

Adaptive Human–Robot Collaboration for Masonry Construction Under Material and Assembly Uncertainty

Jutang Gao¹ and Arash Adel^{1*}

¹Princeton University, USA

*Corresponding Author

jutang.gao@princeton.edu, arash.adel@princeton.edu

Abstract -

Human–robot collaboration in construction is often challenged by limited robot-to-human communication and the need to adapt to tolerance accumulation arising from material and assembly uncertainties. We present an adaptive human–robot collaborative workflow for masonry construction that addresses communication limitations and tolerance accumulation, demonstrated through a brickwork case study in which a robot places bricks while a human applies adhesive. This workflow is enabled by two complementary mechanisms: 1) an end-effector-mounted projector that provides spatially registered, just-in-time projection guidance for manual adhesive application, and 2) laser scanning for feedback-driven grasping and placement pose correction. Together, these mechanisms enable adjustment of human and robotic actions in response to material variability and accumulated assembly tolerances. Full-scale experiments across conventional running-bond and nonstandard configurations demonstrate that projection guidance improves adhesive application consistency and reduces application time, while laser-based correction maintains level courses and avoids collision-prone failures associated with open-loop execution. These results indicate that integrating spatial projection with feedback-driven adaptation, enabled by material and as-built sensing, can mitigate tolerance accumulation and improve precision and robustness in human–robot collaborative construction.

Keywords -

Human–robot collaboration; Construction robotics; Spatial augmented reality; Adaptive fabrication; Brickwork

1 Introduction

Robotic systems are increasingly viewed as a means to improve productivity, safety, and working conditions in construction, while enabling mass customization, improved building performance, and novel architecture [1–4]. Yet the practical deployment of construction robotics remains limited, partly because robots must be integrated into human-centered workflows [5] and operate in unstructured, dynamic environments where material imperfec-

tions, as-built deviations, and environmental uncertainties challenge open-loop execution [1, 6, 7]. Human–robot collaboration has therefore been proposed and demonstrated as a way to improve adaptability, flexibility, and robustness by combining human dexterity, situated judgment, and material handling capabilities with robotic precision and repeatability [8–10]. Nevertheless, such collaboration remains constrained by limited and often non-intuitive methods for information exchange and task handoff [8], as well as limited robotic adaptability to material and as-built deviations [1]. Addressing this gap requires workflows that provide spatially explicit robot-to-human guidance while enabling feedback-driven robotic correction based on measured deviations.

This paper uses adhesive-based brickwork as a case study of human–robot collaborative masonry construction, where a human applies adhesive and the robot places bricks. While open-loop robotic brickwork has been explored in prior research [11], it is prone to tolerance accumulation due to variability in brick dimensions and placement inaccuracies. Unlike mortar-based masonry, where such tolerances can be compensated during placement, adhesive-based brickwork is far less tolerant due to thin bonding layers. This case study, therefore, highlights the necessity of feedback-driven, adaptive construction workflows. Accordingly, the paper’s main contributions are:

- We present a human–robot collaborative brickwork workflow that integrates projection-based robot-to-human guidance for adhesive application.
- We introduce an adaptive brick pose correction loop that uses laser-based sensing to detect placement tolerances and enforce explicit geometric constraints in the presence of uncertainty in brick and adhesive placement.
- We evaluate the workflow through full-scale physical experiments, quantifying performance in terms of adhesive application accuracy, tolerance handling, and cycle times for human and robot actions.

Altogether, this work proposes a transferable workflow for adaptive human–robot collaborative masonry construction under material variability and tolerance accumulation, which are inherent to real-world construction.

2 Related Work

This section contextualizes our work within three bodies of literature: human–robot collaboration in construction (Section 2.1), spatial augmented reality for information sharing (Section 2.2), and feedback-driven adaptive approaches for construction (Section 2.3).

