Co-Adaptation in Human-Robot Collaborative Construction: A Case Study on Timber Joint Assembly

Michihiro Abe[1], Arastoo Khajehee[1] and Yasushi Ikeda[1]

[1]Department of Architecture, Graduate School of Engineering, The University of Tokyo
micabe4510@gmail.com,
arastoo@arch1.t.u-tokyo.ac.jp,
yasushi@arch1.t.u-tokyo.ac.jp

Abstract

In the construction industry, research on Human-Robot Collaboration (HRC) is advancing to leverage the complementary strengths of both humans and robots. While research on HRC in architecture is progressing, most studies focus on task partitioning. There remains a lack of research on co-manipulation, where humans and robots handle the same object, particularly in complex tasks requiring high adaptability, such as timber assembly involving friction and material tolerances. This paper proposes the Co-Adaptation framework for construction HRC, where both the human and the robot mutually learn and adapt during a collaborative task. We developed a system where a 6-DOF robot arm performs a timber joint insertion task with a human, utilizing reinforcement learning that incorporates an observation phase to learn from human force modulation. A preliminary exploratory experiment was conducted with three subjects to explore different adaptation scenarios: unilateral machine adaptation, unilateral human adaptation, and co-adaptation. We also introduced a "Cooperation Index" to evaluate the system's cooperativeness relative to human physical effort. The initial results suggest that a mutual learning system, particularly one where the robot observes and adapts to human haptic cues, has the potential to facilitate higher system adaptability and coordination than unilateral approaches. Furthermore, qualitative feedback from participants indicated the emergence of subjective adaptation strategies, such as role-based adaptation to the robot's movements. While these findings are preliminary and limited by the small sample size, they provide early evidence that explicitly accounting for mutual learning interactions can enhance human-centric robotic systems in construction.

Keywords –
Human-Robot Collaboration; Co-Adaptation; Timber Construction; Reinforcement Learning;

1 Introduction

1.1 Background

In the construction industries of advanced nations, labor shortages and the decline of skilled workers are critical issues, making the improvement of labor productivity an urgent challenge [1]. To address this, R&D on construction automation using robotics has been advancing [2]. However, unlike factory environments in the manufacturing industry, construction sites require bespoke production in unstructured environments [3]. Consequently, the scope of automation, particularly for on-site assembly tasks, has remained limited. Furthermore, it has been pointed out that automation which excludes humans reduces the remaining human tasks to simple labor, potentially leading to a loss of autonomy and intellectual satisfaction in work [4].

Against this backdrop, recent technological trends in construction automation emphasize a transition from rigid, efficiency-driven systems towards human-centric collaborative frameworks [5]. This approach attempts to improve productivity through Human-Robot Collaboration (HRC), leveraging both human perception and adaptability, and the robot's strength and precision.

However, most existing HRC frameworks in construction remain predominantly limited to rigid task partitioning or unilateral assistance, where the robot follows fixed trajectories. Such systems lack the reciprocal adaptability required for complex assembly tasks, such as timber joint insertion, where unpredictable material friction and geometric tolerances necessitate a continuous, bidirectional exchange of haptic information between agents.

1.2 State of the art

In the field of architecture, specifically, HRC systems that fuse human skills with robotic precision are being researched to address the heterogeneity of natural materials and component tolerances.

Mitterberger et al. presented a collaborative workflow for a reciprocal structure made of wooden rods. In their case study, the human handled the rope binding, which is a task difficult to automate, while the robot held the rods [6]. Rogeau and Rad proposed a workflow for assembling timber structures requiring adaptive motion due to tolerances. They used Mixed Reality to visually communicate robot movements and fitting positions to the human [7].

These examples defined robots and humans as agents with different skill sets and partitioned the tasks. However, in the transport and assembly of heavy architectural components, "Co-manipulation", where humans and robots handle the same object and interact through physical forces is essential. The implementation and verification of co-manipulation in the construction field remains insufficient, specifically regarding tasks that mimic the high level of adaptability required on actual construction sites.

