대한건축학회 창립80주년기념대회 및 추계학술발표대회논문집
제45권 제2호(통권 제84집) 2025년 10월 29일 ~ 10월 31일

2025년 추계학술발표대회 : 일반부문

디자인 적응형 로봇 벽돌 쌓기를 위한 온라인 모방 학습

Online imitation learning for design-adaptive robotic bricklaying

○ 메로트라 아비세크* 이 황**
Mehrotra, Abhishek Yi, Hwang

Abstract

Digital fabrication is advancing in construction, yet autonomous robotic realization of architectural designs remains underdeveloped. To address this gap, we propose a design-adaptive vision-based bricklaying framework where robotic agents self-learn to adapt parametrically varied wall typologies. Two wall types—linear and curvilinear—were varied through five control parameters to generate diverse design alternatives. Adaptation was accomplished through online imitation learning using the DAgger framework, followed by reinforcement learning fine-tuning with image-based observations that incorporated demonstrations to enable rapid generalization. Simulation results show 72% success in design-aware bricklaying across unseen parametric variants while baseline struggle to achieve moderate success even after training with many episodes.

Keywords: DAgger, Robotic Fabrication, Soft Actor-critic, Bricklaying, Imitation Learning

1. Introduction

The integration of artificial intelligence (AI) and reinforcement learning (RL) into robotic automation has advanced significantly, yet its application in architectural design and construction remains comparatively underdeveloped. While industrial robotic arms are widely utilized in manufacturing, their extension to autonomous, design-informed fabrication workflows in architecture is still in its formative stages. Deep reinforcement learning (DRL) presents substantial potential for enabling robotic agents to acquire advanced spatial reasoning and adaptive assembly skills through iterative interaction with their environment. However, research explicitly addressing DRL-driven design adaptation and rule-based construction within architectural contexts remains sparse.

A limited body of work has investigated self-learning approaches for robotic assembly in architecture. Apolinarska et al. (2021) introduced RL-based automation for timber joints using Ape-X Deep Deterministic Policy Gradient (Ape-X DDPG) for insertion tasks. Belousov et al. (2022) applied model-free RL to assemble geometrically irregular blocks with complex contact dynamics, coupling Rhino–Python workflows for structurally optimized design generation. Leder et al. (2024) explored RL-driven robotic swarms for collective assembly, aligning multiple agents with architectural objectives while incrementally refining cooperative behaviors. More recently, Mehrotra and Yi (2025) demonstrated rule-based design formation using DRL-driven sequential planning and highlighted how the selection of RL algorithms influences design variability while in another study, Mehrotra and Yi (2025) automated the robot agent for brick picking in 6D pose as the prior task for bricklaying.

In parallel, several studies have explored robotic fabrication without incorporating self-learning mechanisms. Raković et al. demonstrated scripted bricklaying using RAPID, while Song et al. employed augmented reality–based coding for design-specific construction. Esfangareh et al. applied Q-learning to bricklaying tasks, and Amstberg et al. developed a fabrication manager for multi-actor timber assembly through human–robot interaction. Xiong et al. proposed vision-driven robotic form-finding for generative design exploration, whereas Groenewolt et al. and Bruun et al. investigated human–robot collaboration in timber construction and multi-robot adaptation for historical architectural reconstruction, respectively.

A persistent limitation across these approaches is the lack of automated robotic fabrication capable of generalizing across parametric design variations without explicit pre-encoding of individual design instances. To bridge this gap, we present a design-aware robotic bricklaying framework that leverages parametric variation to enable autonomous agent adaptation

* 고려대학교 건축학과 박사과정
** 고려대학교 건축학과 교수, 건축학박사
(Corresponding author : Department of Architecture, Korea University, hwyi@korea.ac.kr)
이 연구는 2024년도 한국연구재단 연구비 지원에 의한 결과의 일부임. 과제번호: RS-2024-00353461

to previously unseen design configurations. The objective of this study is to develop a self-learning framework for design-adaptive bricklaying, leveraging vision-based observations and inference to enable autonomous adaptation to parametric design variations.

