Specification-Driven Dual-Reference Error Control for Sequential Robotic Floor Tiling

Junxi Cai[1] and Yantao Yu[1,*]

***Abstract*— Construction specifications judge floor tiling by relations to adjacent as-built tiles and to a setting-out line, not by a single design pose. This paper shows how such requirements can be organised into step-wise decision constraints for a tiling robot. Dual-reference error control (DREC) writes joint width, tangential joint offset, row-to-line distance and row orientation as linear bounds on the planar correction (x, y, yaw) of the next tile and solves a strictly convex quadratic program before release. Within a no-rework scope, local joint bounds stay hard and global bounds are relaxed only against a non-empty local set. A remaining-capacity allocator lets each correction account for the placements still available in the row. In simulation, the dual reference reduced the median per-run RMS lateral deviation from 2.936 to 0.713 mm under a 10 mm setting-out disturbance. With three image-estimated channels, allocation recovered to ± 0.5 mm in 7/9 runs versus 5/9 for immediate correction and reduced the total variation of correction requests from 10.07 to 1.75 mm. In an as-laid audit, all joint widths stayed within 3–8 mm, and tangential offset magnitudes exceeded 1 mm at 13.0 % of joints with allocation versus 21.3 % without.**

Index Terms—Robotic floor tiling, construction specifications, as-built feedback, dual-reference control, constrained placement optimisation.

I. INTRODUCTION

Tiling quality is judged against two references at once. A joint is accepted by its width and alignment relative to the adjacent bedded tile, while a row is accepted by its position relative to the setting-out line. In a long sequence these references drift apart: bedded tiles leave their design positions, so a robot that follows only the previous tile passes deviations along the row, and a robot that tracks only the line may spoil the joint it is forming. Each placement also changes how much correction later tiles can still provide.

Hong Kong's Construction Industry Council forecasts an industry-wide gap of 10,000–15,000 skilled and semi-skilled workers in 2026, widening to 45,000–50,000 by 2030 [1], which strengthens the case for automating repetitive trades such as tiling. Tiling acceptance is already written as explicit relations: the General Specification for Building [2] refers floor tiling to BS 5385-3 [3] and requires straight, even joints at least 3 mm wide for the laying methods it covers, and BS 5385-3 checks joint straightness along the full joint against a string line or laser (Table I).

[1]The authors are with the Department of Civil and Environmental Engineering, The Hong Kong University of Science and Technology, Clear Water Bay, Hong Kong SAR, China. jcai@connect.ust.hk

*Corresponding author: ceyantao@ust.hk

Supplementary videos of the comparative evaluations and of the full 144-tile robotic execution are available at github.com/CaiSir0075/DREC-Robotic-Tiling.

Tiling robots have used seam and edge vision for placement and in-process quality monitoring [4], [5], visual alignment under a limited field of view [6] and Stewart-platform alignment of heavy tiles [7], and tiling quality has been improved through process parameters and adhesives [8]. As-built measurements have adjusted later placements in timber and masonry construction [9], [10], and tolerance information has been formalised for robot tasks [11]. The step decision that links these elements remains open: which references govern the current tile, what to keep when they conflict, and how to share a correction that cannot be finished now with the tiles that remain.

We therefore ask: *how can quality requirements set out in construction specifications be converted into step-wise decision constraints for the robot, so that quality is controlled over a long sequence of floor-tile placements?* DREC (Fig. 1) answers with three connected elements. Four quality relations with local and global references become linear constraints on one planar correction. A local-priority rule keeps joint bounds hard and relaxes only global rows when both cannot be met, reflecting a process in which released tiles are not reworked. A remaining-capacity allocator (Sec. II-D) turns the row's outstanding offset into a scalar request that enters the objective. Because the requirements are explicit bounds, they are encoded rather than learned, and a standard quadratic program solves each step. Simulations show the value of the dual reference and of allocation under partial image feedback, including an as-laid joint audit.

