Adaptive Meta-Imitation Learning for Design-Aware Robotic Assembly

A Case of Parametric Brick Laying Work

Abhishek Mehrotra[1], Hwang Yi[2] and One Jae Lee[3]
[1,2]*Korea University*
[3]*Haenglim Architects & Engineers Co. Ltd*
[1]*abhi17@korea.ac.kr, 0000-0002-3123-3574*
[2]*hwyi@korea.ac.kr, 0000-0002-9386-2456*
[3]*onelee37@gmail.com, 0000-0003-1300-0168*

Abstract. Amid growing interest in robotic construction, fully autonomous, parametric architectural fabrication remains limited, with existing systems lacking the adaptive intelligence needed for design-aware automation. To address this gap, we present a self-learning framework that enables a robotic arm to adapt to parametrically defined wall geometries using Meta-Imitation Learning (MIL). Two canonical wall typologies—linear and curvilinear—were systematically varied across five design classes using seven control parameters: wall start point, rotation angle, curvature amplitude, curvature frequency, brick displacement, rotational offset, and inter-brick spacing. Experiments were conducted in PyBullet simulation and dataset of 400 tasks were generated, each with 15 trials (6,000 episodes). Of these, 95% were allocated for meta-training and 5% for testing, with an 80–20 train–validation split within the training set. A five-shot MIL setup integrated Model-Agnostic Meta-Learning (MAML) for cross-design generalization and Behaviour Cloning (BC) for inner-loop adaptation. The control policy incorporated a proximity-based gripper activation strategy. Simulation results demonstrate strong generalization to unseen wall geometries, achieving an 82% brick-placement success rate with 42.5 mm positional tolerance and 3.5° yaw error. These findings validate the framework's capability for adaptive robotic fabrication under parametric variability.

Keywords. Meta Learning, Imitation Learning, Few-shot Learning, Robotic Bricklaying, Parametric Architecture

1. Introduction

Robotic fabrication has progressed rapidly across multiple disciplines owing to its precision, efficiency, and ability to realize complex geometries. The convergence of artificial intelligence (AI) and reinforcement learning (RL) with robotic fabrication has further expanded automation capabilities. However, their application within

HUMANISTIC COMPUTATION AND INTELLIGENCE, Proceedings of the 31st International Conference of the Association for Computer-Aided Architectural Design Research in Asia (CAADRIA) 2026, Volume 2, 327-336.

architectural fabrication and construction remains comparatively nascent. Deep reinforcement learning (DRL) offers strong potential for enabling robotics systems to acquire special reasoning and adaptive assembly competencies through iterative interaction with the environment. While DRL-based robotic control has been extensively explored in industrial manufacturing, its extension toward autonomous, design-adaptive fabrication workflows in architecture remains limited, primarily due to learning instability and generalization challenges across varying design conditions. Consequently, automated architectural fabrication demands not only effective self-learning methodologies but also mechanisms for continuous adaptation to parametric design variations.

Digital fabrication has gained growing attention in the architectural and construction domains (Ghiyasi, 2023; Mostafavi, 2023). Raković et al. (2014) demonstrated scripted bricklaying using RAPID programming, whereas Song and Hahm (Song, 2023) employed augmented reality–assisted coding for design-specific construction. Ruttico et al. (2024) developed a lightweight autonomous bricklaying robot for large-scale on-site applications. Maali Esfangareh et al. 2022) introduced Q-learning for robotic brick assembly, and Amtsberg et al. (2023) proposed a fabrication manager coordinating multi-agent timber assembly through human–robot collaboration. Xiong et al. (2025) presented vision-driven robotic form-finding for generative design exploration, while Groenewolt et al. (2023) and Bruun et al. (2024) explored cooperative multi-robot systems for timber construction and architectural reconstruction, respectively. Despite their innovation, these systems were predominantly rule-based or preprogrammed, exhibiting limited generalization and adaptability.

Research explicitly focusing on DRL-driven construction within architectural contexts remains sparse. Apolinarska et al. (2021) implemented an Ape-X Deep Deterministic Policy Gradient (Ape-X DDPG) framework for autonomous insertion of timber joints. Belousov et al. (2022) utilized model-free RL to assemble geometrically irregular blocks exhibiting complex contact dynamics, integrating RL with Rhino–Python–based structural optimization workflows. Leder et al. (2024) investigated multi-agent swarm assembly using RL-driven collective behaviours aligned with architectural objectives. Wibranek et al. (2021) applied heuristic–RL hybrid optimization for curve-following assembly of S- and L-shaped blocks, enabling 2D/3D freeform layouts. Recently, Mehrotra and Yi (2025) demonstrated a rule-based design formation pipeline using DRL-driven sequential planning, highlighting the influence of RL algorithm selection on design variability. Although these studies advance the use of self-learning in architectural fabrication, they lack complete adaptation to diverse design geometries, underscoring a persistent gap in design-aware robotic learning for architecture.

