Dealing with Deep Uncertainties in Digital Twining with Adaptive Agentic AI

ZhiQiang Chen [1] Sailor Usher [2] Chuheng Xiao [3]

[1]Professor of Civil Engineering [2]AR/VR Coordinator, Graduate Student, Computer Science [3]Graduate Student, Data Science and Analytics
University of Missouri – Kansas City

Introduction: Deep Epistemic Uncertainties in Digital Twinning

Digital twinning of civil infrastructure, spanning design, construction, maintenance, operation, and emergency recovery, confront *deep uncertainties* on multiple fronts. At a macro level, these uncertainties include:

- **Physical:** e.g., climate change altering environmental loads.
- **Technical:** e.g., the adoption of emerging technologies that affect simulation fidelity.
- **Socioeconomic:** e.g., policy shifts and planning decisions that impose subjective constraints.

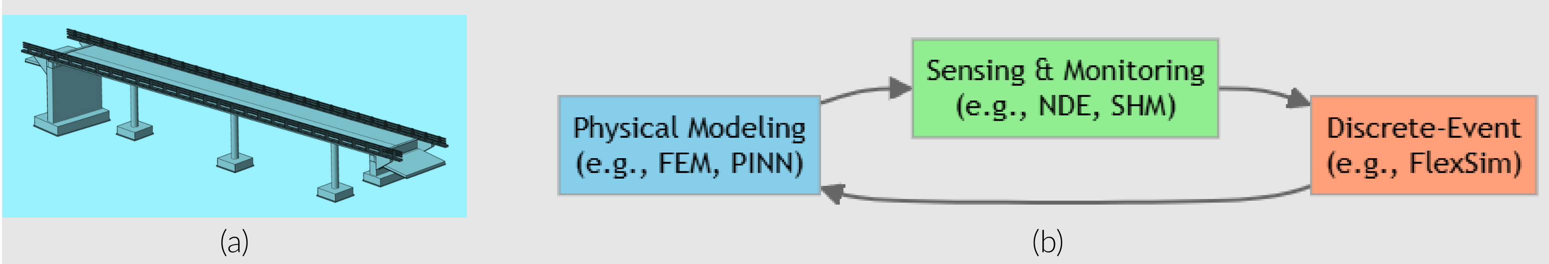

Figure 1. Digital twinning for civil structures: (a) life-cycle bridge as an example; (b) interacting foundational blocks of digital twinning.

At a micro level, many deep uncertainties arise from an overlooked domain - problems that are inherently NP-hard, where global optimal decisions are effectively "lost" in combinatorial complexity. This domain is characterized by a myriad of interwoven logical rules (e.g., design specifications, inspection manuals, and assessment guidelines) that have evolved from decades of civil engineering research. These deep epistemic uncertainties (DEUs) may partly result from professional practices, but they fundamentally stem from a lack of advanced computational intelligence. Consequently, the common approach of relying on engineers or stakeholders to apply rules, logic, and contextual judgment to navigate these complexities is widely accepted yet intrinsically flawed.

Therefore, we assert that while current digital twin architectures are built on three well-developed foundational blocks (Fig. 1), they lack an integrated layer for processing deep epistemic uncertainties and effectively interfacing with human experts.