II. METHOD

A. Quality relations as step constraints

Relations and bounds are compiled manually from specification and project requirements (Table I). The placement planner supplies the nominal candidate pose (t_0, θ_0); in the image-feedback runs, its along-row coordinate is updated from the commanded laying frontier, the current tile-size

TABLE I

FROM SPECIFICATION REQUIREMENTS TO CONTROLLER QUANTITIES

relation	ref.	qty.	bound	source
joint width	tile	g_k	3–8 mm	spec.^{a}/proj.
joint offset	tile	e_k	$\lvert e_k \rvert \leq 1$ mm	control^b
line distance	line	d_k	$\lvert d_k \rvert \leq 5.5$ mm	project
row orientation	line	θ_k	$\lvert \theta_k \rvert \leq 0.01$ rad	project

Tile: adjacent as-built tile; line: setting-out reference. ^{a}3 mm minimum (not $\pm$3 mm): GS for Building 2022 Cl. 18.118(a)(iv),(b)(v), BS 5385-3:2024 Cl. 10.1.6; 8 mm is a project value. ^{b}BS 5385-3 Cl. 10.1.3.2 checks straightness along the whole joint without a per-tile limit.

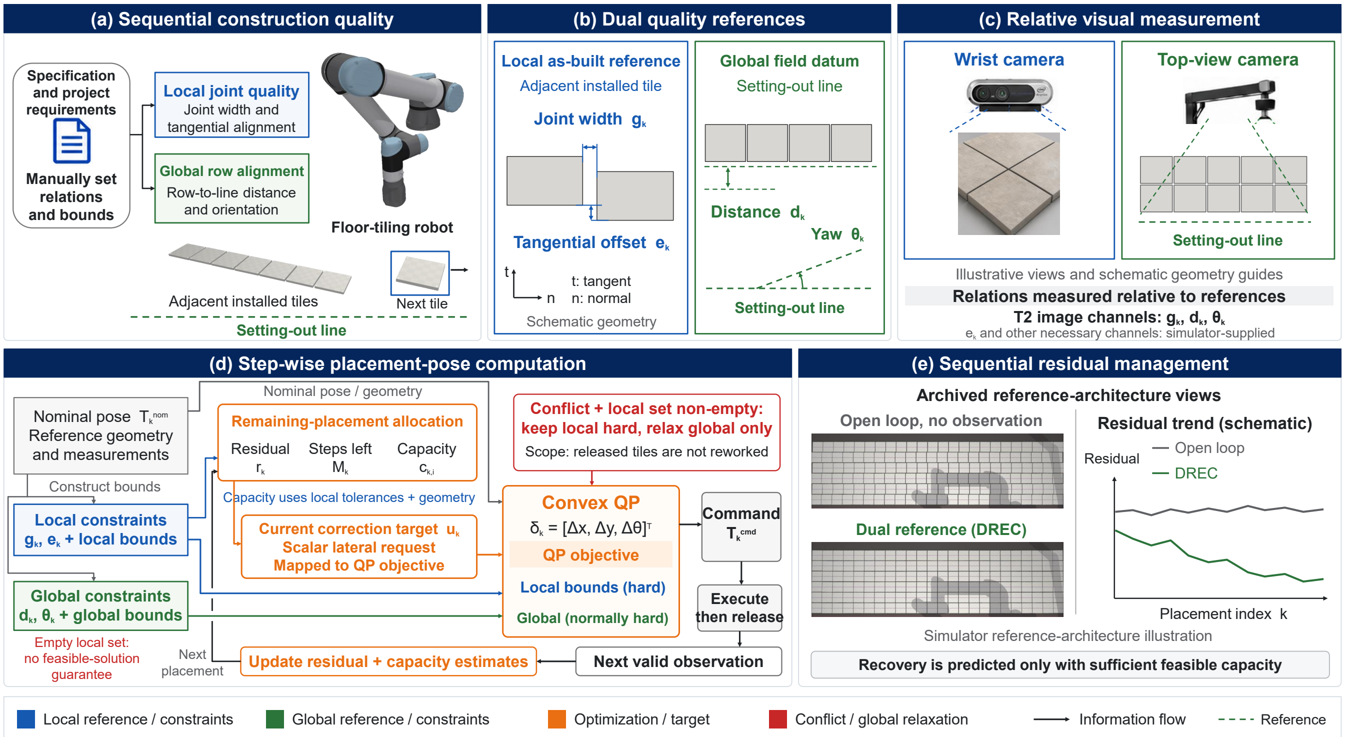

Fig. 1. DREC overview (schematic). (a) Requirements set manually as relations and bounds. (b) Local (g_k, e_k) and global (d_k, θ_k) references. (c) Illustrative views; in Sec. III-C only g_k, d_k, θ_k are image-estimated. (d) Bounds constrain δ_k directly; the allocator's scalar request u_k enters the objective; the red branch keeps local rows hard when the local set is non-empty (Sec. II-C). (e) Archived simulator views and schematic trends, not measured data.

