outputs were reviewed by the project team, who either accepted or rejected the LLM’s issue report.

Table 1. List of RFIs selected for evaluating the AI construction quality agent.

ID	Issue Description	Issue Type
RFI 2	The solid wall is not permissible to install faucets and conduits	Design Coordination
RFI 3	HVAC diffuser will not fit into the available ceiling space	Design Coordination
RFI 5	HVAC line is below the preferred ceiling height	Field non-conformance
RFI 6	Furring out a column for electrical works causes misalignment with drywall	Design Coordination
RFI 14	Artwork supports and wood blocking are missing	Field non-conformance
RFI 22	Light is obstructed by the existing exhaust system	Field non-conformance
RFI 25	Access panel locations are missing	Design Coordination
RFI 30	Ceiling speaker is obstructed by overhead piping	Field non-conformance
RFI 36	Gap between demountable partitions and ceiling	Field non-conformance
RFI 39	The number of faucets in kitchen mismatch across drawings	Design Coordination

RESULTS

The findings are presented in two parts: (A) the agent’s ability to retrieve construction specifications through RAG, and (B) the evaluation of multimodal construction quality assessment through the proposed MCP-based workflow.

Retrieval of Construction Specifications through RAG. Construction specifications and drawing sets issued by clients served as the primary reference for identifying and resolving conflicts during the construction process. Before evaluating whether issues can be autonomously identified, it was necessary to assess the agent’s capability to retrieve relevant specifications. The evaluation employed 15 representative queries, which were developed based on typical questions asked by designers or engineers during design coordination and field inspection. These queries

spanned various disciplines (architectural, electrical, mechanical, plumbing, and fire protection) and building activities (demolition, installation, and inspection).

In this preliminary evaluation, cosine similarity scores ranged from 0.66 to 0.94 (mean = 0.82, median = 0.83). While the sample size is limited, these results suggest promising retrieval performance that warrants validation with larger datasets. Importantly, in all 15 cases, the AI agent retrieved the correct specification page references, typically returning one to four pages per query while summarizing retrieved information. This provides initial evidence of feasibility and practical utility for project applications.

Multimodal Construction Quality Assessment. The construction quality agent was evaluated within the proposed MCP architecture using 10 RFIs, divided into five design misalignments and five field non-conformances.

Identification of design misalignments in construction drawings and specifications. To test the agent on design misalignments, construction drawings (architectural, mechanical, electrical, and plumbing) together with project specifications were utilized. First, the VLM processed each drawing in conjunction with the corresponding specifications to examine building elements and extract requirements. Second, it compared building elements across drawings to identify potential coordination gaps.

- **RFI 2 (electrical boxes and conduits):** The AI agent retrieved the specifications for wall-mounted electrical boxes and conduits. It then inspected the millwork plan, partition details, and power/communication layouts for the office server area. The agent responded: *“Multiple power/communication outlets are located on the wall behind the sink run, but the partition plan/details depict this as a solid masonry/concrete wall with no furred stud cavity. A solid wall does not provide the depth needed for in-wall faucet valves/supplies or for the electrical/data boxes and conduits called out on the millwork plan.”*
- **RFI 3 (mechanical diffuser in IT room):** The AI agent examined the constructability of a specified diffuser. It responded: *“There’s a clash. The IT room has only 12 in. plenum depth (9 ft ceiling to 10 ft joists), but F-1200 is specified with a 24 × 24 × 14 in. plenum (Note 19), which cannot fit, and a neck damper per Note 11 would need additional depth.”*
- **RFI 6 (column–partition misalignment):** The AI agent failed to identify the issue where a column and a parallel partition wall aligned in architectural drawings but did not align in the field due to different partition types. The hypothesis is that the false visual alignment in the drawings led the agent to miss this issue.
- **RFI 25 (missing access panel details):** After determining the locations of access doors modeled in the reflected ceiling plan, the AI agent responded: *“Access panel locations are not specified in the drawings.”*
- **RFI 39 (kitchen/server faucet):** The AI agent analyzed two drawing sets (architectural millwork and plumbing piping/drainage) and one set of specifications with plumbing schedules. It responded: *“The architectural drawings show three faucets (two regular, one filtered) for the S-1 sink, whereas the plumbing drawings specify only one faucet and one filtered faucet.”* (Figure 2)

Overall, the construction quality agent successfully identified 4 of 5 design misalignments. These examples demonstrate its capacity to flag coordination gaps that are often overlooked, enabling earlier interventions and reducing the likelihood of downstream rework.




Figure 2: Example identification of design misalignment in RFI 39 across construction drawings

Field non-conformances. The second evaluation focused on the construction quality agent's ability to detect non-conformances in field conditions. This assessment combined architectural detail drawings with corresponding site photographs. In each case, the VLM first extracted requirements documented in text and CAD representations, then compared them against field photos to determine whether requirements were satisfied.