2.1 Human–Robot Collaboration in Construction

Human–robot collaboration involves humans and robots working together to leverage complementary capabilities [2], combining robotic precision with human decision-making and adaptability. Prior research has emphasized the importance of legible and predictable robot motion for conveying intent [12], as well as communication interfaces that can enhance safety, productivity, and quality in collaborative processes [8, 13, 14]. However, many existing implementations still rely primarily on one-way information flow from humans to robots through direct commands, programming, and supervision [2]. To address this gap, this paper investigates just-in-time robot-to-human communication through spatially registered projection onto the work-in-progress structure to guide human actions in collaborative masonry.

2.2 Projection-Based Augmented Reality for Information Sharing

Prior work on projection-based spatial augmented reality for human–machine communication explores diverse information types, projection devices, and spatial configurations for effective information delivery [13, 15–17]. From this set of work, two key trade-offs emerge: 1) projection reach versus cost: high-power projectors provide extended reach and higher resolution at increased cost and energy demand, whereas low-power projectors are more cost-efficient but constrained in visibility range; 2) flexibility versus accuracy: fixed setups provide higher accuracy, whereas mobile setups provide extra flexibility but introduce localization challenges. This paper approaches this problem by mounting a low-cost, low-power projector on the robot end-effector and leveraging robotic positioning accuracy to enable flexible, spatially registered projection to aid human–robot collaborative masonry.

2.3 Feedback-Driven Adaptive Construction

Adaptive construction allows robots to adjust actions based on measured physical tolerances rather than executing open-loop plans [1]. Discrete assembly tasks in construction, such as brickwork, inherently involve tolerances in materials and processes, leading to accumulative errors if uncorrected [11]. Prior works argue that adaptivity in

such scenarios depends on perception to localize materials, compare as-built and as-planned states, and correct actions based on measurements [1]. Existing robotic systems for aiding masonry construction, such as Hadrian [18], leverage specialized machinery and standardized materials to circumvent this issue. For robot-aided masonry construction, the key challenge remains capturing and correcting geometric tolerances in situ before they propagate. Building on feedback-driven adaptive construction frameworks [1, 7], this paper proposes an adaptive pipeline for human–robot collaborative masonry construction in which laser sensing is used to measure brick poses, brick dimensions, and as-built course heights, thereby enabling precise grasping and placement while maintaining level courses and aligned wall edges.

3 Methods

The methods section describes the implementation of a human–robot collaborative masonry workflow and is organized to detail three components of the system: the workcell setup (Section 3.1); the modeling, mapping, and calibration of the end-effector–mounted projector for delivering spatially explicit visual cues (Section 3.2); and the adaptive pipeline (Section 3.3). Collectively, these components enable human guidance and robot adaptation, improving robustness and construction accuracy.

3.1 Workcell Setup

The robotic system operates in a shared human–robot workcell comprising a brick loading area, pickup platform, construction platform, and designated human safety zone (Figure 1). Two workers participate in the process: Worker A applies adhesive and performs local safety checks, while Worker B monitors system state, conducts global safety checks, and coordinates the robot. The construction process follows a repetitive pick-and-place cycle that combines robotic precision with human dexterity through spatially explicit communication (Figure 2). Before grasping each brick, the robot scans the brick’s top surface with a one-dimensional laser sensor to estimate its pose and dimensions. After grasping, the robot then moves to a projection pose and displays the planned brick footprint as a registered quadrilateral using the end-effector–mounted projector (Section 3.2). Worker A applies adhesive within the projected boundary, confirms completion with Worker B, and the robot resumes placement once the worker returns to the safety zone. Between courses, the robot scans the as-built surface with the laser sensor and applies corrections derived from these measurements to subsequent placements (Section 3.3).




Figure 1. Workcell setup for augmented human-robot collaborative masonry construction.




Figure 2. Collaborative masonry workflow integrating projection for guiding human action.