input and the remaining row length. The adjacent installed tile does not set this pose but provides the local quality reference. The correction is expressed in a work frame F fixed during the solve:

$$\begin{aligned}\delta_k &= [\Delta x,\ \Delta y,\ \Delta\theta]^\top,\\ p_a(\delta_k) &= t_0 + [\Delta x, \Delta y]^\top + R(\theta_0+\Delta\theta)\,a,\\ J_a &= \big[\,I_2\ \ R(\theta_0)\,\Omega a\,\big],\quad \Omega = \big[\begin{smallmatrix}0&-1\\1&0\end{smallmatrix}\big],\end{aligned}\tag{1}$$

where a is a tile-fixed anchor and J_a linearises p_a at $\delta_k = 0$. The neighbour's joint normal n and tangent t and the normal m of the reference line stay fixed during the solve. With the facing-edge midpoint a_f for the local rows and the tile centre $a_c = 0$ for the line row, the four requirements become

$$\begin{aligned}g_{\min} &\le g_k^0 + j_g\delta_k \le g_{\max}, & j_g &= n^\top J_{a_f},\\ -e_{\max} &\le e_k^0 + j_e\delta_k \le e_{\max}, & j_e &= t^\top J_{a_f},\\ -d_{\max} &\le d_k^0 + j_d\delta_k \le d_{\max}, & j_d &= m^\top J_{a_c},\\ -\theta_{\max} &\le \theta_k^0 + j_\theta\delta_k \le \theta_{\max}, & j_\theta &= [0,0,1],\end{aligned}\tag{2}$$

where g_k^0, e_k^0, d_k^0, θ_k^0 are the measured current relation values. Stacking them in q_k^0 and $A = [j_g; j_e; j_d; j_\theta]$ gives $l \le A\delta_k \le u$ with $l = q_{\min} - q_k^0$ and $u = q_{\max} - q_k^0$; the first two rows form the local block and the last two the global block. The reference line is the controller's target line for tile centres, with any design offset from the struck line absorbed into its position. Each row thus ties the current command to one requirement. The rows are linearised and constrain the command; as-laid quality is evaluated separately after release, at the control-side corner of the facing edge, from which perception also derives the midpoint relations.

B. Objective

The objective balances a local and a global pose target,

$$\begin{aligned}J(\delta) &= \tfrac{1}{2}w_l\|W(\delta-\delta_l)\|^2 + \tfrac{1}{2}w_g\|W(\delta-\delta_g)\|^2,\\ H &= (w_l + w_g)\,W^\top W,\\ f &= -W^\top W\,(w_l\delta_l + w_g\delta_g),\end{aligned}\tag{3}$$

with $W = \mathrm{diag}(1, 1, L)$, tile size $L = 100\,\mathrm{mm}$, $w_l = 1$ and $w_g = 0.5$, so $H \succ 0$ and each feasible step has a unique solution. The targets are minimum-norm solutions in scaled coordinates,

$$\begin{aligned}\delta_l &= W^{-1}(B_l W^{-1})^+\,[\,g_{\mathrm{nom}} - g_k^0,\ -e_k^0\,]^\top,\\ \delta_g &= W^{-1}(A_g W^{-1})^+\,[\,b_k,\ -\theta_k^0\,]^\top,\end{aligned}\tag{4}$$

with $g_{\mathrm{nom}} = 5.5\,\mathrm{mm}$, $A_g = [j_d; j_\theta]$, and B_l equal to $[j_g; j_e]$ except that the tangential row has no yaw entry, so the local target does not use yaw to correct offset while the constraints still account for it. Without allocation, $b_k = -d_k^0$.

The allocator issues a scalar lateral request u_k, which is converted into the global target. Without active constraints, the minimiser $\delta^{\mathrm{free}} = (w_l\delta_l + w_g\delta_g)/(w_l + w_g)$ moves laterally by

$$j_d\,\delta^{\mathrm{free}} = \alpha\, b_k + \beta_k,\quad \alpha = \tfrac{w_g}{w_l+w_g} = \tfrac{1}{3},\quad \beta_k = \tfrac{w_l}{w_l+w_g}\,j_d\delta_l,\tag{5}$$

so setting $b_k = (u_k - \beta_k)/\alpha$ makes the unconstrained step deliver u_k; only δ_g and f change. Active constraints may then limit the planned lateral correction.