There is a significant research gap of self-training of robot arm for making parametric wall with geometrical design or parameters adaptation. Figure 1 is depicting the clear difference between previous studies and the current study, which is indicating a gap of exclusion of design adaptation for brick-wall generation. To address this limitation, the present study investigates how robotic agents can autonomously adapt to architectural design variations through self-learning–based robotic bricklaying. Meta-learning and multi-task learning offer promising frameworks for achieving global policy generalization across design domains. For example, Lu et al. (2024) introduced language-conditioned multitask manipulation using dynamic Gaussian splatting (ManiGaussian) for unstructured environments. Cho et al. (2024) proposed a meta imitation learning (MIL) with embodiment-aware state encoding for cross-robot transfer. Kasaei & Kasaei (2025) employed Interactive Few-shot Imitation Learning (IFIL) to enable rapid adaptation from limited demonstrations. Finn et al. (2017) introduced Model-Agnostic Meta-Learning for Imitation Learning (MIL), demonstrating improved sample efficiency and generalization across multiple robotic tasks. Wu et al. (2025) extended this concept by integrating Adaptive Dynamical Primitives into MIL for long-horizon manipulation. Across these studies, meta-learning combined with imitation learning has proven effective for few-shot adaptation in stacking and assembly tasks (Kasaei & Kasaei, 2025; Lu et al., 2024). Given that architectural fabrication represents a geometrically complex and precision-critical extension of tasks, we adopt MIL as the foundation for design-adaptive robotic fabrication.

This work aims to advance AI-enabled digital fabrication in architecture by embedding meta-learning–based adaptability into parametric, design-aware construction workflows. We propose a self-learning framework for robotic additive fabrication (Figure 1) in which a robot autonomously interprets and executes parametric wall designs through adaptive policy optimization, closing the gap between computational design intent and autonomous realization. To the best of our knowledge, this is the first application of MIL to a design-adaptive robotic fabrication task requiring high positional and alignment precision.

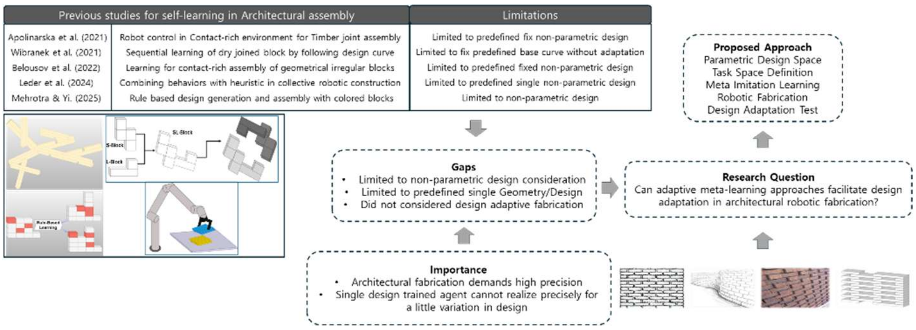

Figure 1. Study overview schematic

2. Material and Methods

The proposed design-aware robotic learning framework was implemented entirely in simulation using the PyBullet physics engine (v3.2.7), which provided forward/inverse kinematics, contact dynamics, and gravity-based stability for expert trajectory generation and policy evaluation. The setup featured an ABB IRB 2600 manipulator with an OnRobot 2FGP20 gripper handling standard 230 × 113 × 59 mm clay bricks. The simulation environment was built on the OpenAI Gym framework (v0.26.2) for RL compatibility, with robot, gripper, and brick models defined via URDF. A virtual Intel RealSense D456 depth camera supplied synchronized RGB-D observations. All learning experiments used PyTorch (v2.8.0).