In the field of Human-Computer Interaction (HCI), the concept of "Co-Learning" has been proposed for co-manipulation tasks. For instance, Shafti et al. demonstrated a framework where the reinforced learning based robot arm improves success rates through trials with a human in a cooperative ball-balancing task [8]. However, these studies often target simple games or low-degree-of-freedom tasks. Transitioning these concepts to the construction domain reveals significant physical barriers; for instance, timber joint assembly involves more complex variables than conventional laboratory tasks, such as non-linear friction, material shrinkage, and unpredictable geometric tolerances. While humans possess the innate haptic adaptability to navigate these uncertainties, there is a lack of research on co-manipulation systems where the robot can explicitly reference and learn from such human-driven strategies in real-time. This gap directly motivates our proposed Co-Adaptation approach, which seeks to enhance systemic performance by allowing the robot to observe and incorporate human haptic intelligence during complex assembly processes.

To establish a consistent terminology for this study, we clarify the relationships between the following terms based on the framework by Van Zoelen et al. [9]: “Co-adaptation” is dynamic and mutual adjustment of behavior between agents during a shared task. “Co-learning” is cumulative acquisition of skills and knowledge derived from iterative co-adaptation processes.

2 Research Objective

Based on the existing research and identified gaps, we hypothesize that "in construction tasks requiring collaboration and adaptability, Co-Adaptation system may yield higher systemic adaptability than unilateral learning by either the human or the machine alone". The objective of this exploratory study is to investigate this hypothesis and gain qualitative and quantitative insights into mutual learning systems for human-robot collaboration in construction tasks.

3 Methodology

To operationalize the research hypothesis, our research is structured into two primary stages: (1) Construct a framework for a human-machine mutual learning system. (2) Define a collaborative task requiring adaptability and the subsequent execution of multi-subject experiments.

First, the mutual learning framework is operationalized as a closed-loop haptic interaction system. The system takes real-time haptic interaction data (force/torque) as input, which is processed via a Multi-Armed Bandit algorithm (Robust Gaussian Thompson Sampling) to optimize robot wiggling parameters, and generates adaptive insertion motion as the system’s output.

Second, a specific collaborative task was defined to validate the framework's effectiveness. The task involves a human and a 6-DOF robot arm (UR10e) jointly holding a timber beam (approx. 1m) to perform a fitting operation at a central joint. This task requires adaptive motions due to friction caused by material tolerances, shrinkage, or expansion and simulates a typical assembly process in timber construction. As illustrated in Figure 1, the task is decomposed into four distinct phases to isolate the learning focus:

1. Transport: Joint gross motion to the target area.
2. Alignment: Human-led fine positioning.
3. Insertion: The core co-adaptation phase where the robot updates its policy based on human haptic cues and insertion progress.
4. Retraction: Completion of the cycle.

This decomposition is critical as it isolates the high-uncertainty insertion process, allowing the reinforcement learning system to focus specifically on haptic coordination while utilizing standard control for other phases.

Finally, the effectiveness of this mutual learning process is evaluated using the "Cooperation Index." This metric quantifies the ratio of task progression (insertion depth) to human physical effort (impulse), providing a human-centric evaluation of the collaborative system's performance.

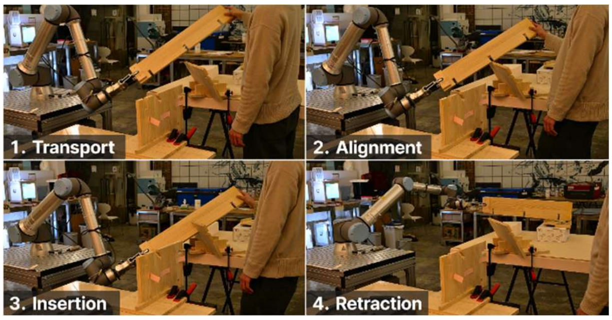

Figure 1. Phases of the human-robot collaborative timber joint assembly task.