C. Local priority

Each step first solves, with OSQP,

$$\min_{\delta} \ \tfrac{1}{2}\,\delta^\top H\delta + f^\top\delta \quad \text{s.t.} \quad l \le A\delta \le u. \tag{6}$$

If this solve is not successful, a relaxed problem keeps the local rows and softens only the global rows,

$$\min_{\delta,\, s\ge 0} \ \tfrac{1}{2}\,\delta^\top H\delta + f^\top\delta + \rho\|s\|^2$$
$$\text{s.t. } l_{\text{loc}} \le A_{\text{loc}}\delta \le u_{\text{loc}}, \ \ l_{\text{glob}} - s \le A_{\text{glob}}\delta \le u_{\text{glob}} + s, \tag{7}$$

with $\rho = 10^4$ and slack in row units. Whenever the local set is non-empty, (7) is feasible and its solution meets every local row, so a conflicting line target can move the command only within the joint-feasible set. This encodes the process choice: without rework a joint defect would remain, whereas a row offset can still be reduced by later tiles. The acceptance limits are unchanged, and an empty local set cannot be recovered by relaxing global rows.

D. Remaining-capacity allocation

The allocator uses r_k, the controller's pre-correction lateral residual estimate at step k, formed from d_k^0 and refreshed at the next valid observation. With M_k tiles left in the row, including the current one, capacity is estimated at zero yaw from nominal geometry and local bounds: a translation $v = a\,n + b\,t$ with admissible joint changes $a \in [l_g, u_g]$ and $b \in [l_e, u_e]$ shifts the tile laterally by

$$m^\top v = (m^\top n)\,a + (m^\top t)\,b, \tag{8}$$

whose reachable interval follows from the endpoints; the QP itself still optimises yaw. In the rectangular layout m is antiparallel to t, giving 1 mm per direction. A modelled execution reserve of 0.835 mm, a bound on the clipped simulation disturbances, leaves a reserved capacity $c_{k,i} = 0.165$ mm per tile. The window and schedule are

$$h_k^\star = \min\Big\{h \le M_k : \textstyle\sum_{i=1}^{h} c_{k,i} \ge |r_k|\Big\},$$
$$\min_{v} \ \tfrac{1}{2}\textstyle\sum_{i=1}^{h_k^\star} v_i^2 \ \ \text{s.t.} \ \ \textstyle\sum_{i=1}^{h_k^\star} v_i = -r_k, \quad v_i \in \mathcal{V}_{k,i}, \tag{9}$$

with $\mathcal{V}_{k,i}$ the reserved reachable interval; only $u_k = v_1$ is requested, and the plan is recomputed at the next step. The shortest window completes the predicted recovery early, while the minimum-energy split keeps each step moderate; spreading over all M_k tiles would give smaller steps but a longer-lasting offset. No reserved window exists when $|r_k| > (0.165\,\text{mm})M_k$; since this test uses the current estimate r_k, not the injected disturbance amplitude, the fallback is decided anew at each step. It then schedules over the whole remaining row with unreserved capacities and records the reserve shortfall; if even these fall short, it schedules the reachable part and predicts a terminal residual. The comparison policy M1s requests $u_k = -r_k$ through (5) and leaves truncation to the constraints. Predicted recovery assumes that the capacities are jointly realisable and that execution and estimation follow the model.

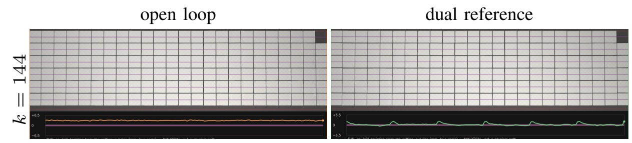

Fig. 2. Completed fields at 10 mm setting-out disturbance, one seed; purple lines are the physical setting-out lines, and strips show settled deviation at true scale.

TABLE II
REFERENCE ARCHITECTURE UNDER SETTING-OUT DISTURBANCE

	setting-out disturbance scale (mm)				
reference architecture	0	1	3	6	10
open loop, no observation	0.170	0.336	0.897	1.770	2.936
local as-built only	0.602	0.660	1.034	1.791	2.910
dual, every tile	0.253	0.266	0.327	0.465	**0.713**
dual, every third tile	0.368	0.407	0.493	0.744	1.154

Median over 12 seeds of the RMS as-laid lateral deviation of 144 tiles (mm); simulator-provided references.