3.2 Projector Integration for Spatially Explicit Visual Cues in Human-Robot Collaboration

Prior work has integrated projector-camera units with robotic mechanisms, using projector-camera calibration, camera-based pose estimation, and kinematic control to register projected information onto physical surfaces [19]. Rather than estimating projector’s relative pose using camera observations, we estimate the projector-end-effector transformation through robotic contact probing and robot-kinematics-based pose estimation. Conceptually, this approach draws on hand-eye calibration [20], in which the fixed transformation between an eye-on-hand camera and

the robot gripper is recovered from camera observations across multiple robot poses. We adapt this calibration logic to estimate the transformation between the gripper-mounted projector and the robot end-effector, enabling just-in-time, spatially registered visual cue display for human-robot collaborative assembly.

We define each spatial visual cue as a set of 3D points $Q = \{\mathbf{q}_j \in \mathbb{R}^3\}$ with $j = 1, 2, 3, \dots$ in the world frame that encodes task-relevant information, e.g., locations of fasteners, centroids of surfaces, and, in the case of this study, a brick footprint for showing the placement plan. Let $\{W\}$ denote the world frame, $\{T\}$ the gripper tool center point (TCP) frame, and $\{P\}$ the projector optical frame, whose origin coincides with the projector’s effective viewing cone’s apex at the lens and whose Z axis aligns with the direction of projection. The projector is rigidly mounted to the side of the gripper, and its pose with respect to the gripper’s TCP is represented by a fixed homogeneous transformation ${}^T\mathbf{T}_P \in SE(3)$, which is obtained from the projector calibration process. The desired and reachable projector pose ${}^W\mathbf{T}_P$ is generated given Q and the robot base pose to ensure full projection coverage and a proper projection distance, which, in this study, is achieved through a fixed translation of Q ’s bounding box centroid for pose origin and directing the pose Z axis to the centroid. In operation, given ${}^T\mathbf{T}_P$ and ${}^W\mathbf{T}_P$, we calculate the target TCP pose ${}^W\mathbf{T}_T$ to position the projector:

$${}^W\mathbf{T}_T = {}^W\mathbf{T}_P {}^T\mathbf{T}_P^{-1} \quad (1)$$

To properly display the spatial visual cue Q through projection, the projector is modeled as a reversed pinhole

device, analogous to an inverse camera [21, 22], with its optical center treated as a point source due to its small aperture relative to the projection distance. The intrinsic parameters of the projector include horizontal and vertical field-of-view angles (FOV), (α_x, α_y) , as shown in Figure 3. For a point in Q expressed in the projector frame, ${}^P\mathbf{q} = [X, Y, Z]^T$, the corresponding normalized 2D coordinates on the projection image plane, $(u, v) \in [0, 1]^2$ are computed via:

$$u = \frac{1}{2} \left(1 + \frac{X}{Z \tan(\frac{1}{2}\alpha_x)} \right), \quad v = \frac{1}{2} \left(1 + \frac{Y}{Z \tan(\frac{1}{2}\alpha_y)} \right) \quad (2)$$

where $u = 0, 1$ correspond to the image plane's left and right bounds, and $v = 0, 1$ correspond to the image plane's top and bottom bounds.

To obtain accurate ${}^T\mathbf{T}_P$ and the FOV (α_x, α_y) , calibration is performed using a sufficiently large planar surface and a contact probe rigidly attached to the end-effector for locating its TCP physically. The process is shown in Figure 3. A full-frame rectangle is cast onto the planar surface from more than two projection poses, ensuring that all four projected corners lie on the surface. For projection $i = 1, 2, 3, \dots$, the robot TCP pose ${}^W\mathbf{T}_{T,i}$ is recorded, and the four projected corner locations are then marked on the surface and measured using the contact probe, yielding 3D coordinates ${}^W\mathbf{X}_{i,k}$ in the world frame, where $k \in \{1, 2, 3, 4\}$. All measured points are transformed into the TCP frame via homogeneous transformation:

$${}^T\mathbf{X}_{i,k} = {}^W\mathbf{T}_{T,i}^{-1} {}^W\mathbf{X}_{i,k} \quad (3)$$

