Dual-Robot Anchor-Bracket Installation: Coordination Challenges and Requirements for Learning-Based Adaptation

Niki Kentroti[1] **and Kepa Iturralde**[1]

[1]Chair of Digital Transformation in Construction, Institute of Construction Management, Faculty of Civil and Environmental Engineering, University of Stuttgart, 70569 Stuttgart, Germany

niki.kentroti@ibl.uni-stuttgart.de, kepa.iturralde@ibl.uni-stuttgart.de

Abstract -

Multi-robot construction tasks remain difficult to automate using pre-programmed control due to environmental variability, tight tolerances, and the need for continuous coordination during execution. This paper presents preliminary work toward adaptive dual-robot façade assembly, using anchor-bracket installation as a representative case. Through simulation-based studies, we validate the feasibility of pre-determined dual-robot installation while identifying key limitations, including sensitivity to variation, execution-level coordination failures, and limited recovery from deviations. These findings motivate the need for adaptive coordination mechanisms beyond fixed control logic. Building on the validated installation sequence, we outline and partially formalize how the task structure can be extended toward learning-based execution. While learning policies are not implemented in this work, the paper establishes a simulation-grounded problem structure and research direction for future development of adaptive multi-robot construction systems.

Keywords -

Construction robotics; Façade installation; Dual-robot coordination; Learning-based execution; Anchor-bracket assembly

1 Introduction

Robotic construction tasks on site are characterized by high environmental variability, sequential dependencies between actions, and the need for coordination between multiple agents. Unlike controlled industrial settings, construction workflows unfold incrementally, with each action altering the physical state of the environment. Task feasibility, accuracy requirements, and coordination constraints continuously change during execution, creating conditions that fundamentally differ from traditional manufacturing automation. Simulating such conditions is challenging due to necessary abstractions and approximations of real-world physical phenomena [1].

Pre-programmed robotic systems have become well-established in the architecture, engineering, and construction (AEC) industry over recent decades. These systems have been successfully deployed across diverse applications including bricklaying [2], scaffold-free assembly of building elements [3], prefabrication assembly [4, 5] and façade installation [6], demonstrating high accuracy and reliability in executing predetermined task sequences. Such approaches excel when operating conditions remain consistent with their programmed parameters. However, construction environments often present variable conditions and unforeseen deviations that require frequent reprogramming or manual intervention. These inherent complexities have resulted in limited adoption of automation in the construction industry [7], where robotic systems struggle to match the adaptability and problem-solving capabilities required for on-site operations.

These challenges are particularly evident in façade assembly, where the installation of mounting brackets 1, anchor points, and connection elements forms the critical infrastructure for curtain wall systems [6]. Their installation requires high positional accuracy and involves complex sequential dependencies, as each anchor bracket must be precisely placed to create a mounting grid for the subsequent installation of the façade panels. In practice, the installation of anchor-brackets is not strictly a linear process but requires continuous adjustment to substrate irregularities, tolerance accumulation, and alignment constraints.

From a robotic perspective, such tasks expose the limitations of single-agent, pre-defined execution models. Sequential dependencies arise not only from task logic but also from interactions between collaborating agents. The actions of one robot can directly alter the environment or partially complete a task, thereby influencing the feasible action space of other agents. Managing these interdependencies through conventional planning and control methods remains challenging. Reinforcement learning (RL) has been explored as a means of enabling robots to adapt their motion planning and decision-making through interaction with dynamic environments [8]. Learning-based approaches offer the potential to cope with uncertainty and variability that are difficult to model explicitly.

This paper examines the problem space of sequential,

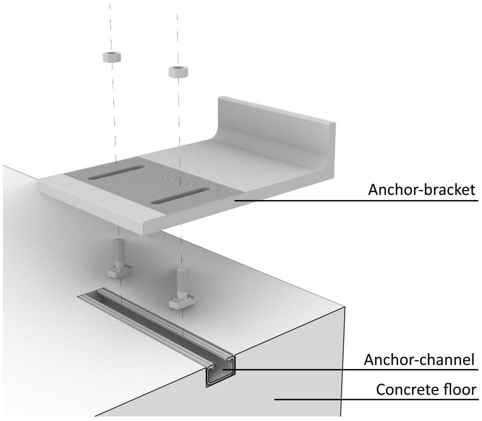

Figure 1. Industry-standard anchorage system.

dual-robot coordination in façade anchor-bracket installation through simulation-grounded task design and execution analysis, and reformulates the task as an execution-level adaptation problem.