- **RFI 5 (boardroom HVAC installation):** The AI agent extracted HVAC installation requirements from demolition and new installation drawings. From a ceiling photo, it responded: *"A rectangular insulated duct is installed too low beneath the structure, which will conflict with the planned boardroom ceiling plane and lighting."* This issue had initially been missed by the general contractor.
- **RFI 14 (conference room wall finish):** The AI agent identified four requirements for the wall: joint lines to be mudded, substrate suitable for the specified finish, panel alignment, and blocking for wall-mounted fixtures. On analysis of the input photo, it responded: *"None of the four requirements have been met."* (Figure 3).
- **RFI 22 (mechanical-electrical conflict in open office):** The AI agent inspected existing overhead mechanical components and new electrical layouts. It responded: *"There is a conflict between the existing exhaust for the old printer system and the ceiling lights."*
- **RFI 30 (ceiling speaker installation):** The AI agent first reported the observed building elements and activities in the field photo. It then retrieved the ceiling speaker submittal and responded: *"Speaker rough-in conflicts with vertical pipe; no 5½ in. (140 mm) unobstructed clearance."* This case illustrates how continuous monitoring can both flag potential issues and provide diagnostic guidance to help crews resolve them more quickly.
- **RFI 35 (demountable partition installation):** The AI agent retrieved plan details showing required acoustic sealant at wall intersections. It then assessed installation

images and responded: “A continuous dark gap is visible at the top of the white partition and along the plywood wall where they meet the ceiling/track. No acoustical (or fire) sealant is evident. This conflicts with the drawings calling for acoustic sealant on both sides at perimeter joints and is a likely acoustic flanking path.”

Across these tests, the AI agent accurately detected all five field non-conformances. Combined with the four out of five design misalignments identified, the evaluation demonstrates that the construction quality agent can both retrieve correct specification page references and apply them in multimodal reasoning over drawings and site imagery. These findings highlight the feasibility of automated issue detection through cross-referencing specifications, drawings, and field photographs. The scope of the study is limited; however, the consistency of results across both retrieval and multimodal tasks indicates that agentic AI can support design coordination and field quality inspection in real-world construction projects. The evaluation suggests that such systems have the potential to reduce overlooked coordination gaps, provide earlier issue diagnostics, and streamline communication between field and design teams, thereby reducing downstream rework and improving overall project efficiency.




Figure 3. Determine field compliance based on construction requirements.

CONCLUSION

This study presents an MCP architecture for an AI quality agent to continuously retrieve, analyze, and document construction quality issues. Specifically, it contributes (1) MCP architecture in a construction context, (2) an evaluation through a case study on an office renovation project, which establishes (3) the feasibility of agentic AI for construction issue identification, analysis, and management.

In terms of performance, the RAG component achieved a mean cosine similarity of 0.83 across 15 representative queries, with correct specification pages retrieved in all cases, demonstrating reliable specification retrieval. For multimodal construction quality assessment, the agent was evaluated on 10 RFIs covering both design misalignments and field non-conformances. It successfully identified 4 out of 5 design misalignments and 5 out of 5 field non-conformances, accurately cross-referencing specifications, drawings, and field imagery to detect coordination

gaps and quality issues. These findings highlight both the potential and current limitations of agentic AI in construction quality workflows.

While the evaluation is limited in scope and draws on a small dataset, it provides empirical insights into the operation of agentic workflows under real project conditions. Future research directions will expand the evaluation through additional case studies, integrate augmented/virtual reality for linking geotagged field images with BIM data, and explore integration with digital twins of construction.

The results provide practical guidelines for designing agentic workflows that target design coordination and quality inspection. Embedding automated AI tools into daily project workflows can surface issues that are often overlooked, providing instant analysis of the issue and thereby equipping designers and engineers with the information to resolve them faster and with fewer downstream changes. More broadly, the approach suggests a new mode of addressing construction quality that augments and partially automates processes that were previously manual, thereby improving construction management productivity.

ACKNOWLEDGMENTS

The authors gratefully acknowledge the support of the Intelligent Construction team at Autodesk Research, particularly Brian Lee, Fikret Kalay, James Donnelly, Jamie Nicholson, Josh Aigen, Leonardo Rodriguez, Matthieu Gillet, Merry Wang, Navita Singh, Nikhil Dhamne, Peter Storey, Prudence Robinson, Richard Pedley, Tomas Lay Herrera, Xin Yang, for their contributions to the MCP framework implementation and evaluation. Special thanks to WZMH Architects, specifically Zenon Radewych, Ted DuArte, and Nazanin Salimi, for providing access to the office renovation project and for their valuable feedback and domain expertise throughout the case study. Initial investigations into RAG for construction documentation were conducted in collaboration with the Berkeley SAAS team of Tyler Zhao, Brandon Xu, Akhil Agarwal, Alex Xu, Aliva Ding, Christing Prieto, Jinn Lim, Mithil Srungarapu, Rithwik Nukala, and Sanay Bordia.