The proposed two canonical wall typologies; linear and curvilinear-served as the generative base for design-aware learning. These typologies were parametrically varied to synthesize a continuum of design alternatives, from which five representative wall configurations were defined as meta-task classes for adaptive policy training. Parametric variability was encoded through seven geometric parameters (Figure 2): wall start coordinate (x: -0.25 to 0.1, y: 0.9 to 1.1), angular rotation with respect to the global y-axis (-100 to 10), curvature amplitude (-0.15 to 0.15), curvature frequency (wave factor) (0.2 to 0.5), brick displacement along the wall normal (-0.05 to 0.05), brick rotational offset relative to the wall surface (-200 to 200), and inter-brick spacing (0.06 to 0.11). All these 7 parameters control the design of wall and make them any of the 5 types of walls which then can be used for bricklaying trajectories.

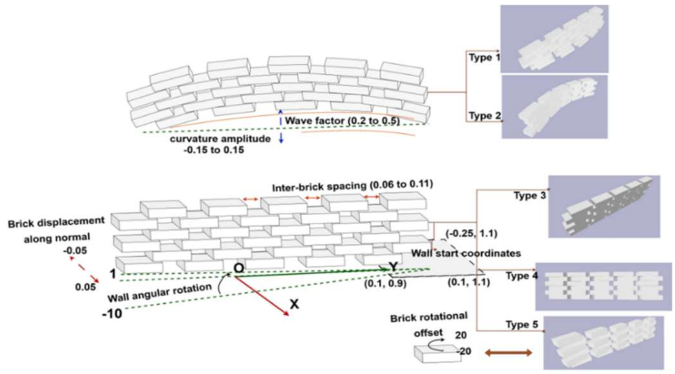

Figure 2. Wall design parameters description and ranges

Latin Hypercube Sampling (LHS) was used to ensure statistically uniform coverage of the parametric wall design space. Each sampled parameter vector generated a distinct wall configuration in PyBullet, providing diverse and structured conditions for meta-learning. Expert demonstrations were produced using a Bi-directional RRT (Bi-RRT) motion planner. For each sampled configuration, the target brick placement pose was first converted to joint angles via PyBullet's inverse kinematics (IK). These IK outputs, together with the robot's initial joint state, were passed to the Bi-RRT planner to compute collision-free joint trajectories for pick-and-place execution. The resulting joint sequences formed temporally ordered expert action paths. Near the wall, brick collision was evaluated for extra collision distance (1 cm) from neighbouring bricks to make the trajectory fully collision free.

Each trajectory was stored as paired state–action samples. States contained RGB-D frames (GIF-encoded), current and target gripper poses, and wall parameters. Actions were defined as incremental joint displacements between successive configurations. Figure 3 summarizes the planning pipeline and environment setup. During imitation learning, the policy network received encoded states and predicted the next joint increment, minimizing the discrepancy from expert actions to achieve accurate placement across varied wall geometries.

Design adaptation was achieved through a Meta-Imitation Learning (MIL) framework integrating demonstration-based learning with few-shot meta-learning.

Model-Agnostic Meta-Learning (MAML) provided outer-loop optimization to extract transferable priors across parametric wall classes, while Behavior Cloning (BC) served as the inner-loop adaptation mechanism for task-specific brick placement. This hierarchical setup enabled the agent to learn generalizable assembly strategies and adapt efficiently to new geometries with limited data. Meta-loss followed the standard formulation over trajectory observations and actions $T = \{o_1, a_1, \ldots, o_T, a_T\}$. Expert demonstrations—including end-effector poses and wall parameters—served as inputs to MIL. As shown in Figure 4, meta-tasks were uniformly sampled across five design classes to maintain balanced representation of typologies. Each task's trajectories were used for supervised inner-loop BC adaptation, while the meta-learner optimized its initialization to minimize post-adaptation loss across all tasks, improving generalization to unseen wall designs.