4 System Implementation

This section describes the co-adaptation system. The setup utilizes a 6-DOF robot arm UR10e (Universal Robots)

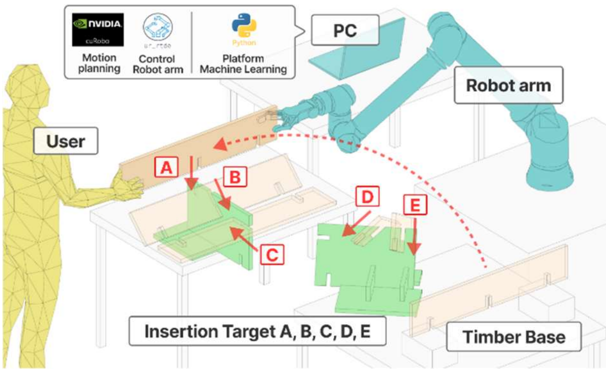

Figure 2. System architecture for co-adaptation experiment including UR10e and control platform.

4.1 Robot Arm Control

Since timber has natural shape tolerances, position control is not suitable for tasks involving contact. We adopted force control, maintaining external contact forces at desirable values while the end-effector follows a target position/orientation. To ensure the high-frequency real-time performance required for force control, communication between the robot and PC is performed via Real-Time Data Exchange (RTDE), which is a protocol designed for Universal Robots to facilitate the real-time communication between the robot controller and external systems, using the ur_rtde Python library [10].

4.2 Machine Learning Formulation

The robot employs Reinforcement Learning. Learning adaptive insertion motions from scratch in a real-world environment would require an unrealistic number of trials. Therefore, we preset a wiggling motion which applies and releases force relative to the insertion direction and robot optimized its three parameters through trials: "Minimum Force (F_{min})", "Force Amplitude (F_{amp})" and "Frequency (f)". The reward was defined as the insertion distance achieved within a fixed time (4 seconds). This dense reward formulation was adopted based on a preliminary pilot study (n=4), which revealed that binary "success/failure" metrics provided insufficient feedback for convergence within the trial limits imposed by participant fatigue. While initial tests used 5-second rounds, the duration was optimized to 4 seconds to increase the frequency of learning updates (up to 15 per 60-second trial), ensuring the number of updates significantly exceeds the number of discrete arms for effective policy refinement.

To optimize the wiggling parameters in real-time, the adaptive insertion task was formulated as a Multi-Armed Bandit (MAB) problem. As a MAB framework, the system is inherently stateless; however, haptic data from the Observation Phase is incorporated as a "Hint Bonus" (selection bias) to guide the search toward human-intuitive actions.

The action space comprises 64 discrete arms, defined by the Cartesian product of the following parameter sets:

$$F_{min} \in \{0,5,10,15\}\ \mathrm{N}$$
$$F_{amp} \in \{10,15,20,25\} N$$
$$f \in \{0.00,0.33,0.67,1.00\}\ \mathrm{Hz}$$

4.2.1 Robust Gaussian Thompson Sampling

Considering the inherent noise in timber assembly, such as material heterogeneity and human intervention, we employed Robust Gaussian Thompson Sampling, a probabilistic algorithm based on Bayesian inference to efficiently identify high-reward strategies under noisy feedback [11]. The reward r_t, defined as the insertion distance achieved per 4-second round, is modeled following a Gaussian distribution. To handle the non-stationarity of the task where the optimal strategy shifts as insertion progresses, an Exponential Moving Average update rule was implemented:

$$\mu_{i,t+1} = (1-\alpha)\mu_{i,t} + \alpha r_t$$

$$\sigma^2_{i,t+1} = (1-\alpha)\sigma^2_{i,t} + \alpha(r_t - \mu_{i,t})(r_t - \mu_{i,t+1})$$

where $\mu_{i,t}$ and $\sigma^2_{i,t}$ represent the estimated mean and variance of the reward, and α is the smoothing factor (learning rate), set to 0.15 to balance stability and adaptability. Additionally, a reward clipping mechanism ([-2.0, 10.0] mm) was applied to suppress the influence of outliers caused by physical collisions or sensor noise.

4.2.2 Observation Phase and Selection Bias

A core feature of the proposed Co-Adaptation is the Observation Phase. During this phase, the robot maintains a constant force while recording human-induced force modulations (Figure 3). The system performs a Fast Fourier Transform on the time-series force data to extract the human's minimum pushing force, force amplitude and frequency.