2 State of the art

Existing research in robotic façade assembly, RL in construction applications and sequential task learning with RL has primarily addressed these challenges in isolation. This section reviews the state of the art in each domain.

2.1 Robotic façade assembly

Early work on robotic façade construction explored mechanized assistance for curtain wall installation. Yu et al. [9] developed a curtain wall installation robot combining a commercial excavator with a custom manipulator to handle and position unitized façade elements on site. More recent work has investigated robotic systems for curtain wall module installation using large-scale positioning mechanisms. Iturralde et al. [6] presented a cable-driven parallel robot platform carrying a robotic arm and task-specific tools, enabling semi-automated anchor bracket installation and curtain wall module placement. The system was demonstrated through full-scale prototyping and testing in an experimental setup corresponding to a building structure. Recent work has also addressed robotic façade assembly from a design-stage optimization perspective. Ali et al. [10] proposed a robot-based façade spatial assembly optimization framework that automatically determines robot base locations and pick-and-place positions using genetic algorithms. Heuer et al. [11] presented a crane-guided, semi-automated platform for façade panel assembly, relying on a pre-installed substructure and combining crane-based coarse positioning with robotic fine alignment. Both [10] and [11] focus on façade panel placement and rely on a pre-installed substructure, leaving earlier installation steps—such as bracket placement—outside the scope of robotic execution. Commercial robotic systems addressing façade installation operations have also emerged in industry, such as the solution proposed by Raise Robotics [12].

2.2 Reinforcement learning in construction applications

Applications of RL in construction robotics have demonstrated successful learning for timber assembly [13], block placement [14], and ceiling tile installation [15], though typically focusing on isolated tasks rather than complete construction sequences [16]. Apolinarska et al. [13] demonstrated one of the first successful applications of RL for assembly tasks using industrial robots in architectural construction. They developed a sim-to-real transfer approach for assembling timber lap joints, training an Ape-X deep deterministic policy gradient (DDPG) policy entirely in simulation based on force/torque and pose observations. Belousov et al. [17] demonstrated RL–based control for the contact-rich placement of discrete building elements in architectural structures.

Recent advances in multi-agent Reinforcement Learning (MARL) have begun to address the coordination challenges in construction robotics. Duan and Zou [18] developed a multi-agent Proximal Policy Optimization (PPO) framework combining MARL-based trajectory planning with inverse kinematics, demonstrating successful multi-robot coordination in pick-and-place and beam installation tasks. Sartoretti et al. [14] demonstrated centralized training with decentralized execution for multi-robot block assembly. Limitations emerged, however, when agents faced strategic construction sequences or complete scaffolding cleanup, highlighting the challenges of applying distributed RL to more complex assembly tasks.

In the domain of RL for sequential tasks, several approaches have demonstrated success in assembly problems. Huang et al. [16] proposed augmenting RL training with virtual reality demonstrations to help robots learn complete pick-and-install sequences for window panels , showing that expert demonstrations can overcome exploration challenges in sparse-reward, long-horizon tasks. Their work remained limited to single-robot scenarios with fixed task sequences. Yang et al. [19] proposed a hierarchical RL framework for single-robot sequential manipulation tasks, while Sun et al. [20] addressed subtask ordering through curriculum-guided exploration, both demonstrating improved stability in multi-step manipulation behaviors. For assembly sequence planning specifically, Shu et al. [21] proposed an RL approach that uses Graph Neural Networks to plan feasible assembly sequences in multi-part assemblies. However, all of these methods focus on single-robot scenarios. While they effectively handle task decomposition, sequential dependencies, and even complex assembly planning, they lack the multi-agent coordination capabilities essential for construction automation where multiple robots must synchronize their actions in dynamic, unstructured environments.