Adaptive Agentic AI Framework

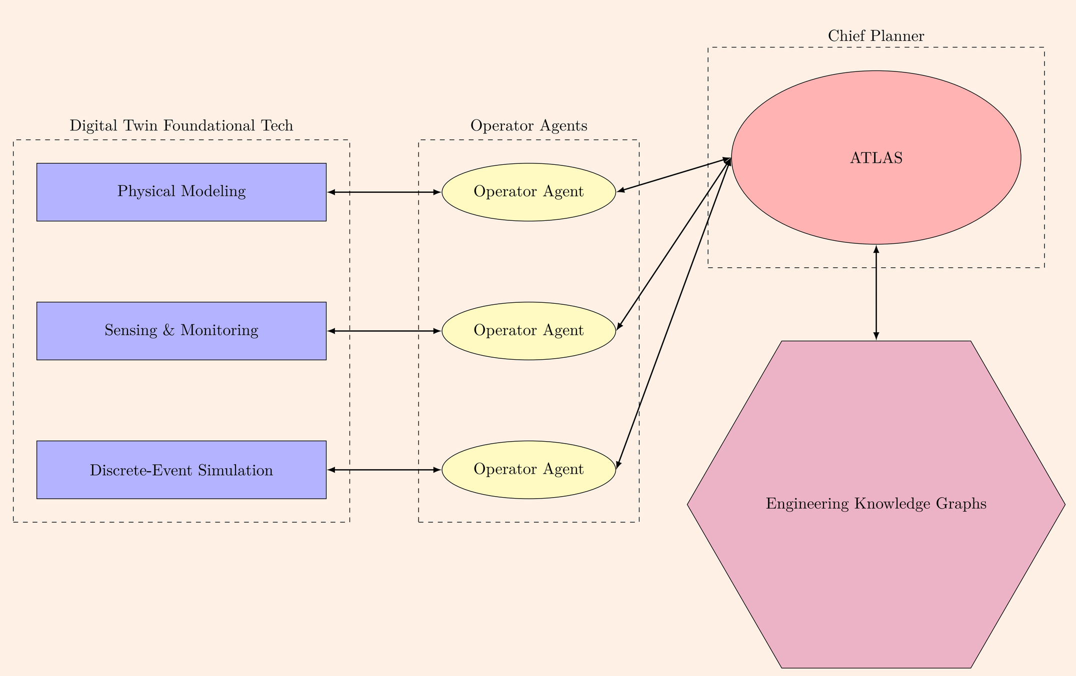

Figure 2. Adaptive Agentic AI interfacing digital twinning, engineering knowledge, and human users.

We propose an **Adaptive Agentic AI Workflow** that integrates a hierarchical, adaptive decision-making process into digital twinning (Fig. 2):

- **ATLAS Integration:** The chief planner agent, **ATLAS: Adaptive Tactical Leadership and Agentic Symbiosis**, leverages symbolic AI (including STRIPS-based planning) to formulate robust high-level strategies and directly interfaces with human users to incorporate expert knowledge and contextual judgment.
- **Operator Agents:** Generative AI-enabled operator agents interact with the digital twin's foundational technologies—physical modeling, sensing and monitoring, and discrete-event simulation—to continuously simulate, update, and adapt system representations in real time.

This framework is uniquely suited to manage multi-modal deep epistemic uncertainties in digital twinning through:

- **Hierarchical Adaptivity:** A dual-layer structure that provides both strategic oversight and operational flexibility.
- **Human-AI Symbiosis:** Direct interaction with human experts ensures that contextual insights and domain-specific nuances are embedded within the decision-making process.

In this preliminary work, we showcase:

1. ATLAS for implementing a digital twinning-assisted structural inspection process.
2. Implementation of an Operator Agent that transforms engineering documents into knowledge graphs for design automation.

Digital Twinning Enabled Structural Inspection

Structural inspection has long been a critical yet predominantly manual process for the routine assessment of structures. To date, engineers employ advanced equipment—such as drones, LiDAR, and various nondestructive evaluation methods—to perform inspections. Although thorough inspection, marking, and coding steps are well documented (Fig. 3 [3]), executing a standardized, systematic plan in the field remains a significant challenge.

Our proposed methodology leverages the Adaptive Agentic AI Framework to transform this process:

- **ATLAS Integration:** The Chief Planner, **ATLAS**, interfaces with an inspection-centric Knowledge Graph to gather context-specific information and best practices.
- **AI-Driven Execution:** AI-enabled Operator Agents seamlessly execute these plans by coordinating various physical and discrete-event simulations, thereby enhancing the accuracy and efficiency of structural assessments.