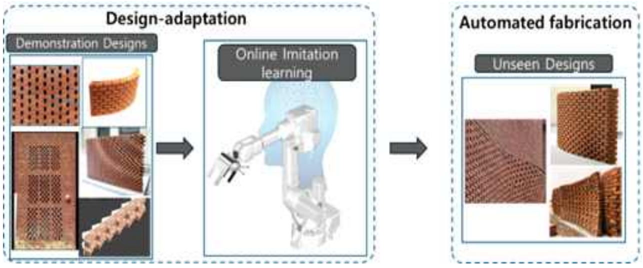

Figure 1. Automated Design-Aware Robotic Bricklaying

2. Materials and Methods

This study investigates DRL-based robotic self-learning for design-adaptive bricklaying, implemented within the PyBullet physics engine (v3.2.6) integrated with the OpenAI Gym framework (v0.26.2). The simulation environment incorporates the ABB IRB2600-20/1.6 robotic arm and standard modular bricks (230 × 113 × 59 mm), both modeled through URDF files. An OnRobot 2FGP20 gripper was simulated as the end-effector for brick placement, also defined via URDF. Visual observations for RGB and depth inference were obtained using an Intel RealSense D456 camera. Brick-picking from predefined piles was scripted, while placement was learned through imitation learning. Parametric wall designs were generated by varying two canonical templates—linear and curvilinear—using Latin Hypercube Sampling (LHS) across five design parameters: wall starting point, angular orientation relative to the y-axis, inter-brick spacing, sine-wave amplitude, and wave factor.

A five-layer Convolutional Neural Network (CNN) was employed for visual feature extraction, producing a 128-dimensional latent representation. These features were concatenated with a 17-element state vector comprising six current joint positions, six target pose parameters, and five wall design parameters. The combined input was passed through three fully connected layers to output joint-angle actions. Training was conducted in two sequential phases (Figure 2). In the first phase online imitation learning was done while in second phase, Soft Actor-Critic (SAC) fine tuning was done for precise placement. Dagger is utilized for online imitation learning, in which for initial 1000 episodes trajectories were saved and utilized for robot learning with simple behavior cloning, while in another 700 episodes policy training is controlled by query based supervised learning. In these 700 episodes, policy was used for rollouts (behavior) while expert actions were utilized for policy comparison. After the Dagger training, policy was fine tuned with SAC by keeping policy as actor and ran for 1300 episodes, when policy almost converged. During SAC fine tune, initially in no learning period only critic was updated only by keeping actor freeze and after 5000 steps only normal SAC operation started.

DAgger performance was benchmarked against a baseline Soft Actor-Critic (SAC) agent, both employing an identical policy architecture. The policy input consisted of concatenated visual features and state vectors (128,17), with the output being six robot joint angles. For DAgger, both observations and expert actions were used for supervised updates, whereas the SAC baseline relied solely on observations, directly mapping them to joint-angle actions. The SAC reward function was designed to address the high-precision requirements of bricklaying, which involves 6-DoF pose alignment. A shaped short-term reward combined