For each corner k , a line in 3D space is fitted to the corresponding point set $\{{}^T\mathbf{X}_{i,k}\}$ using principal component analysis, yielding four lines that approximate the edges of the projector's viewing cone in the TCP frame (Figure 3). The projector optical center is estimated as the point $\mathbf{o}^*$ that minimizes the sum of squared distances to these lines:

$$\mathbf{o}^* = \arg \min_{\mathbf{o}} \sum_{k=1}^4 \left\| (\mathbf{I} - \mathbf{d}_k \mathbf{d}_k^T) (\mathbf{o} - \mathbf{x}_k) \right\|^2 \quad (4)$$

where $\mathbf{x}_k$ is a point on the k -th line and $\mathbf{d}_k$ is its unit direction vector. The projector optical axis is computed as the normalized average direction of the fitted lines. The projector's principal image axes are reconstructed by averaging the directions of opposing line pairs. The two principal image axes are then orthonormalized with the optical axis to form the projector pose ${}^T\mathbf{T}_P$ in the TCP frame. The projection FOV (α_x, α_y) is obtained geometrically from the reconstructed viewing cone.




Figure 3. Projector calibration: finding the accurate transformation from the gripper TCP pose to the projector pose.

3.3 Adaptive Construction Pipeline

The adaptive approach here is formulated as pose correction enabled through perception. The pipeline operates at two levels, as presented in Figure 4: brick-level adaptation to estimate individual brick pose and dimensions before grasping; and course-level adaptation, reflected in brick-placement pose correction to compensate for cumulative tolerances across courses. In this section, let $\{W\}$ denote the world frame, $\{b\}$ the brick's pose, $\{g\}$ the gripper TCP pose when grasping and $\{r\}$ when placing and releasing the brick.

3.3.1 Brick-Level Pose Estimation and Adaptive Grasping

For each brick, a predefined scan path is used to acquire a point cloud measurement of the brick's top face. The point cloud is then fitted with a minimum-area bounding box to estimate the brick's pose and dimensions. Let ${}^W\mathbf{T}'_b \in SE(3)$ denote the updated brick pose through scanning. This pose defines the grasping frame via:

$${}^W\mathbf{T}'_g = {}^W\mathbf{T}'_b {}^b\mathbf{T}_g \quad (5)$$

where predefined ${}^b\mathbf{T}_g$ encodes a fixed TCP grasping pose in the brick's frame.

3.3.2 Course-Level Alignment and Placement Correction

Before placing each course of bricks, the robot scans the top of the current course to estimate the as-built height by averaging the heights of the scanned points. Nominal placement poses ${}^W\mathbf{T}_r$ are precomputed from the design model, including reference vertical planes defining wall-edge alignments. Vertical adaptation adjusts the placement height based on the measured course top height, the

scanned brick height (measured in Section 3.3.1), and a constant adhesive thickness. Let z' be the measured surface height, h'_b the scanned brick's height, and h_a the predefined adhesive thickness (0.8 mm). The adjusted placing height is:

$$z'_r = z' + h'_b + h_a \quad (6)$$

Horizontal adaptation enforces wall-edge alignments by adjusting in-plane placement to compensate for perceived brick length variations (Section 3.3.1). To accommodate this correction, a 3-mm lateral gap is introduced between the nominal brick geometries. The final corrected placement pose for each brick is:

$${}^wT'_r = {}^wT_r \Delta T_z \Delta T_{xy} \quad (7)$$

where ΔT_z is a vertical correction shared within a course and ΔT_{xy} is a brick-specific lateral correction. Defining brick poses at the brick top center and apply vertical corrections enforces level courses, while lateral corrections maintaining wall-edge alignment under bricks' dimensional variability (Figure 4).




Figure 4. Brick pose correction based on brick-level and course-level scanning. Left: Placing correction. Right: Grasping correction.

4 Experimental Setup

An initial experiment without brick pose correction failed due to a collision resulting from accumulated grasping and placement tolerances (Figure 5), underscoring the necessity of the adaptive pipeline for collaborative masonry. Four successful experiments, integrated with brick pose correction, were then conducted, varying brickwork configurations and projection usage. Two configurations were tested: a nonstandard assembly of 28 bricks over seven courses and a conventional running-bond assembly of 21 bricks over seven courses, as presented in Figure 6. For each, two conditions were evaluated: *projection-augmented*, with the footprint projected during adhesive application, and *projection-verified*, with projection used only for verifying adhesive application quality. The experiment processes are documented in Figure 7.