REFERENCES

- Anthropic (2024). “Introducing the Model Context Protocol.” Retrieved from <https://www.anthropic.com/news/model-context-protocol>
- Assouline, J. (2024). “CrewAI: Framework for orchestrating role-playing, autonomous AI agents.” Retrieved from <https://github.com/joaomdmoura/crewAI>
- Changali, S., Mohammad, A., & Nieuwland, M. (2021). “The construction productivity imperative.” McKinsey & Company. <https://www.mckinsey.com/business-functions/operations/our-insights/the-construction-productivity-imperative>
- Gumusburun Ayalp, G., & Arslan, F. (2025). “Modeling Critical Rework Factors in the Construction Industry: Insights and Solutions.” *Buildings*, 15(4), 606. <https://doi.org/10.3390/buildings15040606>
- Jung, Y., Cho, I., Hsu, S. H., & Golparvar-Fard, M. (2024). “VisualSiteDiary: A detector-free Vision-Language Transformer model for captioning photologs for daily construction reporting and image retrievals.” *Automation in Construction*, 165, 105483. <https://doi.org/10.1016/j.autcon.2024.105483>

- LangChain. (2024). "LangGraph: Build resilient language agents as graphs." Retrieved from <https://github.com/langchain-ai/langgraph>
- Montoya-Zapata, S., Klement, N., Silva, C., Gibaru, O., & Lafou, M. (2024). "Multi-agent system for perturbations in the kitting process of an automotive assembly line." *Engineering Applications of Artificial Intelligence*, 135, 108679. <https://doi.org/10.1016/j.engappai.2024.108679>
- OpenAI. (2024). "New embedding models and API updates. OpenAI." Retrieved from <https://openai.com/index/new-embedding-models-and-api-updates/>
- OpenAI. (2025). "Introducing GPT-5." Retrieved from <https://openai.com/index/introducing-gpt-5/>
- OpenSpace. (2024). "AI-powered 360° construction site capture and progress tracking." <https://www.openspace.ai>
- Pal, A., Murguia, D., and Middleton, C. (2025). "AI-powered application for construction schedule management using natural language." *Proc., European Conf. on Computing in Construction (EC³)*.
- Pan, Y., & Zhang, L. (2023). "Integrating BIM and AI for smart construction management: Current status and future directions." *Archives of Computational Methods in Engineering*, 30(2), 1081-1110. <https://doi.org/10.1007/s11831-022-09830-8>
- PlanRadar. (2024). "Enhancing inspections with digital documentation tools." <https://www.planradar.com>
- Sapkota, R., Roumeliotis, K. I., & Karkee, M. (2025). "AI agents vs. agentic AI: A conceptual taxonomy, applications and challenges." arXiv preprint arXiv:2505.10468. <https://arxiv.org/abs/2505.10468>
- Viact.ai. (2024). "AI-driven construction monitoring using 360-degree imagery." <https://www.viact.ai>
- Wang, L., Ma, C., Feng, X., Zhang, Z., Yang, H., Zhang, J., Chen, Z., Tang, J., Chen, X., Lin, Y., Zhao, W. X., Wei, Z., Wen. J. (2024). "A survey on large language model based autonomous agents." *Frontiers of Computer Science*, 18(6), 186345. <https://doi.org/10.1007/s11704-024-40231-1>
- Wu, C., Ding, W., Jin, Q., Jiang, J., Jiang, R., Xiao, Q., Liao, L., Li, X. (2025). "Retrieval augmented generation-driven information retrieval and question answering in construction management." *Advanced Engineering Informatics*, 65, 103158. <https://doi.org/10.1016/j.aei.2025.103158>
- Wu, Q., Bansal, G., Zhang, J., Wu, Y., Li, B., Zhu, E., Jiang, L., Zhang, X., Zhang, S., Liu, J., Awadallah, A. H., White, R. W., Burger, D., & Wang, C. (2023). "AutoGen: Enabling next-gen LLM applications via multi-agent conversation." arXiv preprint arXiv:2308.08155. <https://doi.org/10.48550/arXiv.2308.08155>
- Yang, B., Zhang, B., Han, Y., Liu, B., Hu, J., & Jin, Y. (2024). "Vision transformer-based visual language understanding of the construction process." *Alexandria Engineering Journal*, 99, 242-256. <https://doi.org/10.1016/j.aej.2024.05.015>
- Zhang, P., Jin, M., Wang, M., Zhang, J., He, J., & Zheng, P. (2025). "A clustering-aided multi-agent deep reinforcement learning for multi-objective parallel batch processing machines scheduling in semiconductor manufacturing." *Measurement and Control*, 58(5), 614–631. <https://doi.org/10.1177/00202940241269643>