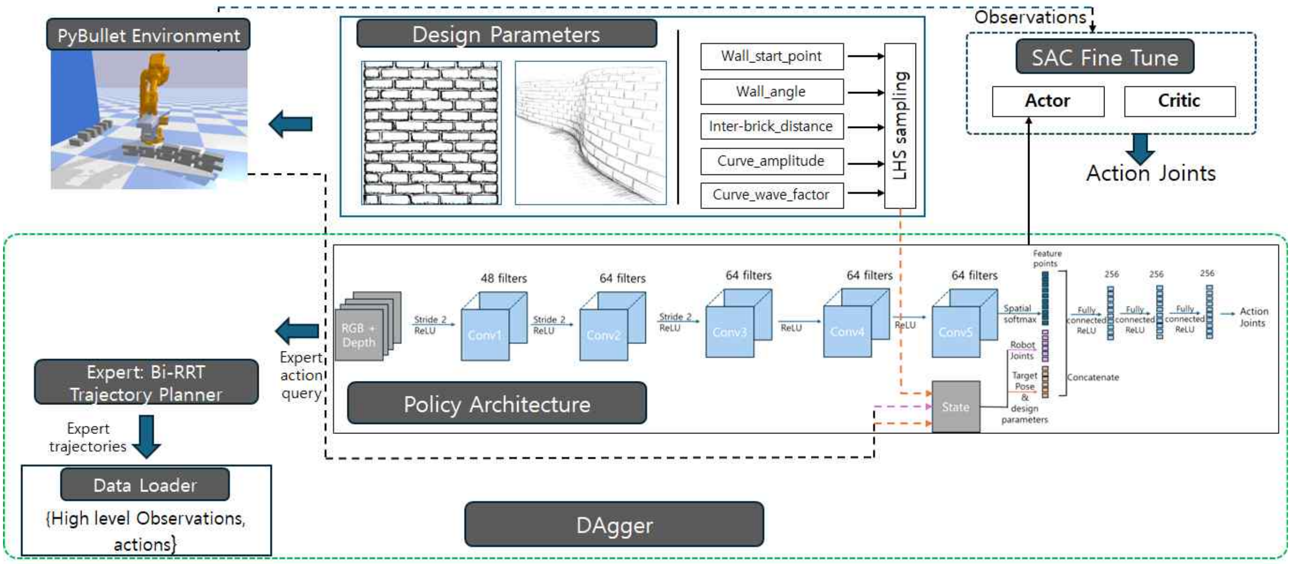

Figure 2. DAgger scheme

distance and orientation components (Eq. 1–4), while a sparse terminal reward of +1 was provided upon successful brick placement. A placement was considered successful if the positional error was below 0.15 m and the orientation error within 7° across all axes.

$$r_{dis} = -\|current\_pos - target\_pos\|_2| \qquad \text{(eq. 1)}$$

$$X_{angle} = -\left(1 - e^{\left(\frac{-angle\_diff_x}{0.35}\right)}\right) \qquad \text{(eq. 2)}$$

$$Y_{angle} = -\left(1 - e^{\left(\frac{-angle\_diff_y}{0.35}\right)}\right) \qquad \text{(eq. 3)}$$

$$r_{angle} = 0.5 \times \left(X_{angle} + Y_{angle}\right) + 0.5\left(Z_{angle}\right) \qquad \text{(eq. 4)}$$

3. Results and discussion

We evaluated performance by comparing the proposed DAgger scheme with SAC fine-tune against a baseline SAC agent in terms of average episodic rewards and success rates. Figure 3(a) illustrates the mean episodic rewards over batches of 20 episodes. During behavior cloning, rewards remained consistently high (>-10) for the first 50 batches (≈ 1000 steps). Although rewards declined when student rollouts began, the DAgger agent quickly recovered, surpassing baseline performance during SAC fine-tune and stabilizing at higher rewards after ~100 batches, whereas the baseline agent only reached comparable levels after ~150 batches. Figure 3(b) presents the average success rate over 25-episode batches. The DAgger agent achieved >80% success within the first 50 batches, dropped to ~40% during student rollouts and fine-tuning, but regained high performance within 150 batches just like baseline.

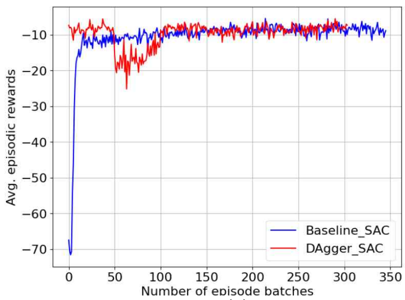

(a)

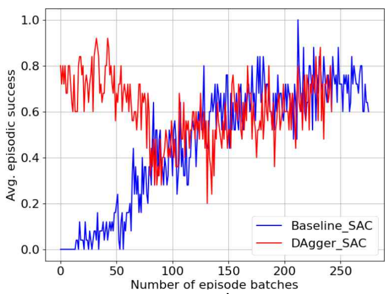

(b)

Figure 3. (a) Average episodic rewards, (b) Average episodic success for baseline and DAgger

However proposed scheme and baseline are showing similar success during training, adaptation of baseline was poor with new unseen tasks while our scheme showed 72 % success rate (Figure 4). These results confirm that the DAgger+SAC Fine-tuning framework enables more efficient and reliable adaptation for design-aware bricklaying. Figure 5 presents some screen shots of successful placement of brick with this learned design-aware policy.