Figure 5. The experiment without brick pose correction leads to collision and early failure.



Figure 6. Variations in brick configuration, including conventional (straight running bond) and non-standard (running bond with brick rotations).

4.1 Evaluation Metrics

The effectiveness of projection augmentation in supporting manual adhesive application was first evaluated. For each brick, images of the adhesive patch with the projected footprint were analyzed using a vision pipeline to compute: 1) region coverage (%), the fraction of the intended bonding area covered by identified adhesive pixels; 2) exposed adhesive (%), the fraction of identified adhesive pixels applied outside the target region; and 3) largest missed patch (%), defined as the largest contiguous uncovered area ratio within the target bonding region. These metrics jointly capture adhesive completeness, cleanliness, and continuity. The assemblies' geometric accuracy was evaluated from course-level laser scans. The scanned point clouds were segmented into brick faces, and the mean absolute error of point heights relative to the updated target heights was computed, reflecting the effectiveness of brick pose correction in enforcing level courses. Process efficiency was assessed by logging timestamps for each pick-and-place cycle and sub-processes, including robot traversal, laser scanning, data collection, adhesive application, and user-input wait time. This breakdown identifies time-consuming stages and quantifies the effects of projection guidance and adaptive sensing on total task duration.

5 Results and Discussion

Quantitative results on adhesive quality and application time are summarized in Table 1 and Figure 8. Across the experiments, effective adhesive coverage generally




Figure 7. Documented experiment processes: a. robotic manipulator and a finished nonstandard construct; b. course scanning; c. brick scanning; d. brick grasping; e. footprint projection; f. adhesive application; g. projection-augmented condition; h. unaugmented condition (human collaborator trying to predict brick pose); i. brick placement; j & k: brick placement in the tightly constrained conventional configuration.

Table 1. Adhesive Coverage Analysis

Experiment condition	Region coverage(%)	Exposed adhesive(%)	Largest missed patch(%)	Application time (s)
Nonstandard, Projection Augmented	49.1 ± 9.1	2.4 ± 2.1	7.2 ± 13.3	25.2 ± 11.7
Nonstandard, Projection Verified	39.9 ± 12.6	19.6 ± 22.0	17.8 ± 23.2	26.9 ± 9.0
Conventional, Projection Augmented	40.5 ± 5.2	0.6 ± 0.8	10.2 ± 6.4	15.1 ± 8.4
Conventional, Projection Verified	40.1 ± 4.2	0.7 ± 0.6	9.4 ± 5.1	24.7 ± 9.0

stabilized between 40–50%. However, the nonstandard projection-verified case showed large fluctuations, including some near-zero values. Together with the higher coverage observed in the nonstandard projection-augmented case, these results suggest that projection guidance improves the consistency of human-applied adhesive coverage in the nonstandard case. Similar trends are observed in exposed adhesive and the largest missed patch metrics. Compared with the projection-augmented case, the projection-verified nonstandard case exhibits highly scattered exposed-adhesive ratios (up to $> 90\%$) and higher missed-patch ratios, indicating imprecise application and material waste. These findings suggest that projection augmentation improves the quality of human-applied adhesive for nonstandard brickwork. For the conventional running-bond configuration, adhesive quality was comparable with and without projection. However, the reduced application time in the augmented case indicates faster execution without loss of quality. Overall, the method is shown to effectively augment human performance by improving adhesive application accuracy in nonstandard brickwork and application speed in conventional brickwork.

As shown in Figure 9, the brick pose correction process limited height deviations to approximately 1.0 mm. This demonstrates effective cancellation of cumulative assembly tolerances. With brick pose correction, collision failures due to inaccurate placement (Figure 5) were not observed, indicating improved placement quality.