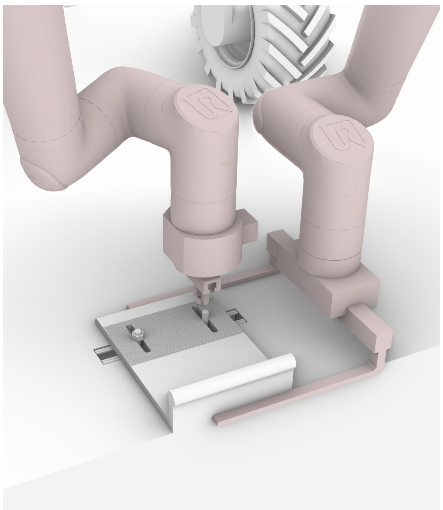

Figure 2. Robots operating in close proximity with minimal clearance

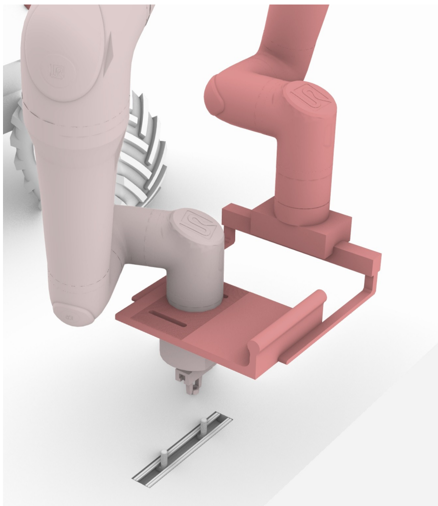

Figure 3. Collision caused by timing or positioning deviations between the two manipulators.

3 Problem Definition

3.1 Task description: Anchor-bracket installation

Anchor-bracket installation constitutes a foundational step in curtain wall construction, establishing the structural interface between the building substrate and façade elements. Each anchor must be positioned accurately relative to the building geometry and neighboring anchors, forming a mounting grid that constrains all subsequent installation steps. Small deviations in early placements propagate through the system, affecting alignment, load distribution, and the feasibility of later installation stages.

3.2 Sequential dependencies and coordination

Sequential dependencies arise from both task logic and physical constraints inherent to anchor-bracket installation. While a nominal execution sequence can be engineered through offline planning and simulation, its validity depends on a set of tightly coupled assumptions regarding reachability, collision-free motion, alignment accuracy, and tool access. Certain actions require partial completion of preceding steps, such as stabilizing an anchor bracket before fastening or maintaining clearance for tool engagement during tightening. Although the overall order of operations can be predefined, local adjustments in timing, positioning, and execution are often required in response to accumulated tolerances, minor deviations in placement, or interaction effects between the two manipulators. Importantly, the execution of one robot's actions directly affects the feasibility of subsequent actions by the other robot, even when the nominal sequence remains unchanged. Examples of these dependencies are illustrated in Figures 2 and 3.

4 System grounding through simulation and task sequence design

4.1 Simulation-Based Task Design

The robotic setup comprised two Universal Robots UR30 manipulators mounted on a Robotnik RB-FIQUS mobile platform. The RB-FIQUS is a large-footprint industrial mobile base, whose extended frame enables a dual-arm configuration with a center-to-center inter-base distance of 905 mm. Each UR30 has a maximum reach of 1300 mm; with bases positioned 905 mm apart, the two workspaces overlap substantially over the installation zone, while the mobile platform body restricts downward arm configurations and defines the lower boundary of the operational envelope. This configuration corresponds to a robotic installation system under active development within the AMALTEA project, targeting on-site façade installation in real building construction workflows. Simulations were developed in Rhinoceros 3D and Grasshopper using the Robots plugin [22] to evaluate process feasibility and define candidate task sequences relative to the CAD layout.

These studies revealed that the installation steps are tightly interdependent, with each action directly constraining subsequent ones. Alternative handling strategies for anchor brackets and threaded fasteners were therefore evaluated, including variations in grasp points, gripper jaw geometries, and part presentation methods. To ensure reliable handling under these constraints, batch-loaded magazines integrated into the mobile platform were introduced. These magazines consist of fixed slot arrays that hold brackets and fasteners in predefined orientations, enabling consistent and repeatable pick actions without active feeding or singulation mechanisms. Safe approach and retreat trajectories were defined based on the simulation results to avoid inter-arm and self-collisions, and appropriate standoff distances were established to ensure controlled tool engagement.

A custom collision-detection Python script was developed within the Grasshopper environment to systematically identify interferences across planned motion sequences. Inter-arm proximity was additionally monitored throughout the sequence using a custom Python script, which computed the minimum joint-center distance between the two manipulators at each frame. The analysis shows that minimum inter-arm proximity occurs during the coordinated fastening phase, reaching 97.49 mm at the end of the recorded sequence. Table 1 reports sampled measurements across the sequence, with Joint R_1 and Joint R_2 indicating the indices of the closest joint pair at each recorded moment. Repeated t_{R2} values indicate Robot 2 holding position while Robot 1 continues its trajectory.