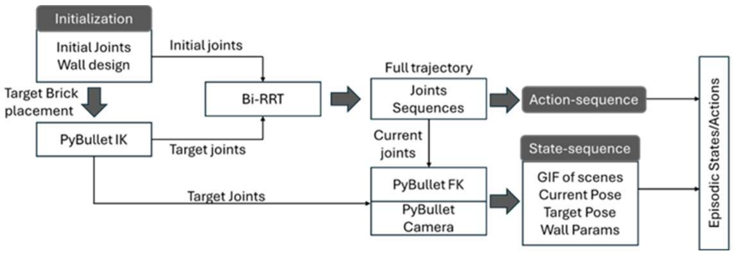

Figure 3. Process flow for expert trajectories generation process

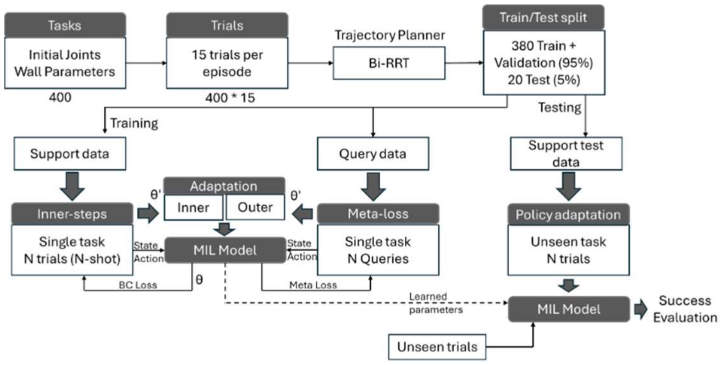

Figure 4. Meta learning framework with Imitation learning

In total, 400 tasks were generated, each with 15 trajectories under fixed wall parameters. Of these, 380 tasks (95%) were used for meta-training, split into 80% training and 20% validation, while 20 tasks (5%) were held out for evaluation. Each meta-batch sampled multiple tasks with N = 5-shot support and query trajectories. Training performed inner-loop updates via BC per task, followed by outer-loop meta-updates from aggregated query losses. After 60 meta-epochs, the policy was evaluated on unseen wall configuration using only inner-loop adaptation.

The policy network employed a CNN-based encoder–decoder architecture composed of four convolutional layers with integrated Squeeze-and-Excitation (SE) blocks and spatial attention modules placed after the first and second layers, respectively, followed by two fully connected layers [19]. The CNN extracted spatial and geometric features from RGB-D inputs, which were fused with parametric wall descriptors to generate continuous joint-angle commands for the 6-DOF manipulator. This multimodal fusion enabled the policy to leverage both visual and structural cues, improving robustness and precision during brick placement. The full MIL training pipeline is illustrated in Figure 5. Hyperparameters were optimized using Optuna, which performed iterative search and validation-driven selection after predefined training epochs. The resulting optimal configuration is summarized in Table 1.

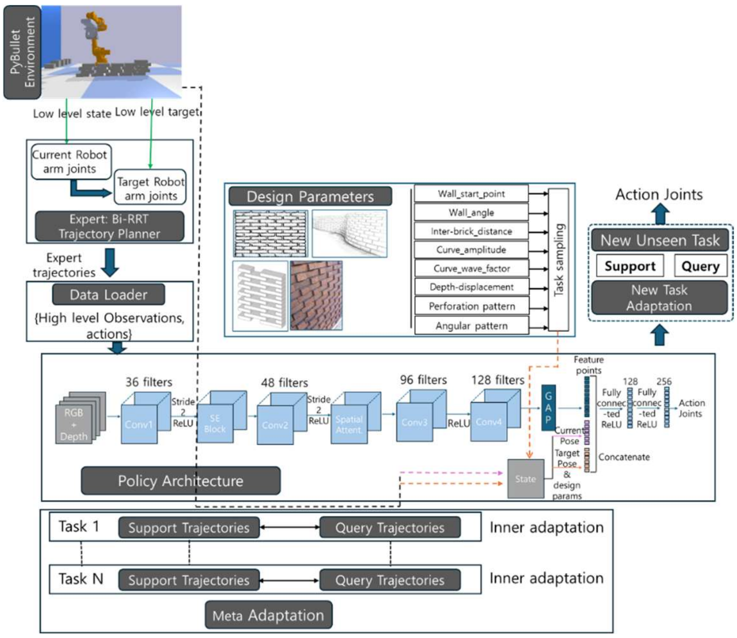

Figure 5. Network architecture of MIL model

Table 1. Hyper-parameters of MIL model

S. No.	Hyper-parameter	Value	Ranges
1	Epochs	60	--
2	Meta batch size	7	3 - 8
3	Inner gradients steps	1	1 - 5
4	Inner loop learning	0.0011	0.0001 – 0.005
5	Outer loop learning	0.00078	0.0001 - 0.001
6	First_order	True	True, False
7	Pose_loss	True	--
8	Position_loss_weight	1.482	0.1 – 1.5
9	Orientation_loss_weight	1.054	0.1 – 1.5

For evaluation, the trained MIL agent was tested on 20 unseen wall designs, each containing 15 demonstration trajectories. For every task, five trajectories (5-shot) were used for inner-loop adaptation, and the remaining trajectories were used for open-loop rollout evaluation. A placement was considered successful if it satisfied the predefined tolerance thresholds governed by the gripper signal: 42.5 mm positional accuracy and 3.5° yaw alignment.