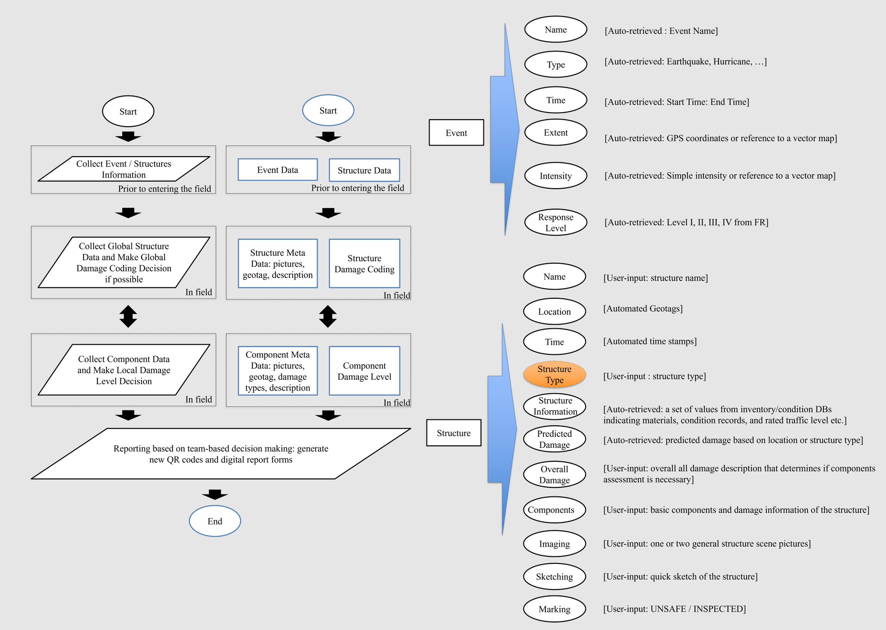

Figure 3. Curated engineering knowledge: structural inspection, marking, and coding for highway bridges.

This integrated approach aims to standardize and optimize structural inspection processes within a digital twin environment, effectively managing the deep uncertainties inherent in field engineering operations.

STRIPS Modeling for Optimal Structural Inspection Planning

ATLAS, our chief planner, generates a STRIPS-based planning model that defines the sequential actions and state transitions required to transform raw inspection data into actionable reports (a partial code snippet is shown below):

```
;; Predicates representing workflow states (:predicates
(event_data_collected)(structure_data_collected)(global_damage_decided)(component_data_collected)(local_damage_decided)(report_generated))
;; ACTION A: Collect event/structure info (:action collect_event_structure_info :parameters() :
precondition(and(not(event_data_collected))(not(structure_data_collected))) :effect(and(event_data_collected)(structure_data_collected)))
;; ACTION B: Global damage decision (:action global_structure_damage_decision :parameters() :
precondition(and(event_data_collected)(structure_data_collected)(not(global_damage_decided))) :effect(global_damage_decided))
;; ACTION C: Local damage decision (:action local_component_damage_decision :parameters() :
precondition(and(global_damage_decided)(not(component_data_collected))(not(local_damage_decided))) :
effect(and(component_data_collected)(local_damage_decided)))
;; ACTION D: Reporting (:action
generate_final_report :parameters() :precondition(and(local_damage_decided)(not(report_generated))) :effect(report_generated)))
```

Advanced algorithms for optimal planning (e.g., reinforcement learning [1]) can ensure that the plan adapts dynamically to field conditions, thereby facilitating robust structural assessment and reporting.

Finite-state Machine Simulation of Inspection

We envision that, for training and interactive augmentation of digital twinning-enabled structural inspection, a specialized Operator Agent can be designed to simulate inspection events based on a tentative optimal inspection plan. By interfacing with ATLAS, it further supports interaction with end-users (Fig. 4a). Preliminary testing was conducted in a video game scenario using the Unity 3D engine and the OpenAI GPT-3.5 API, enabling natural language-based prompting and plan enhancement through ATLAS (Fig. 4b).