Table 1. Minimum inter-arm joint-center distance across the dual-robot installation sequence

t_{R1} (s)	t_{R2} (s)	Dist. (mm)	Joint R_1	Joint R_2
6.16	14.06	1000.54	3	1
9.93	17.83	763.39	7	1
13.81	21.71	680.77	7	1
17.64	25.54	598.37	7	4
21.58	29.47	435.43	7	5
23.87	29.86	284.63	7	5
27.32	29.86	143.91	7	7
28.20	29.86	**97.49**	7	7

Collectively, these analyses demonstrated that even minor variations in task ordering or approach strategy could invalidate subsequent actions, reinforcing the need to treat the process as a tightly coupled sequential problem. While these studies establish the feasibility of the nominal sequence, the simulation environment is limited to kinematic analysis and does not capture force-torque dynamics, contact mechanics, material dimensions and weight, or sensing requirements such as force/torque feedback, visual perception, and proximity sensing. Addressing these limitations and the associated reality gap [1] is planned as the next step, beginning with transition to a physics-based simulation environment.

4.2 Operational Task Sequence

Based on the simulation outcomes, a synchronized dual-robot operational sequence was defined. Robot 1 is responsible for anchor-bracket placement and stabilization, while Robot 2 handles fasteners and tool-based operations. From verified standoff poses, both manipulators execute predefined approach trajectories within the validated collision envelope. Robot 2 retrieves bolts from an on-platform magazine, seats them into the anchor channel, and rotates them into the locking orientation using wrist-axis rotation 4. Robot 1 then retrieves an anchor bracket, aligns it with the target installation pose, and maintains a stabilizing grasp throughout the fastening process. Fasteners are pre-threaded using wrist rotation, after which Robot 2 exchanges to a high-torque impact wrench via a wrist-mounted tool changer system and tightens each joint to the specified torque. Upon completion, robot 1 releases the anchor, and both robots return to standoff positions before proceeding to the next installation location 5.

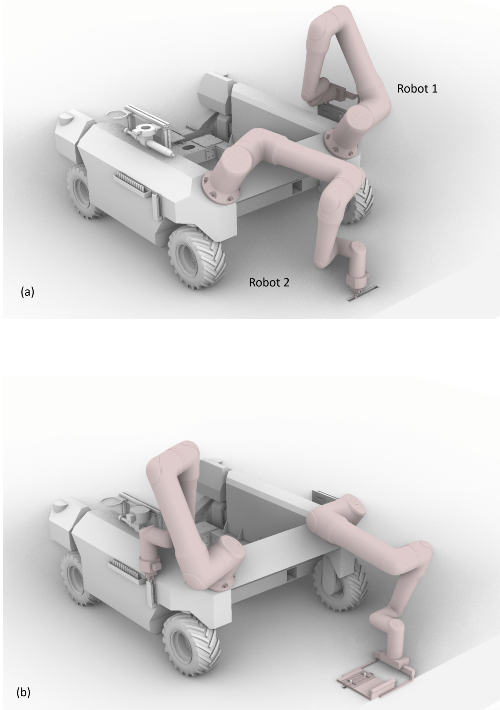

Figure 4. Operational task sequence for dual-robot anchor installation (a-b).

5 Toward Learning-Based Execution

5.1 From Validated Sequence to Execution-Level Learning

The simulation studies presented in Section 4 confirm the kinematic feasibility of the dual-robot anchor instal-