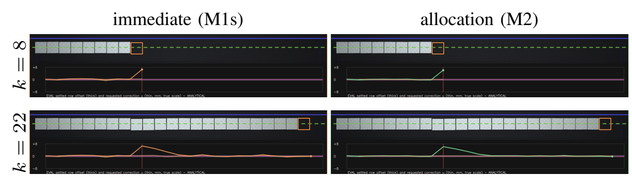

Fig. 3. 6 mm lateral disturbance at zero-based column 8, one seed, shown at zero-based $k = 8$ and $k = 22$. Bands: as-laid row offset (thick) and request u_k (thin), true scale.

III. EXPERIMENTS

A. Setup

All experiments use MuJoCo with a UR5e and vacuum end effector on a transport rail, laying a 6×24 field of 144 tiles of 100 mm with 5.5 mm joints under modelled tile-size, grasp and release disturbances. Compared configurations share the disturbance configuration through common random numbers. The architecture sweep uses simulator-provided references and an earlier controller. In the allocation runs and in T2, d_k, θ_k and g_k are estimated from rendered top and wrist images, while e_k, the joint reference axes and the picked tile's size and grasp are simulator-generated; when the wrist estimate lost the joint, a declared nominal-design value was used. The allocation runs use a legacy row-end handling that differs from T1/T2.

B. Dual reference reduces lateral deviation under setting-out disturbance

Table II and Fig. 2 compare four architectures over 12 seeds. Observing the line at every tile keeps the median per-run RMS lateral deviation below 0.5 mm up to a 6 mm setting-out disturbance and cuts it by about three quarters at 10 mm, where open loop and local-only control remain near 2.9 mm; observing the line at every third tile retains much of this gain. The local-only arm is worse than open loop without disturbance, consistent with deviations propagating along neighbours. The line channel is thus what brings an offset site datum back into the decision.

TABLE III

RECOVERY POLICY UNDER PARTIAL IMAGE FEEDBACK (18 RUNS)

metric	M1s	M2	pairs
runs recovered to ± 0.5 mm	5/9	7/9	—
placements to recover[a]	17	3	—
both recovered[a]	17	2	5/5
$TV(u)$ of requests (mm)[a]	10.07	1.75	9/9
as-laid $g \notin [3, 8]$ mm[b]	0/207	0/207	—
as-laid $\lvert e \rvert > 1$ mm[b]	44/207	27/207	—
slack / non-solved / degraded	0/0/0	0/0/0	—

Pairs: (seed, position) pairs with M2 lower. [a]Medians; recovery steps over recovering runs ($n = 5$, $n = 7$) or over the 5 pairs where both recovered. [b]Offline audit of within-row joints; counts are descriptive, as joints are correlated.

C. Allocation improves recovery under image feedback

Eighteen runs (2 policies $\times$ 3 seeds $\times$ 3 positions) lay all 144 tiles with a 6 mm lateral disturbance at zero-based column 2, 8 or 20 of one row; recovery statistics use that 24-tile row. A run recovers when the as-laid row offset enters ± 0.5 mm and stays there to the row end.

Allocation recovered more often and with smoother correction requests (Table III, Fig. 3). M2 recovered in 7 of 9 runs and M1s in 5; in the five pairs where both recovered, M2 needed fewer placements in every pair. Its pre-QP requests u_k varied far less, with lower total variation $TV(u) = \sum_k |u_{k+1} - u_k|$ in all nine pairs, as expected when each request is a bounded share of a schedule rather than the full current estimate. Five of the six non-recovering runs had the latest disturbance, where few tiles remain. The global-relaxation branch was not triggered in any run, and every step was solved.

An offline audit of these runs checked the 23 within-row joints of the disturbed row per run at the control-side facing-edge corner, using settled poses and actual tile sizes and including joints next to the disturbance. All 414 joint widths stayed within 3–8 mm, and M2 lowered the share of tangential offset magnitudes exceeding 1 mm from 21.3 % to 13.0 %. The remaining exceedances show that a feasible command does not by itself fix the as-laid offset, and they give the tangential channel a concrete target.