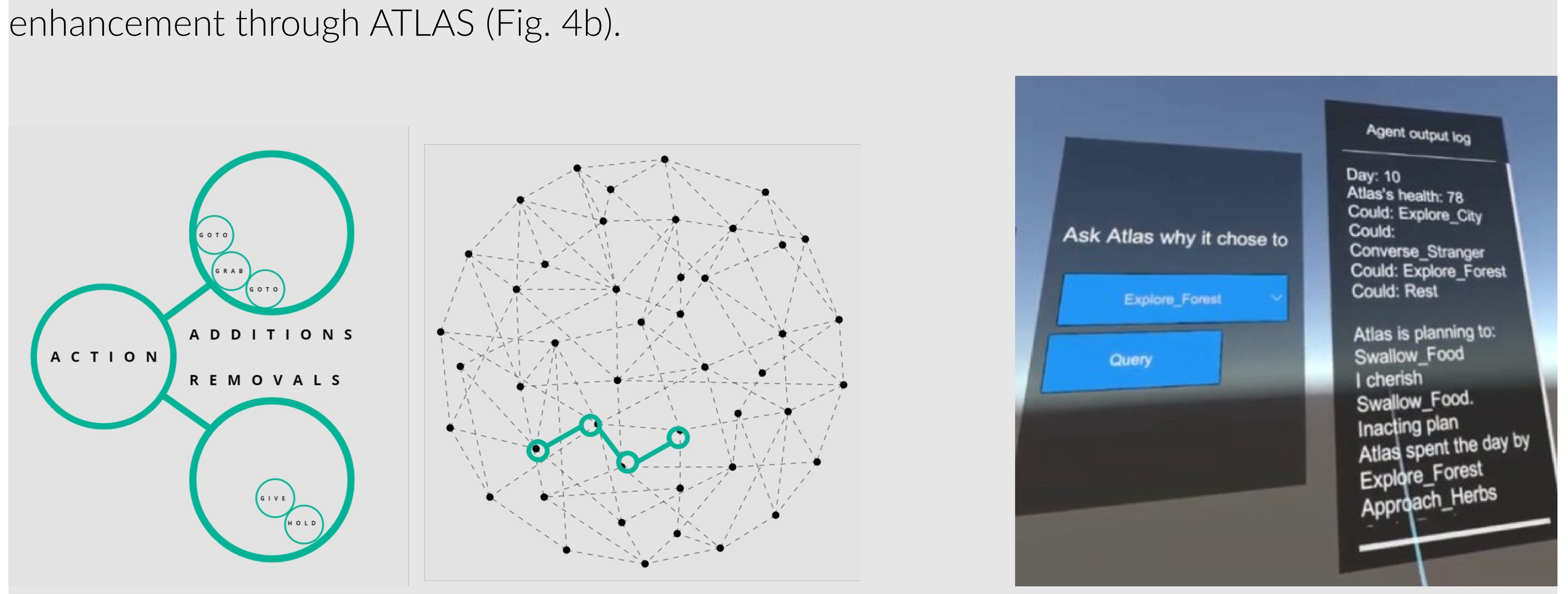

Figure 4. (a) Finite-state machine space for executing an optimal inspection plan; (b) ATLAS interfacing with end-users (simulated in a gaming context).

Digital Twinning Enabled Design Automation

Geotechnical and structural design is a well-established process performed by engineers in numerous design firms. Yet, automating this process has long been deemed impractical, largely due to the challenges of handling unstructured field reports and an overwhelming number of design specifications and rules. We argue that digital twinning—built on classical foundational technologies—can support design automation; however, deep uncertainties in managing field information and design logic must first be addressed. Initially, the digital twin may manifest only as a 3D geometric model (e.g., from architectural outputs), but it evolves as it integrates with the design process.

Leveraging recent advances in Generative AI, we propose an agentic workflow for digital twinning-enabled design automation, focusing on highway bridge structures. Highway bridge design is characterized by a limited selection of superstructure types, while substructure systems are highly dependent on site conditions. In practice, site reports can span hundreds or even thousands of pages (Fig. 5a).

Figure 5b illustrates our agentic workflow, where multiple Operator Agents collaborate with the universal **ATLAS** (Chief Planner) Agent.