To achieve implicit coordination without corrupting the robot's objective statistics, a Selection Bias mechanism was introduced. The robot identifies the closest discrete arm i to the human's motion and assigns a temporary Hint Bonus b_i to the selection score S_i:

$$S_i = \theta_i - \beta\sigma_i + b_i$$

where θ is the sample drawn from the posterior distribution, and $\beta\sigma_i$ is a risk-averse penalty term (with $\beta = 0.3$) to suppress unstable candidates. This bias effectively narrows the search space toward human-intuitive strategies, accelerating system convergence. Based on the calculated scores, the robot executes the action associated with the arm i_t that yields the maximum selection score:

$$i_t = \mathrm{argmax}_{i \in \{1,\ldots,K\}} S_i$$

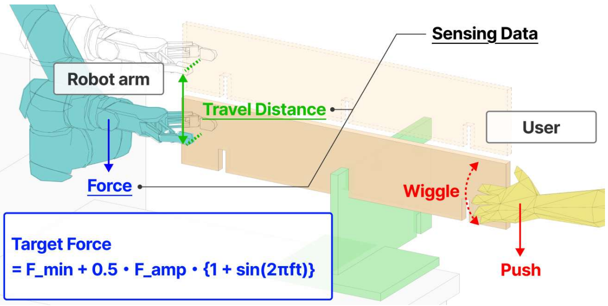

Figure 3. Definitions of force parameters and wiggling motion for robot adaptation.

5 Experiment

We conducted the insertion task with three subjects. As shown in the figure 2, five insertion targets (A–E) with different positions and angles were set up. The task involved moving from A to E and back to A, performing the insertion twice at each location (10 insertions total). All subjects provided informed consent. Safety measures included emergency stops upon contact detection and manual monitoring.

5.1 Experiment Overview

To verify the hypothesis that "the Co-Adaptation system yields higher systemic adaptability than unilateral learning," we prepared four scenarios (Table 1). Each subject performed the task under two control conditions to allow for comparison.

Table 1. Summary of the four experimental scenarios for comparative adaptation analysis.

Scenarios	Machine Learning	Human Learning	Observation
① Unilateral Machine Adaptation	○	×	-
② Unilateral Human Adaptation	×	○	-
③ Co-Adaptation with Observation	○	○	○
④ Co-Adaptation without Observation	○	○	×

1. Unilateral Machine Adaptation:
   A constant load is applied using weights instead of a human, so that the human side force is fixed. The robot learns adaptive behavior. This scenario is used only as a reference baseline to characterize machine performance under simplified conditions and is not treated as directly comparable to human-involved scenarios.
2. Unilateral Human Adaptation:
   The robot applies a constant insertion force without adaptive behaviours, while the human adapts their strategy
3. Co-Adaptation with Observation:
   Both human and robot adapt, and the robot incorporates the observation phase described in Section 4.2.
4. Co-Adaptation without Observation:
   Both human and robot adapt, but the robot does not observe and incorporate human force characteristics.

5.2 Result & Analysis

5.2.1 Pilot Study: Unilateral Machine Adaptation vs. Unilateral Human Adaptation

This experiment was conducted with one subject. Figure 4 illustrates the relationship between time and insertion depth. Since Unilateral Machine Adaptation alone involves suspending weights in place of humans, experiments are conducted only at Target A where the insertion direction is vertically downward.

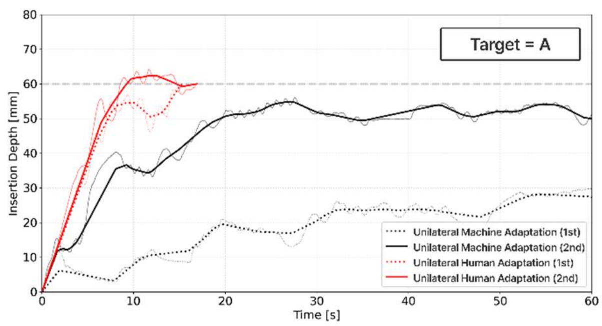

Figure 4. Comparison of insertion depth over time: Unilateral Machine (black) vs. Unilateral Human adaptation (red). Dotted lines represent the first trial; solid lines represent the second trial. 60mm is the insertion completion position.