D. Image joint width affects its own relation

T1 estimates d_k and θ_k from images; T2 also estimates g_k. Over three seeds of 144 placements, switching to the image joint width raised the median paired planned joint-width RMSE by 0.057 mm, in all three seeds, while the row-offset RMSE changed by only $+0.0009$ mm. In these runs, the added measurement cost stayed in the relation that the new channel measures. In the T2 focus row, the top channel was available at every placement and the wrist width at 95.7 % of applicable joints, and all runs completed 144 placements.

IV. DISCUSSION AND CONCLUSION

The experiments support the design of DREC. The line reference returns site-datum errors to the decision, and allocation converts a row offset into moderate, scheduled correction requests that recovered more runs under partial image feedback. The local block keeps feasible commands within the joint bounds by construction. The as-laid audit links these decisions to released joints: widths stayed in range, and allocation reduced the proportion of audited joints with tangential offset magnitudes exceeding 1 mm.

These results come from simulation with synthetic images, uncalibrated disturbance models and three seeds in the policy comparison, so they describe the tested conditions rather than field performance or statistical significance. The tangential channel, joint axes and picked-tile geometry were simulator-generated; lighting, dust, reflective mortar, wet-adhesive sliding and suction compliance were not modelled. No global-only or penalty-only baseline was run, the global-relaxation branch was not triggered in the image runs, and the audit covers within-row joints at one anchor. Allocation is per row, although from the second row on a tile also meets the previous row, whose joints add constraints and reduce later capacity.

DREC shows that specification-derived quality relations can drive each placement: joint and line requirements become explicit step constraints, local joint bounds are kept hard within the no-rework scope, and outstanding global correction is shared with the remaining tiles through a capacity-limited request. In simulation this reduced lateral deviation under setting-out disturbance, improved recovery with smoother correction requests under partial image feedback, and kept all audited joint widths in range. The next step is to close the tangential image channel and validate cross-row capacity on physical tiling trials.

REFERENCES

[1] Construction Industry Council, *Manpower Forecast for Hong Kong Construction Industry: Construction Manpower Forecast Results*. Hong Kong, Apr. 2026.

[2] Architectural Services Department, *General Specification for Building*, 2022 Ed. Hong Kong: Government of the Hong Kong SAR, 2022.

[3] British Standards Institution, *BS 5385-3:2024 Wall and Floor Tiling—Part 3: Design and Installation of Internal and External Ceramic and Mosaic, Natural Stone and Terrazzo Flooring, and Internal Agglomerated Stone Flooring in Normal Conditions—Code of Practice*. London: BSI, 2024.

[4] D. Apostolopoulos, H. Schempf, and J. West, "Mobile robot for automatic installation of floor tiles," *Proc. IEEE ICRA*, 1996, pp. 3652–3657.

[5] R. Navon, "Process and quality control with a video camera, for a floor-tilling robot," *Autom. Constr.*, vol. 10, pp. 113–125, 2000.

[6] S. Wang, H. Zhou, Z. Zhang, X. Zheng, and Y. Lv, "Robot floor-tiling control method based on finite-state machine and visual measurement in limited FOV," *Adv. Civ. Eng.*, vol. 2021, art. no. 8372815, 2021.

[7] S. Chang, Z. Lyu, J. Chen, T. Hu, R. Feng, and H. Liang, "Automated six-degree-of-freedom Stewart platform for heavy floor tiling," *Autom. Constr.*, vol. 170, art. no. 105932, 2025.

[8] K. Wu, Y. Zhang, X. Kong, S. Zhang, and L. Gao, "Quality control of robotic floor-tiling by the modifications on technology parameters and adhesive properties," *J. Field Robot.*, vol. 42, pp. 356–372, 2025.

[9] A. Adel, D. Ruan, W. McGee, and S. Mozaffari, "Feedback-driven adaptive multi-robot timber construction," *Autom. Constr.*, vol. 164, art. no. 105444, 2024.

[10] J. Gao and A. Adel, "Adaptive human-robot collaboration for masonry construction under material and assembly uncertainty," in *Proc. 43rd Int. Symp. Automat. Robot. Constr. (ISARC)*, 2026, pp. 618–625.

[11] S. K. Mohammed, M. H. Arbo, and L. Tingelstad, "Using semantic geometric dimensioning and tolerancing information from STEP AP242 neutral exchange files for robotic applications," *Int. J. Interact. Des. Manuf.*, vol. 18, pp. 6587–6603, 2024.