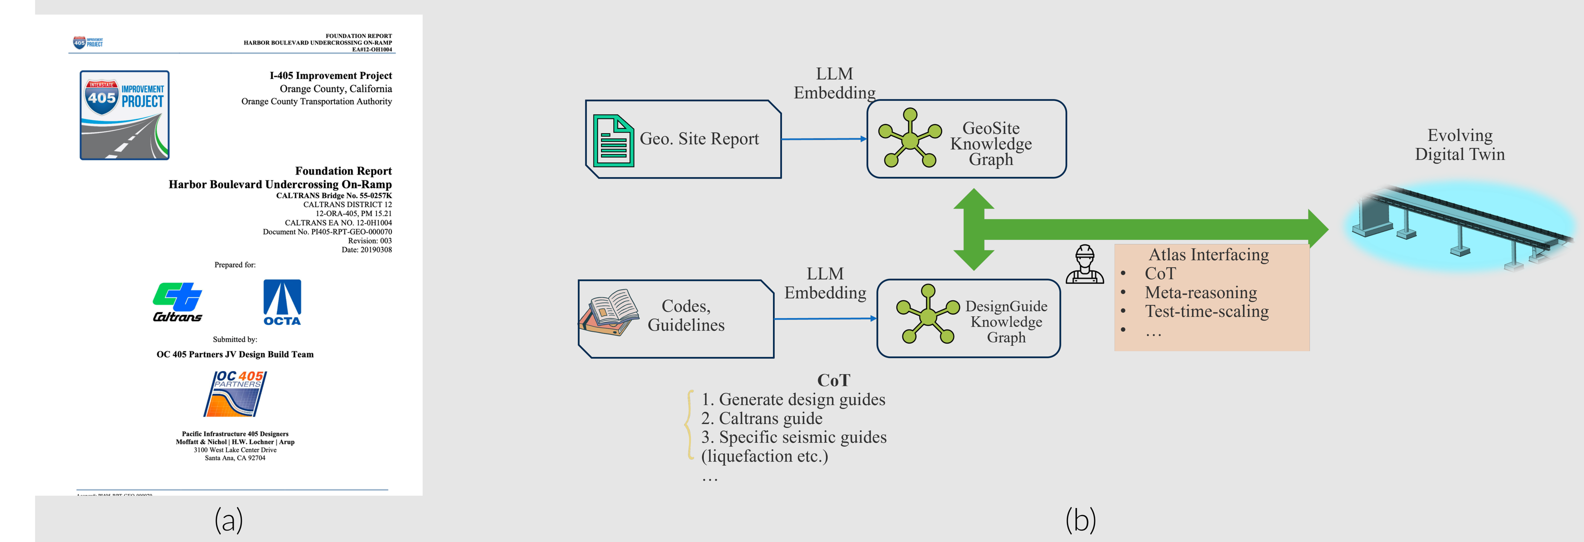

Figure 5. Conceptual schematics for digital twinning-enabled design automation for bridge foundations.

Design Automation via Test-time Scaling

We recognize that the design process is a highly structured "chain-of-thought" (CoT) process [6]. With modern reasoning-ready LLMs, engineering design automation can be realized through test-time scaling—that is, by leveraging inference compute. A few studies have been proposed for materials and drug design [5, 2]. In this work, we present preliminary results demonstrating the use of LLMs to generate engineering reports from knowledge graphs that support CoT generation of design processes. This approach will serve as the basis for further research toward developing a design automation-capable ATLAS system.