Training effectiveness was quantified by (1) reduction in meta-loss across meta-epochs and (2) success rate during test rollouts. Five rollouts were executed per test task, and a success factor (0–1) was computed as the fraction of accurate placements. Overall performance was reported as the mean success factor across all tasks. All experiments were conducted on a workstation with an Intel® Core™ i9 (13th Gen) CPU, 64 GB RAM, and an NVIDIA® RTX 4070 GPU running Windows™ 11.

3. Results and Discussion

The learning performance of the MIL-based policy was evaluated both during training and during open-loop rollout on unseen tasks. Figure 6(a) and 6(b) present the evolution of the meta-loss (outer adaptation loss) and the corresponding mean support and query losses for tasks within each meta-batch over 60 training epochs. The meta-loss exhibited a smooth and stable reduction from an initial value of 3.0 to 1.5, while the mean support and query losses followed a similar trend. Notably, the support and query losses remained closely matched throughout training, converging even further after approximately 50 epochs, indicating effective inner-loop adaptation and consistent generalization across tasks.

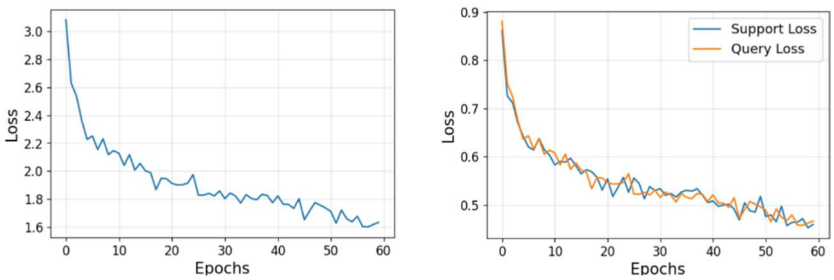

(a) (b)

Figure 6. (a) Meta loss and (b) Support and Query loss with epochs

Figure 7 presents the trial-level performance of the proposed MIL framework across 20 unseen test tasks under three gripper open-signal precision settings. The highest precision condition corresponds to the most challenging scenario, where the gripper is triggered to open upon reaching 42.5 mm proximity to the target position with a yaw deviation limited to 3.5°. The moderate precision setting relaxes this criterion slightly, activating the gripper at 42.5 mm distance with 4.7° yaw variation, while the least stringent condition allows opening at 57.5 mm with 4.7° yaw variation. In the figure, blue markers represent individual trial successes for the most precise placement, with the blue curve denoting the aggregated task-level success rate. Orange and green markers and curves depict the corresponding results for medium and lowprecision placements, respectively. Under the highest precision setting, most tasks achieved success rates between 40% and 100%, whereas for the relaxed precision conditions, success rates ranged from 60% to 100%, reflecting the increasing task feasibility as placement tolerance is relaxed. Overall, the MIL-based agent demonstrated strong generalization to unseen wall geometries, achieving an average success rate of 81% under the strictest precision criteria and exceeding 85% under more relaxed placement conditions.

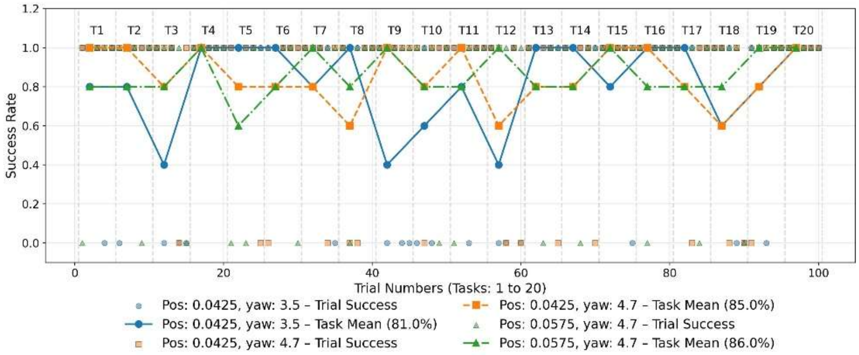

Figure 7. Success rate of each trial of each task

Simulation experiments demonstrated that the agent achieved more than 80% success rate during meta-testing despite relying on only five demonstration trajectories per meta-class, highlighting the efficacy of few-shot adaptation. These results indicate that meta-imitation learning can substantially reduce dependence on extensive demonstration datasets and preprogrammed routines, enabling autonomous robotic fabrication to efficiently accommodate parametric variations in architectural design. Previous studies also showed higher success rate but that was for assembly of predefined geometry without design adaptation. Figure 8 shows robotic successful placement for various parametric designs.