In the Unilateral Machine Adaptation scenario, although the insertion depth increased in the second trial compared to the first, the task could not be completed in sixty seconds. In contrast, the Unilateral Human Adaptation scenario resulted in successful completion in both trials. These results indicate that human adaptability is crucial for achieving this task.

It should be noted that, as shown in Figure 4, the insertion depth in this study was measured based on the travel distance of the robot arm's gripper. However, if the timber component rotates relative to the insertion direction, a discrepancy may arise between the gripper's travel distance and the actual insertion depth. To accurately evaluate the progression of the insertion depth, oscillatory errors were corrected in the analysis using the extreme value interpolation method.

5.2.2 Proposal of the Cooperation Index:

Before proceeding to the experiment of co-adaptation scenarios, we discuss the metrics used to evaluate the effectiveness of the mutual learning system. If the system aimed solely at improving productivity, a metric such as insertion depth over time, as used previously, would suffice. However, since this study addresses a human-centric HRC system, it is essential to introduce a metric focused on human contribution.

In a post-experiment interview, one subject stated, "I adapted my movements so that the insertion would proceed even with less applied force". Based on this insight, the ratio of the achieved insertion result relative to the amount of force applied by the human can be considered a valid indicator for capturing collaborative system.

In the current task, the force applied by the human was not directly measured; therefore, estimating the human's applied force from the robot arm's force sensor reading is needed. Figure 5 illustrates the relationship between the target force and the measured values in 5.2.1 Pilot Study (① Unilateral Machine Adaptation vs. ② Unilateral Human Adaptation). The red line represents the value measured by the force sensor, while the yellow dotted line indicates the control target value.

In the ① Unilateral Machine Adaptation scenario, where the force on the human side is constant due to the counterweight, a correlation between the yellow (target) and red (measured) graphs is observed. In contrast, in the ② Unilateral Human Adaptation scenario where the human modulates force, a deviation from the target force occurs due to human intervention. Furthermore, considering the control mechanism and the fact that the timber acts like a seesaw between the human and the machine, it can be inferred that if the robot arm detects a force exceeding the target value, it is mainly due to the influence of the force applied by the human.

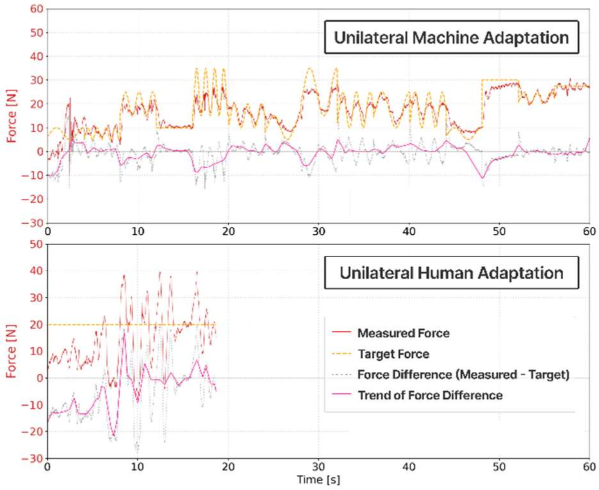

Figure 5. Measured force vs. target force profiles in Unilateral Machine Adaptation and Unilateral Human Adaptation scenarios.

Based on this reasoning, we propose that the cooperation between the human and robot can be measured by defining the Cooperation Index E as follows:

$$E := \frac{D}{\int_{\{t|F(t)>0\}} F(t)dt}$$

In this formula, D represents the final insertion depth. The denominator represents the impulse calculated only when the measured force F(t) exceeds the target force, identifying the specific moments of active human intervention. It is intended as a comparative performance metric. Specifically, E functions as a "normalized efficiency score" that quantifies the system's gain, the extent of task progression achieved relative to the human's physical impulse investment. This allows for an evaluation of whether the robot is effectively reducing the physical burden on humans while maintaining task performance.