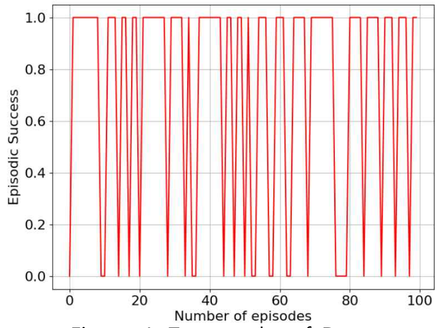

Figure 4. Test results of Dagger

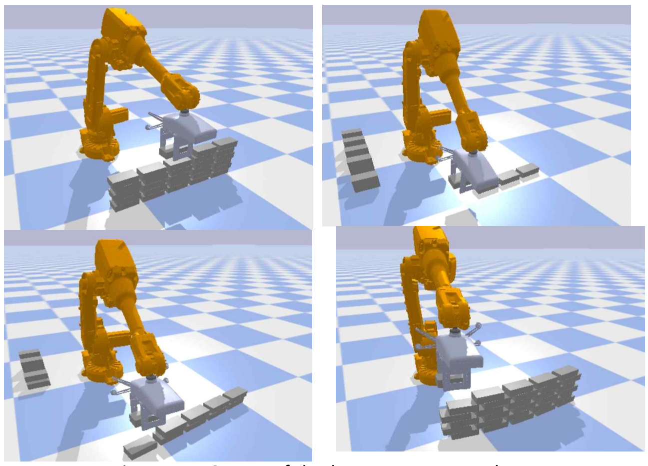

Figure 5. Successful placement snap shots

4. Conclusion

This study demonstrates that the proposed agent achieved significantly faster adaptation to design variations compared to baseline RL approaches. By integrating parametric design principles with demonstration-guided self-learning, the framework enabled robotic agents to perform design-aware bricklaying with robust adaptability, advancing the role of intelligent systems in digital fabrication workflows. Nonetheless, further investigations are required to evaluate effectiveness across broader parametric variations and to explore alternative learning algorithms for enhanced generalization and precision.

References

1. Apolinarska, A. A.; Pacher, M.; Li, H.; Cote, N.; Pastrana, R.; Gramazio, F.; Kohler, M. Robotic Assembly of Timber Joints Using Reinforcement Learning.

Automation in Construction **2021**, *125*, 103569. https://doi.org/10.1016/j.autcon.2021.103569.

2. Belousov, B.; Wibranek, B.; Schneider, J.; Schneider, T.; Chalvatzaki, G.; Peters, J.; Tessmann, O. Robotic Architectural Assembly with Tactile Skills: Simulation and Optimization. *Automation in Construction* **2022**, *133*, 104006. https://doi.org/10.1016/j.autcon.2021.104006.

3. Leder, S. Approaching the Augmentation of Heuristic Behaviors with Reinforcement Learning in Collective Robotic Construction. In *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) - ISBN 978-9915-9635-2-5, iBAG-UIC Barcelona, Spain, 13-15 November 2024, pp. 1283–1294*; CUMINCAD, 2024.

4. Mehrotra, A. and Yi, H. Performance Comparison of Deep Reinforcement Robot-Arm Learning in Sequential planning of Rule-based Design Form. Frontiers of Architectural Research 2025 (In Press)

5. Mehrotra, A. and Yi, H. Self-learning of Robot Arm using a Modified Soft-Actor-Critic (SAC) Approach: Complex 6D Pose Brick Grasping Test. *Journal of Construction Automation and Robotics* 2024, 3(2), pp. 1-10.

6. Raković, M.; Jovanović, M.; Borovac, B.; Tepavčević, B.; Nikolić, M.; Papović, M. Design and Fabrication with Industrial Robot as Brick-Laying Tool and with Custom Script Utilization. In *2014 23rd International Conference on Robotics in Alpe-Adria-Danube Region (RAAD)*; 2014; pp 1–5. https://doi.org/10.1109/RAAD.2014.7002251.

7. Song, Y. and H. Augmented Robotic Bricklaying. In *Immanuel Koh, Dagmar Reinhardt, Mohammed Makki, Mona Khakhar, Nic Bao (eds.), HUMAN-CENTRIC - Proceedings of the 28th CAADRIA Conference, Ahmedabad, 18-24 March 2023, pp. 323–332*; CUMINCAD, 2023.

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

9. Amtsberg, F. Multi-Actor Fabrication for Digital Timber Construction. In *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, Graz, 20-22 September 2023, pp. 417–426*; CUMINCAD, 2023.

10. Xiong, Y. and W. Placing Nature: Interactive Form-Finding Using Computer Vision and Robotic Arm Collaboration for Natural Timber Structure. In *Dagmar Reinhardt, Nicolas Rogeau, Christiane M. Herr, Anastasia Globa, Jielin Chen, Taro Narahara (eds.), ARCHITECTURAL INFORMATICS - Proceedings of the 30th CAADRIA Conference, Tokyo, 22-29 March 2025, Volume 2, pp. 295–304*; CUMINCAD, 2025.

11. Groenewolt, A. Collaborative Human-Robot Timber Construction. In *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, Graz, 20-22 September 2023, pp. 407–416*; CUMINCAD, 2023.

12. Bruun, E. P. G.; Oval, R.; Al Asali, W.; Gáspár, O.; Paris, V.; Adriaenssens, S. Automating Historical Centering-Minimizing Masonry Vaulting Strategies: Applications to Cooperative Robotic Construction. *Developments in the Built Environment* **2024**, *20*, 100516. https://doi.org/10.1016/j.dibe.2024.100516.