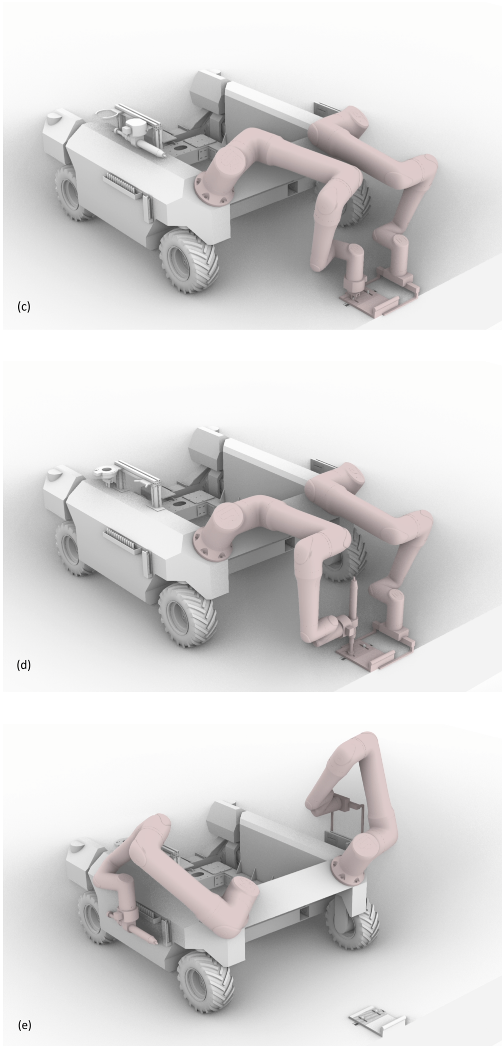

Figure 5. Operational task sequence for dual-robot anchor installation (c-e).

lation sequence under nominal conditions. However, execution is highly sensitive to small deviations in timing and positioning between the two manipulators. Minor discrepancies, such as millimeter-scale misalignments and slight delays in synchronization, were observed to propagate across phases and result in collision risks or failure to maintain the required coordination constraints.

- **Limits of Pre-Programmed Coordination**

  The execution challenges identified above are closely linked to the degree of coupling between the two robotic agents. During multiple phases of the installation process, including bracket stabilization and torque-controlled fastening, the manipulators operate in close proximity. In these configurations, the motion applied by one robot directly influences the feasible action space of the other. Consequently, execution success depends not only on individual robot accuracy, but on precise coordination in time and space.

  A common response to such challenges is the introduction of rule-based coordination mechanisms, including explicit synchronization points, conditional branching, and hand-coded recovery behaviors, corresponding to what the multi-robot systems literature describes as static or strongly coordinated approaches based on predefined rules and explicit coordination protocols [23]. However, as interaction complexity increases, such approaches scale poorly and rely on discrete state transitions that are ill-suited to capturing gradual shifts in dynamics, such as the transition from compliant stabilization to torque-controlled fastening, and become brittle when variability in component placement, accumulated tolerances, and contact conditions is unavoidable.

  This is further compounded by indirect mechanical coupling through the shared workpiece, a challenge anticipated from the task structure but not directly observable in kinematic simulation. During fastening, torque applied by one manipulator is expected to induce positional disturbances in the other's grasp if compliance is not continuously regulated. These conditions collectively motivate learning-based approaches, as the combination of tight spatial tolerances, unavoidable execution variability, and indirect mechanical coupling through the shared workpiece exceeds what can be reliably captured through explicit specification.

- **Reframing the Problem: Adaptive Execution Within a Fixed Task Structure**

  Rather than treating the task as one of discovering or optimizing task sequences, the core challenge lies in enabling robust execution of a known and validated sequence under conditions of tight coupling and uncertainty. From this perspective, the installation process is best understood as an adaptive execution problem.

  In the proposed formulation, the sequence of task phases remains fixed and externally defined by construction logic and safety considerations. Learning is not introduced to alter task ordering or decision logic, but to support execution-level adaptation within each phase 6. The specific execution-level adaptation required is formalized in Section 5.2. By constraining adaptation to the execution layer, the system preserves interpretability, safety, and alignment with construction requirements, while gaining flexibility where explicit specification is most difficult. This reframing