5.2.3 Unilateral Human Adaptation vs. Co-Adaptation with Observation

One subject performed both scenarios. Table 2 shows the calculated results of the Cooperation Index. Since the insertion task was performed twice at each target location, the improvement in cooperation was evaluated by calculating the ratio between the average values of the first and second trials.

The results showed a ratio of 0.92 for ②Unilateral Human Adaptation, whereas ③ Co-Adaptation with Observation achieved a higher ratio of 1.81. This result suggests the adaptive effect of mutual learning between the human and the machine in this specific task.

Table 2. Comparison of Cooperation Index (E): Unilateral Human Adaptation vs. Co-Adaptation with Observation scenarios

scenario	Target	A	B	C	D	E	Ave	Ratio
Unilateral Human Adaptation	E_1st	2.98	23.85	0.62	8.13	0.22	7.16	-
	E_2nd	4.4	21.03	0.58	6.74	0.33	6.62	0.92
Co-Adaptation with Observation	E_1st	1.14	7.58	0.65	4.7	0.31	2.88	-
	E_2nd	16.95	5.21	0.57	3.00	0.26	5.20	1.81

5.2.4 Co-Adaptation with Observation vs. Co-Adaptation without Observation

Next, to verify the hypothesis that mutual learning interactions are critical within the co-adaptation system, a comparative experiment between ③ Co-Adaptation with Observation and ④ Co-Adaptation without Observation was conducted with two subjects. To mitigate potential order effects, the condition order was counterbalanced across participants: one participant performed the with-observation condition first, whereas the other participant performed the without-observation condition first. Table 3 presents the calculated Cooperation Index results.

Table 3. Comparison of Cooperation Index (E) for two subjects: Co-Adaptation with Observation vs. Co-Adaptation without Observation scenarios

Subject	scenario	Target	A	B	C	D	E	Ave	Ratio
X	Co-Adaptation without Observation	E_1st	5.97	8.43	0.41	0.91	0.21	3.19	-
		E_2nd	38.1	4	0.45	38.4	0.19	16.2	5.09
	Co-Adaptation with Observation	E_1st	0.8	2.39	0.34	3.55	0.21	1.46	-
		E_2nd	21.9	2.52	0.3	23.4	0.26	9.66	6.66
Y	Co-Adaptation without Observation	E_1st	1.76	25.09	0.32	0.66	0.42	5.65	-
		E_2nd	3.09	6.71	0.3	1.14	0.38	2.32	0.41
	Co-Adaptation with Observation	E_1st	3.68	27.28	0.21	4.79	0.39	7.27	-
		E_2nd	11.2	7.87	0.43	0.48	0.43	4.07	0.56

For Subject X, comparing the first and second halves of the trials revealed a greater increase in the index value under the with Observation condition. Similarly, for Subject Y, the system with Observation showed a tendency to suppress the decline in the evaluation metric. These results suggest the importance of designing systems that explicitly account for mutual learning interactions

5.2.5 Qualitative Feedback and Subjective Adaptation

Post-experiment interviews provided further insight into the diverse strategies human subjects employed during co-adaptation.

Subject X reported a subjective improvement in task fluidity during the latter half of ③Co-Adaptation with Observation trials. The subject noted, "I felt the insertion progressing even with reduced force in the later stages, which allowed me to relax my physical effort. My overall perception was that the coordination improved over time, and the insertion speed appeared to increase." In contrast, Subject X perceived no significant temporal change in ④ without observation scenario. This feedback aligns with the measured increase in the Cooperation Index for Subject X, suggesting that robots' observation of human cues may have fostered a more intuitive collaborative environment.

Subject Y, however, did not perceive a distinct difference between the observation-based and non-observation-based scenarios. Instead, Subject Y focused on a role-based adaptation strategy regardless of the machine's learning mode. The subject stated, "Since the robot arm was providing force in the insertion direction, I adapted my role to prioritize alignment and orientation over pushing. I assumed that if the position and angle were correct, the robot would handle the rest." This indicates that while the robot learns from the human, the human also autonomously redefines their role to optimize the partnership, a key characteristic of co-evolution in HRC.

6 Conclusion & Contributions

This exploratory study proposed the Co-Adaptation framework for human-robot collaborative tasks in timber construction, specifically addressing joint assembly where friction and tolerances require adaptive force modulation.