Projection-augmented nonstandard cases achieved a lower average mean absolute error than projection-verified cases, suggesting the significance of precise adhesive application in reducing course height deviations.

The time analysis shows that, across all experiments, the average time to process a single brick is 123.1 s. Aside from data collection (14.2%) and adhesive application (15.9%), robot traversal accounts for over 42.6% of the total cycle time, largely due to the conservative motion speed (250 mm/s) adopted for safety. User input waiting was the second-largest contributor (20.3%), revealing limitations of terminal-based interaction and motivating future multimodal communication. At this preliminary stage, a 1D laser sensor was used for ease of implementation; data collection and scanning overheads are expected to decrease (currently 7.0%) with higher-dimensional sensing, as shown in prior work [1].

6 Conclusion

This work presents a human–robot collaborative masonry workflow that integrates projection-based robot-to-human guidance with feedback-driven adaptive brick placement. Using adhesive-based brickwork as the case study, an end-effector-mounted projector displays desired adhesive regions directly on the work-in-progress structure, while laser-based sensing supports brick-level grasp correction and course-level placement adaptation. Comparative experiments show that projection guidance im-





Figure 8. Adhesive coverage identification and analysis. Numerical values are reported in Table 1.

proves adhesive application accuracy in nonstandard configurations and reduces application time in conventional running-bond patterns. At the same time, the adaptive pipeline maintains course-height accuracy and avoids collision-prone failures associated with open-loop execution. Together, these results demonstrate how projection-augmented interaction and feedback-driven adaptation can improve precision and robustness in human-robot collaborative masonry construction.




Figure 9. Height mean absolute error (MAE) analysis based on course-level scanning results.

Current system includes limited integration of active safety monitoring and a constrained human-to-robot communication interface, which currently relies on indirect terminal-based input. Future work will expand the system in several directions: incorporating direct communication modalities, such as gesture and speech, to support more intuitive collaboration; integrating 2D/3D perception to reduce scanning overhead; and adding perception-based safety modules for more robust operation. In addition, a user study with worker feedback is needed to evaluate system performance from a human-centered perspective.

Acknowledgments

This research was supported by Princeton University. The work presented in this paper builds on a course project developed in collaboration with Yilin Lin and Samet Yilmaz. We thank Daniel Ruan and Salma Mozaffari for their feedback and technical support, and the members of the Adel Research Group for their assistance during system development and experiments.

References

- [1] Arash Adel, Daniel Ruan, Wesley McGee, and Salma Mozaffari. Feedback-Driven Adaptive Multi-Robot Timber Construction. *Automation in Construction*, 164:105444, 2024.
- [2] Ci-Jyun Liang, Xi Wang, Vineet R. Kamat, and Carol C. Menassa. Human-Robot Collaboration in Construction: Classification and Research Trends. *Journal of Construction Engineering and Management*, 147(10):03121006, 2021.