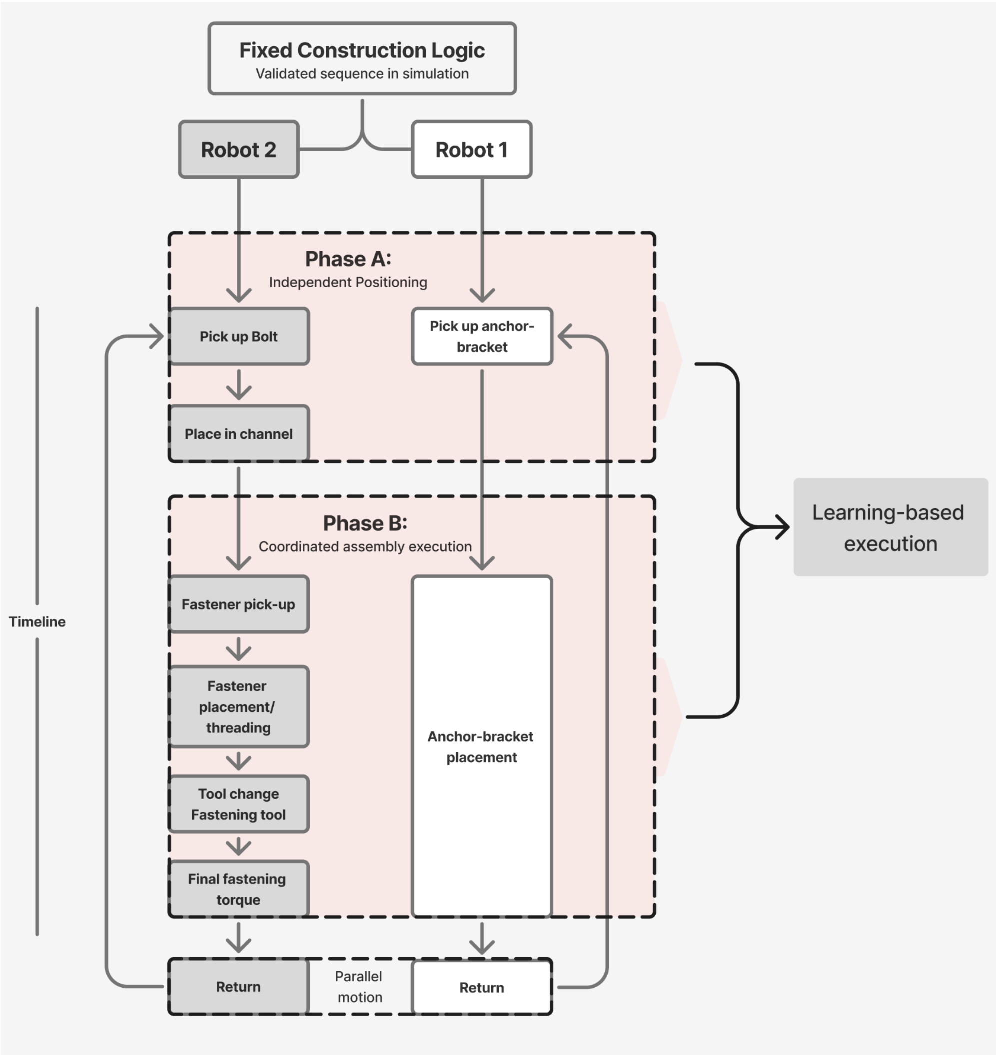

Figure 6. Dual-robot anchor installation sequence structured into Phase A(positioning) Phase B(coordinated execution).

distinguishes the proposed approach from task and motion planning frameworks that focus on symbolic decision-making or high-level sequencing [24]. Instead, the emphasis is placed on maintaining feasibility of a predefined sequence in the presence of execution variability, a problem that is particularly pronounced in multi-robot construction scenarios [25].

5.2 Toward Learning-Based Execution Policies

The adaptive execution problem for Phase B can be partially formalized as a cooperative MARL problem. The core challenge is not learning full motion trajectories — nominal trajectories are pre-programmed and executed by the built-in robot controllers — but rather learning when and how to adjust execution in response to deviations that arise during contact-rich operations such as bolt insertion, bracket alignment, and coordinated fastening.

Each robot $i \in \{1, 2\}$ is modeled as an independent agent observing a local state and producing execution adjustments. The tight coupling between agents arises because the valid action space of one robot depends directly on the current execution state of the other: Robot 1 cannot safely proceed to anchor placement while Robot 2 is still correcting bolt positioning. This dependency cannot be resolved through fixed timing rules, since the duration of each correction varies with the magnitude of the deviation encountered.

The state observed by each agent is defined as $s_i^t = (\mathbf{f}_i^t, \mathbf{e}_i^t, \mathbf{f}_j^t, \mathbf{e}_j^t, \phi^t)$, where $\mathbf{f}_i^t \in \mathbb{R}^6$ is the end-effector force/torque measurement of agent i, $\mathbf{e}_i^t \in \mathbb{R}^6$ is the pose error relative to the nominal target, $\mathbf{f}_j^t$ and $\mathbf{e}_j^t$ are the corresponding observations of the other agent j, and $\phi^t \in \Phi$ is a discrete phase indicator encoding the current sub-step within Phase B. Each agent produces a continuous action $a_i^t \in \mathbb{R}^6$ representing Cartesian end-effector velocity corrections applied on top of the nominal trajectory, within bounds imposed by the pre-validated collision envelope.