6.1 Primary Contributions

The primary contributions of this paper are summarized as follows:

- Proposal of the Co-Adaptation Framework: We developed a physical HRC framework where both the human and the robot adapt to each other in real-time, moving beyond simple task partitioning.
- Implementation of "Observation-based" Learning: We demonstrated a mechanism where the robot observes human haptic cues to accelerate its own reinforcement learning, facilitating "implicit coordination"
- Human-Centric Evaluation Metric: We proposed

the Cooperation Index to evaluate HRC systems by the efficiency of the collaborative output relative to human physical effort

- Preliminary Findings: Our preliminary results suggest that mutual learning, supported by observation capabilities, can enhance systemic adaptability and foster diverse human adaptation strategies, such as role-redefinition

6.2 Limitations and Future Work

While these initial findings are promising, they should be interpreted as preliminary evidence within the scope of this exploratory case study. This research is primarily limited by its small sample size (n=3) and the simplified nature of the assembly task. However, in the nascent field of physical co-manipulation for construction, this scale facilitated a detailed, qualitative analysis of individual "cognitive and physical adaptation" strategies that might be obscured in larger aggregate datasets. For instance, the proactive role-redefinition observed in participants provides critical insight into how humans adjust their agency in response to robotic learning.

Nevertheless, factors such as potential order effects may not yet be generalizable to all construction scenarios. Future work is required to validate these findings with a larger participant pool to statistically confirm the observed trends. Furthermore, we aim to investigate how long-term collaboration and increased task complexity influence the stability and efficiency of the Co-Adaptation process in unstructured construction environments.

References

[1] FIEC. Shortage of Labour in the Construction Industry: 2025 Edition. European Construction Industry Federation. On-line: https://www.fiec.eu/application/files/1517/6331/6876/202511_FIEC_Position_Paper_labour_shortages.pdf, Accessed: 12/01/2026.

[2] T. Bock. The future of construction automation: Technological disruption and the upcoming ubiquity of robotics. Automation in Construction, 59:113–121, 2015.

[3] L. Koskela. Application of the new production philosophy to construction, volume 72(39). Stanford University, Stanford, 1992.

[4] Braverman, Harry. Labor and monopoly capital: The degradation of work in the twentieth century. nyu Press, 1998.

[5] A. G. Burden, G. A. Caldwell and M. R. Guertler. Towards human–robot collaboration in construction: current cobot trends and forecasts. Construction Robotics, 6(3):209-220, 2022.

[6] D. Mitterberger, L. Atanasova, K. Dörfler, F. Gramazio, M. Kohler. Tie a knot: human–robot cooperative workflow for assembling wooden structures using rope joints. Construction Robotics, 6(3):277–292, 2022.

[7] N. Rogeau and A. R. Rad. Collaborative timber joint assembly: Augmented reality for multi-level human-robot interaction. International Journal of Architectural Computing, 22(4):642–660, 2024.

[8] A. Shafti, J. Tjomsland, W. Dudley, and A. A. Faisal. Real-world human-robot collaborative reinforcement learning. In Proceedings of the 2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pages 11161–11166, Las Vegas, USA, 2020.

[9] E. M. Van Zoelen, K. Van Den Bosch, and M. A. Neerincx. Becoming Team Members: Identifying Interaction Patterns of Mutual Adaptation for Human-Robot Co-Learning. Frontiers in Robotics and AI, 8:692811, 2021.

[10] A. P. Lindvig, I. Iturrate, U. Kindler, and C. Sloth. ur_rtde: An Interface for Controlling Universal Robots (UR) using the Real-Time Data Exchange (RTDE). In Proceedings of the 2025 IEEE/SICE International Symposium on System Integration (SII), pages 1118–1123, Munich, Germany, 2025.

[11] S. Agrawal and N. Goyal. Analysis of Thompson sampling for the multi-armed bandit problem. In Proceedings of the 25th Annual Conference on Learning Theory, pages 39.1–39.26, Edinburgh, Scotland, 2012.

Acknowledgement: Google Gemini 3 was used to edit and proof read this paper.