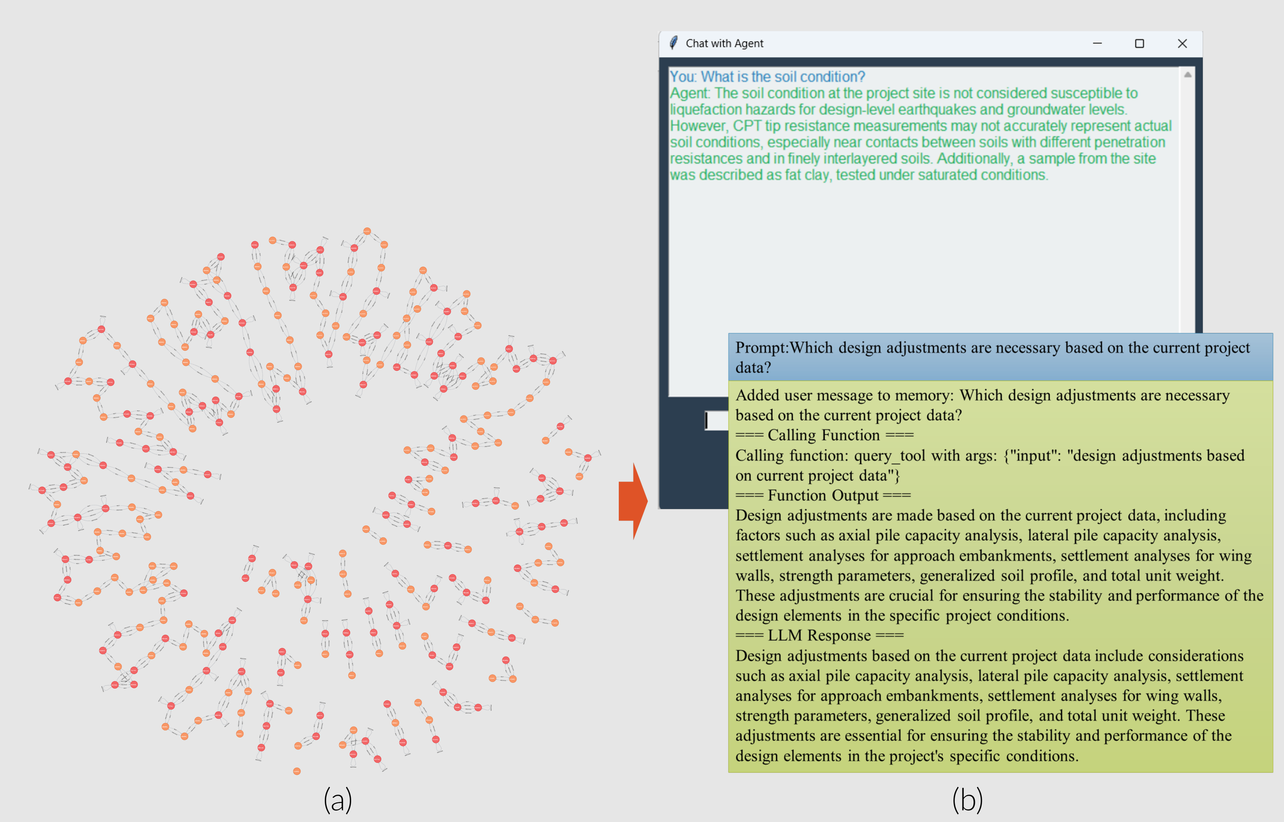

Figure 6. (a) Neo4j graph of the site report in Fig. 5a and (b) CoT prompting.

References

[1] Masataro Asai and Christian Muise. Learning neural-symbolic descriptive planning models via cube-space priors: The voyage home (to strips). *arXiv preprint arXiv:2004.12850*, 2020.

[2] Dela Quarme Gbadago, Sejin Go, and Sungwon Hwang. A leap forward in chemical process design: Introducing an automated framework for integrated ai and cfd simulations. *Computers & Chemical Engineering*, 192:108906, 2025.

[3] Michael J Olsen, Andre Barbosa, Patrick Burns, Alireza Kashani, Haizhong Wang, Marc Veletzos, Zhiqiang Chen, Gene Roe, and Kaz Tabrizi. *Assessing, coding, and marking of highway structures in emergency situations, Volume 1: Research overview*. Number Project 14-29. Transportation Research Board, National Academies, 2016.

[4] Claude E. Shannon. A mathematical theory of communication. *Bell System Technical Journal*, 27(3):379–423, 1948.

[5] Morgan Thomas, Albert Bou, and Gianni De Fabritiis. Test-time training scaling for chemical exploration in drug design. *arXiv preprint arXiv:2501.19153*, 2025.

[6] Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35:24824–24837, 2022.