$$r^t = r^t_{\text{task}} - \lambda_c\, r^t_{\text{collision}} - \lambda_d\, r^t_{\text{deviation}} \tag{1}$$

where r^t_{task} rewards progress toward phase completion within tolerance, $r^t_{\text{collision}}$ penalizes inter-arm proximity violations, and $r^t_{\text{deviation}}$ penalizes proceeding to the next phase before the current operation is within the required tolerance bounds. The coefficients λ_c and λ_d balance the relative contribution of each penalty term and will be tuned during training. The phase indicator ϕ^t gates which reward components are active across sub-steps. This formulation does not constitute a complete implementation specification, but establishes the structural elements necessary to scope the learning problem and connect the validated task sequence to future implementation.

5.3 Conclusion

This paper analyzed dual-robot anchor-bracket installation as a tightly coupled, sequential construction task and investigated its execution through simulation-grounded task design. The studies demonstrate that a predefined installation sequence can be validated in terms of reachability, collision avoidance, and tool access, confirming the feasibility of coordinated dual-robot execution under nominal conditions. At the same time, the simulations show that successful execution is highly sensitive to small deviations in timing and positioning, even when the overall task sequence remains fixed. The task structure further anticipates sensitivity to force interactions during fastening, a factor not directly observable in kinematic simulation but expected to compound execution challenges in physical deployment. Together, these observations and anticipated interaction challenges expose the limitations of purely pre-programmed coordination for multi-robot construction tasks. Based on these findings, the paper reframes anchor-bracket installation as an execution-level adaptation problem within a validated task structure. This formulation motivates and organizes future investigation of learning-based methods to support robust coordination during execution, while preserving construction-driven task logic and safety constraints. Future work will implement and evaluate learning-based execution policies in a physics-based simulation environment, with physical validation planned as part of the ongoing AMALTEA project.

Acknowledgments

The research presented in this paper was originally carried out in the context of the Horizon project **AMALTEA**. **AMALTEA** has received funding from the EU's Horizon Europe program under Grant Agreement No. **101189678**.

References

[1] J. Collins, D. Howard, and J. Leitner. Quantifying the reality gap in robotic manipulation tasks. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 6706–6712, 2019. doi:10.1109/ICRA.2019.8793591.

[2] M. Giftthaler, T. Sandy, K. Dörfler, I. Brooks, M. Buckingham, G. Rey, M. Kohler, F. Gramazio, and J. Buchli. Mobile robotic fabrication at 1:1 scale: the in situ fabricator. 2017. doi:10.48550/arXiv.1701.03573.

[3] E. P. G. Bruun, S. Adriaenssens, and S. Parascho. Structural rigidity theory applied to the scaffold-free (dis)assembly of space frames using cooperative robotics. *Automation in Construction*, 141:104405, 2022. doi:10.1016/j.autcon.2022.104405.

[4] P. Eversmann, F. Gramazio, and M. Kohler. Robotic prefabrication of timber structures: towards automated large-scale spatial assembly. *Construction Robotics*, 1, 2017. doi:10.1007/s41693-017-0006-2.

[5] K. Orlowski. Automated manufacturing for timber-based panelised wall systems. *Automation in Construction*, 109:102988, 2020. doi:10.1016/j.autcon.2019.102988.

[6] K. Iturralde, M. Feucht, R. Hu, M. Schlandt, T. Linner, T. Bock, J.-B. Izard, I. Eskudero, M. Rodriguez, J. Gorrotxategi, J. Astudillo, J. Santos, M. Gouttefarde, M. Fabritius, C. Martin, T. Henninge, S. Nornes, Y. Jacobsen, and L. Elia. A cable driven parallel robot with a modular end effector for the installation of curtain wall modules. pages 1472–1479, 2020. doi:10.22260/ISARC2020/0204.

[7] S. Parascho. Construction robotics: From automation to collaboration. *Annual Review of Control, Robotics, and Autonomous Systems*, 6:183–204, 2023. doi:10.1146/annurev-control-080122-090049.

[8] Í. Elguea-Aguinaco, I. Inziarte-Hidalgo, S. Bøgh, and N. Arana-Arexolaleibo. A review on reinforcement learning for motion planning of robotic manip-

ulators. *International Journal of Intelligent Systems*, 2024:1636497, 2024. doi:10.1155/int/1636497.