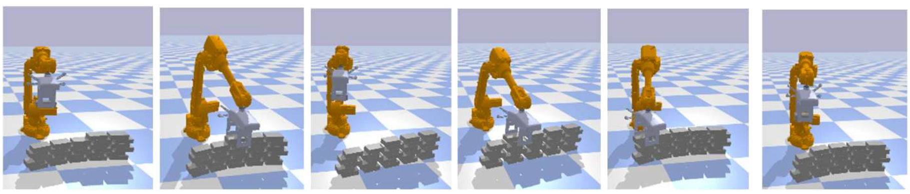

Figure 8. Robot arm design aware placement during evaluation

4. Conclusion

The study concludes that MIL offers a promising solution for bridging architectural design variability and autonomous robotic fabrication. By combining parametric design principles with demonstration-based self-learning, robotic agents can achieve meaningful adaptation in bricklaying tasks, advancing the integration of intelligent systems into digital fabrication workflows. The work specifically addresses a limitation in prior research by embedding adaptiveness within parametric design–to–fabrication workflows. Despite these contributions, several challenges remain, including enhancing placement accuracy, expanding the spectrum of constructible geometries, and migrating from simulated environments to deployment on physical construction sites. The proposed research trajectory holds clear relevance for architectural and construction applications involving complex additive assembly, yet demands focused investigation into precision control, structural stability, and multi-agent coordination. Future work will extend this framework to physical fabrication, addressing challenges such as material heterogeneity, tolerance management, and real-time error correction.

References

Amtsberg, F., Yang X., Skoury L., Calepso A. S., Wortmann T., Sedlmair M., & Menges A. (2023). Multi-Actor Fabrication for Digital Timber Construction. Dokonal, W, Hirschberg, U and Wurzer, G (Eds.), Digital Design Reconsidered - *Proceedings of the 41st Conference on Education and Research in Computer Aided Architectural Design in Europe, eCAADe 2023, Volume 1* (pp. 417–426).

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

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

Bruun, E. P. G., Oval, R., Al Asali, W., Gáspár, O., Paris, V., & Adriaenssens, S. (2024). Automating historical centering-minimizing masonry vaulting strategies: Applications to cooperative robotic construction. *Developments in the Built Environment*, 20, 100516. https://doi.org/10.1016/j.dibe.2024.100516

Cho, S., Kim, D., Lee, J., & Hong, S. (2024). Meta-Controller: Few-Shot Imitation of Unseen Embodiments and Tasks in Continuous Control (No. arXiv:2412.12147). arXiv. https://doi.org/10.48550/arXiv.2412.12147

Finn, C., Yu, T., Zhang, T., Abbeel, P., & Levine, S. (2017). One-Shot Visual Imitation Learning via Meta-Learning (No. arXiv:1709.04905). arXiv. https://doi.org/10.48550/arXiv.1709.04905

Ghiyasi, T., Zargar S. H., & Baghi A. (2023). Layer-by-Layer Pick and Place Collaboration Between Human and Robot Using Optimization. Dokonal, W, Hirschberg, U and Wurzer, G (Eds.), Digital Design Reconsidered - *Proceedings of the 41st Conference on Education and Research in Computer Aided Architectural Design in Europe, eCAADe 2023, Volume 1* (pp. 769–778). https://papers.cumincad.org/cgi-bin/works/paper/ecaade2023_423

Groenewolt, A., Krieg O., & Menges A. (2023). Collaborative Human-Robot Timber Construction. Dokonal, W, Hirschberg, U and Wurzer, G (Eds.), Digital Design Reconsidered - *Proceedings of the 41st Conference on Education and Research in Computer Aided Architectural Design in Europe, eCAADe 2023, Volume 1* (pp. 407–416). https://papers.cumincad.org/cgi-bin/works/paper/ecaade2023_414

Kasaei, H., & Kasaei, M. (2025). VITAL: Interactive Few-Shot Imitation Learning via Visual Human-in-the-Loop Corrections (No. arXiv:2407.21244). arXiv. https://doi.org/10.48550/arXiv.2407.21244