- [3] Arash Adel. Co-Robotic Assembly of Nonstandard Timber Structures. In *Hybrids & Haecceities, Proceedings of the 42nd Annual Conference of the Association for Computer Aided Design in Architecture (ACADIA)*, pages 604–613. CumInCAD, 2022.
- [4] Ryan Craney and Arash Adel. Engrained Performance: Performance-Driven Computational Design of a Robotically Assembled Shingle Facade System. In *Distributed Proximities, Proceedings of the 40th Annual Conference of the Association of Computer Aided Design in Architecture (ACADIA)*, pages 604–613. CumInCAD, 2020.
- [5] Delgado Davila, Manuel Juan, Lukumon Oyedele, Anuoluwapo Ajayi, Lukman Akanbi, Olugbenga Akinade, Muhammad Bilal, and Hakeem Owolabi. Robotics and Automated Systems in Construction. *Journal of Building Engineering*, 26:100868, 2019.
- [6] Zhengyi Chen and Arash Adel. Advancing Robotic Assembly in Construction: Innovations, Challenges, and Opportunities. *Automation in Construction*, 178: 106370, 2025.
- [7] Daniel Ruan, Wes McGee, and Arash Adel. Reducing Uncertainty in Multi-Robot Construction through Perception Modelling and Adaptive Fabrication. In *Proceedings of the 40th International Symposium on Automation and Robotics in Construction (ISARC)*, pages 25–31, 2023.
- [8] Ondrej Kyjanek, Bahar Al Bahar, Lauren Vasey, Benedikt Wannemacher, and Achim Menges. Implementation of an Augmented Reality AR Workflow for Human Robot Collaboration in Timber Prefabrication. In *Proceedings of the 36th International Symposium on Automation and Robotics in Construction (ISARC)*, pages 1223–1230, 2019.
- [9] Arash Adel. *Computational Design for Cooperative Robotic Assembly of Nonstandard Timber Frame Buildings*. PhD thesis, ETH Zurich, 2020.
- [10] Xiliu Yang, Felix Amsberg, Michael Sedlmair, and Achim Menges. Challenges and potential for human–robot collaboration in timber prefabrication. *Automation in Construction*, 160:105333, 4 2024. ISSN 09265805.
- [11] Tobias Bonwetsch. *Robotically Assembled Brickwork: Manipulating Assembly Processes of Discrete Elements*. PhD thesis, ETH Zurich, 2015.
- [12] Anca D. Dragan, Kenton C.T. Lee, and Siddhartha S. Srinivasa. Legibility and Predictability of Robot Motion. In *2013 8th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, 2013.
- [13] Jie Zhang, Shuxia Wang, Weiping He, Jianghong Li, Zhiwei Cao, and Bingzhao Wei. Projected Augmented Reality Assembly Assistance System Supporting Multi-Modal Interaction. *The International Journal of Advanced Manufacturing Technology*, 123(3):1353–1367, 2022.
- [14] Cynthia Brosque, Elena Galbally, Oussama Khatib, and Martin Fischer. Human-Robot Collaboration in Construction: Opportunities and Challenges. In *2020 International Congress on Human-Computer Interaction, Optimization and Robotic Applications (HORA)*, pages 1–8, 2020.
- [15] SangJun Ahn, SangUk Han, and Mohamed Al-Hussein. 2D Drawing Visualization Framework for Applying Projection-Based Augmented Reality in a Panelized Construction Manufacturing Facility: Proof of Concept. *Journal of Computing in Civil Engineering*, 33(5):04019032, 2019.
- [16] Amir Degani, Wen Bo Li, Rafael Sacks, and Ling Ma. An Automated System for Projection of Interior Construction Layouts. *IEEE Transactions on Automation Science and Engineering*, 16(4):1825–1835, 2019.
- [17] Pedro Tavares, Carlos M. Costa, Luís Rocha, Pedro Malaca, Pedro Costa, António P. Moreira, Armando Sousa, and Germano Veiga. Collaborative Welding System using BIM for Robotic Reprogramming and Spatial Augmented Reality. *Automation in Construction*, 106:102825, 2019.
- [18] FBR Ltd. Hadrian. URL <https://www.fbr.com.au/view/hadrian>. access date: Mar 22, 2026.
- [19] Ahyun Lee, Joo-Haeng Lee, and Jaehong Kim. Data-Driven Kinematic Control for Robotic Spatial Augmented Reality System with Loose Kinematic Specifications. *ETRI Journal*, 38(2):337–346, 2016.
- [20] Roger Y. Tsai and Reimar K. Lenz. A New Technique for Fully Autonomous and Efficient 3D Robotics Hand/Eye Calibration. 5(3):345–358, 1989.
- [21] R. Tsai. A Versatile Camera Calibration Technique for High-Accuracy 3D Machine Vision Metrology Using Off-The-Shelf TV Cameras and Lenses. *IEEE Journal on Robotics and Automation*, 3(4):323–344, 1987.
- [22] Li Zhang and Shree Nayar. Projection Defocus Analysis for Scene Capture and Image Display. In *ACM SIGGRAPH 2006*, page 907, Boston, Massachusetts, 2006. ACM Press.