[9] S. Yu, S. Lee, C.-S. Han, K. Y. Lee, and S.-H. Lee. Development of the curtain wall installation robot: Performance and efficiency tests at a construction site. *Auton. Robots*, 22:281–291, 2007. doi:10.1007/s10514-006-9019-2.

[10] A. K. Ali, O. J. Lee, and H. Song. Robot-based facade spatial assembly optimization. *Journal of Building Engineering*, 33:101556, 2021. doi:10.1016/j.jobe.2020.101556.

[11] C. Heuer, A. Lentes, Z. Zhao, and S. Brell-Çokcan. Robotic platform for the (semi-) automated assembly of façade panels. *Proceedings of the International Symposium on Automation and Robotics in Construction (IAARC)*, pages 1057–1064, 2024. URL `https://api.semanticscholar.org/CorpusID:270079168`.

[12] Raise Robotics. Precision automation for vertical construction. `https://raiserobotics.ai/`. Accessed: January 2026.

[13] A. A. Apolinarska, M. Pacher, H. Li, N. Cote, R. Pastrana, F. Gramazio, and M. Kohler. Robotic assembly of timber joints using reinforcement learning. *Automation in Construction*, 125:103569, 2021. doi:10.1016/j.autcon.2021.103569.

[14] G. Sartoretti, Y. Wu, W. Paivine, T. K. S. Kumar, S. Koenig, and H. Choset. Distributed reinforcement learning for multi-robot decentralized collective construction. 2018. ISBN 978-3-319-77503-6. doi:10.1007/978-3-030-05816-6_3.

[15] C.-J. Liang, V. R. Kamat, and C. C. Menassa. Teaching robots to perform quasi-repetitive construction tasks through human demonstration. *Automation in Construction*, 120:103370, 2020. doi:10.1016/j.autcon.2020.103370.

[16] L. Huang, Z. Zhu, and Z. Zou. To imitate or not to imitate: Boosting reinforcement learning-based construction robotic control for long-horizon tasks using virtual demonstrations. *Automation in Construction*, 146:104691, 2023. doi:10.1016/j.autcon.2022.104691.

[17] B. Belousov, B. Wibranek, J. Schneider, T. Schneider, G. Chalvatzaki, J. Peters, and O. Tessmann. Robotic architectural assembly with tactile skills: Simulation and optimization. *Automation in Construction*, 133:104006, 2022. doi:10.1016/j.autcon.2021.104006.

[18] K. Duan and Z. Zou. Enhancing construction robot collaboration via multiagent reinforcement learning. *Journal of Intelligent Construction*, 3(2):1–16, 2025. doi:10.26599/Jic.2025.9180089.

[19] X. Yang, Z. Ji, J. Wu, Y.-K. Lai, C. Wei, G. Liu, and R. Setchi. Hierarchical reinforcement learning with universal policies for multistep robotic manipulation. *IEEE Transactions on Neural Networks and Learning Systems*, 33(9):4727–4741, 2022. doi:10.1109/TNNLS.2021.3059912.

[20] Z. Sun, B. Pang, X. Yuan, X. Xu, Y. Song, R. Song, and Y. Li. Hierarchical reinforcement learning with curriculum demonstrations and goal-guided policies for sequential robotic manipulation. *Engineering Applications of Artificial Intelligence*, 153:110866, 2025. doi:10.1016/j.engappai.2025.110866.

[21] C. Shu, A. Kim, and S. Park. Subassembly to full assembly: Effective assembly sequence planning through graph-based reinforcement learning. pages 14946–14952, 2025. doi:10.1109/ICRA55743.2025.11127710.

[22] V. Soler. Robots (v1.9.0). `https://github.com/visose/robots`, 2025. Accessed: January 2026.

[23] J. K. Verma and V. Ranga. Multi-robot coordination analysis, taxonomy, challenges and future scope. *Journal of Intelligent and Robotic Systems*, 102(1): 10, 2021. doi:10.1007/s10846-021-01378-2.

[24] H. Guo, F. Wu, Y. Qin, R. Li, K. Li, and K. Li. Recent trends in task and motion planning for robotics: A survey. *ACM Computing Surveys*, 55(13s):289, 2023. doi:10.1145/3583136.

[25] D. Ruan, W. McGee, and A. Adel. Reducing uncertainty in multi-robot construction through perception modelling and adaptive fabrication. In *Proceedings of the International Symposium on Automation and Robotics in Construction (ISARC)*, 2023. doi:10.22260/ISARC2023/0006.