Leder, S., Kannenberg F., Siriwardena L., Schwinn T., & Menges A. (2024). Approaching the Augmentation of Heuristic Behaviors with Reinforcement Learning in Collective Robotic Construction. Herrera, Pablo C., Gómez, Paula, Estevez, Alberto T., Torreblanca-Díaz, David A. Biodigital Intelligent Systems - *Proceedings of the XXVIII Conference of the Iberoamerican Society of Digital Graphics, SIGraDi 2024* (pp. 1283–1294). https://papers.cumincad.org/cgi-bin/works/paper/sigradi2024_298

Lu, G., Zhang, S., Wang, Z., Liu, C., Lu, J., & Tang, Y. (2024). ManiGaussian: Dynamic Gaussian Splatting for Multi-task Robotic Manipulation (No. arXiv:2403.08321). arXiv. https://doi.org/10.48550/arXiv.2403.08321

Maali Esfangareh, A. (2022). ADOPTING AI TECHNIQUES IN ROBOTIC FABRICATION IN ARCHITECTURE: INTELLIGENT ROBOTIC BRICKLAYING USING REINFORCEMENT LEARNING ALGORITHMS [Master Thesis, Middle East Technical University]. https://open.metu.edu.tr/handle/11511/97933

Mehrotra, A., & Yi, H. (2025). Performance comparison of deep reinforcement robot-arm learning in sequential fabrication of rule-based building design form. *Frontiers of Architectural Research*, 14(6), 1654–1680. https://doi.org/10.1016/j.foar.2025.08.008

Mostafavi, S., Kemper B. N., Kretzer M., Etemadi A., Mahmoud H. A., Yaseen A., Nabizadeh M., Balazadeh T., & Chatterjee S. (2023). Cobotic Matters – Collaborative Robots and Discrete Assembly Design: From Stacking to Self-Interlocking of Reciprocal Components. Immanuel Koh, Dagmar Reinhardt, Mohammed Makki, Mona Khakhar, Nic Bao (Eds.), HUMAN-CENTRIC - *Proceedings of the 28th CAADRIA Conference* (pp. 241–250). https://papers.cumincad.org/cgi-bin/works/paper/caadria2023_

Raković, M., Jovanović, M., Borovac, B., Tepavčević, B., Nikolić, M., & Papović, M. (2014). Design and fabrication with industrial robot as brick-laying tool and with custom script utilization. *2014 23rd International Conference on Robotics in Alpe-Adria-Danube Region, RAAD* (pp. 1–5). https://doi.org/10.1109/RAAD.2014.7002251

Ruttico, P., Pacini, M., & Beltracchi, C. (2024). BRIX: An autonomous system for brick wall construction. *Construction Robotics*, 8(2), 10. https://doi.org/10.1007/s41693-024-00123-z

Song, Y. & Hahm, S. (2023). Augmented Robotic Bricklaying. Immanuel Koh, Dagmar Reinhardt, Mohammed Makki, Mona Khakhar, Nic Bao (Eds.), *HUMAN-CENTRIC - Proceedings of the 28th CAADRIA Conference* (pp. 323–332). https://papers.cumincad.org/cgi-bin/works/paper/caadria2023

Wibranek, B., Liu, Y., Funk, N., Belousov, B., Peters, J., & Tessmann, O. (n.d.). Reinforcement Learning for Sequential Assembly of SL- Blocks.

Wu, S., Wang, Y., & Huang, Y. (2025). One-Shot Robust Imitation Learning for Long-Horizon Visuomotor Tasks from Unsegmented Demonstrations [Proceedings Paper]. IROS 2025: IEEE/RSJ International Conference on Intelligent Robots and Systems. *Proceedings of IROS 2025*; IEEE. https://eprints.whiterose.ac.uk/id/eprint/230399/

Xiong, Y. & Wang, Y. (2025). Placing Nature: Interactive form-finding using computer vision and robotic arm collaboration for natural timber structure. Dagmar Reinhardt, Nicolas Rogeau, Christiane M. Herr, Anastasia Globa, Jielin Chen, Taro Narahara (Eds.), ARCHITECTURAL INFORMATICS - *Proceedings of the 30th CAADRIA Conference, Volume 2* (pp. 295–304). https://papers.cumincad.org/cgi-bin/works/paper/caadria2025_542