Retention Is Not Selection: An Audited Failure Analysis of Multi-Hypothesis Homography Release

XuanHao Liang

Shenzhen MSU-BIT University, Shenzhen, China
1793407303@qq.com

Abstract Repetitive planar appearance can produce coherent but spatially wrong registrations. Multi-Hypothesis Ambiguity Rejection (MHAR) retains distinct homography candidates in a fixed-capacity bank and releases only when the bank selected by its frozen rule is SINGLETON. On the predeclared 1,620-pair LinTrack registry, M-V8 contains a reference-consistent mode in 1,183 pairs versus 881 for B-R0; all 338 MHAR releases match B-TOP1 in candidate ID, matrix, and support set. At the same coverage, MHAR has 64 wrong homographies versus 34 for evaluator-only B-CONF-Match@338. Alias-positive periodic pairs yield empty valid pools. Historical robot outputs provide retrospective diagnostics; a coordinate audit identifies a unit mismatch and no corrected downstream replay is claimed. The study separates candidate retention, release quality, adapter implementation, and downstream utility.

Keywords homography estimation; multi-hypothesis registration; ambiguity rejection; repetitive planar appearance; robotic surface inspection

1. Introduction

Wall-climbing platforms enable close-range inspection where manual access is costly or unsafe. Prior systems combine robotic acquisition with mosaicking, localization, and defect mapping [1,2]. In our pipeline, the mosaic is the spatial index for revisiting wall regions and projecting frame-level observations.

Repetitive appearance can produce false pairwise registrations and globally inconsistent alignments [3], potentially displacing mapped regions or projected observations. Figure 1 illustrates the acquisition context and competing-mode motivation.

RANSAC-based estimators fit one model robustly [4,5], while multi-model methods estimate several instances [6,7] and stitching methods may retain multiple registrations [3,8]. We ask which pairwise modes should survive under fixed bank capacity and when a legacy single-edge back end should abstain.

MHAR retains several homography modes, applies a frozen bank-selection rule, and releases an edge only when the selected bank is SINGLETON. Retention, release, and support-preserving Sim(2) adaptation are evaluated separately at $K = 8$.

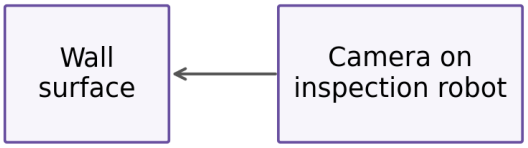
We make three contributions:

1) A fixed-capacity multi-hypothesis interface that retains complementary geometric cells and correspondence supports before a legacy single-edge back end.

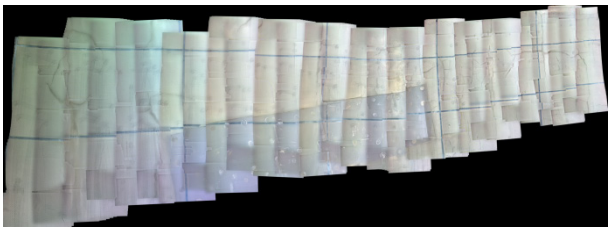
2) A deterministic bank-selection and withholding rule coupled to support-preserving Sim(2) adaptation.

3) A controlled audit showing that additional retained modes do not replace Top-1 on the tested registry and that MHAR has no matched-coverage release-quality advantage, together with periodic-texture and robot diagnostics.

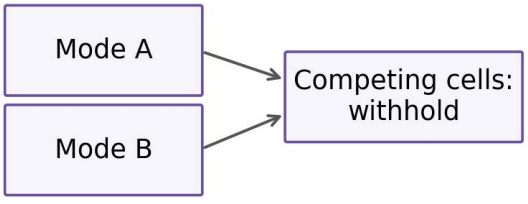
ISSN 3106-7018 (Print) ISSN 3106-5791 (Online)



(a) Acquisition



(b) Repetitive wall



(c) Competing modes



(d) Multi-strip mosaic example

Fig. 1. Task and registration motivation: (a) acquisition schematic; (b) repetitive-wall mosaic; (c) competing-mode schematic; (d) multi-strip mosaic. Photographs are qualitative examples.

2 Related Work

2.1 Robotic surface inspection

Sano et al. combine a wall-climbing robot with mosaicking and localization [1], while Yang et al. integrate robotic acquisition, localization, defect perception, and reconstruction [2]. Our contribution is limited to deciding whether a pairwise constraint should enter an existing mapping back end.

2.2 Robust, multi-model, and global registration

RANSAC and Graph-Cut RANSAC fit a single model robustly [4,5], while Progressive-X and CONSAC estimate multiple models [6,7]. Other work selects globally consistent registrations [3], handles multiple registrations and seams [8], or suppresses inconsistent graph constraints [9–12]. These lines of work distinguish model fitting, local alternative retention, and global inconsistency resolution; MHAR addresses only the zero-or-one-edge interface for an unchanged Sim(2) back end.

2.3 Selective prediction and abstention

Selective prediction evaluates error jointly with acceptance coverage [13]. We borrow this risk–coverage framing: MHAR withholds a pair unless the selected bank contains one observed geometric mode, then refits its support to Sim(2).

3 Problem Definition and Design Objectives

3.1 Interface at a glance

For each registered risk pair, the method receives ordered correspondences, generates a pre-bank homography pool, retains at most eight candidates, and labels the selected bank as EMPTY, SINGLETON, or AMBIGUOUS. Only EMIT+SINGLETON proceeds to support-preserving Sim(2) adaptation. Thus K = 8 bounds the output bank, not generator work or memory. Figure 2 summarizes this release interface.

3.2 Registered risk population

Images are represented in a canonical 640´480 pixel-center frame, with $\bar{x}_t \sim H_{s\circledast t}\bar{x}_s$. A deterministic front end produces

$$\mathbb{C}_{st} = \{(x_i^s, x_i^t, \mathrm{id}_i)\}_{i=1}^{N}. \quad (1)$$

All loop, cross-column, and relocalization pairs are risky. Temporal pairs are included when any frozen pre-method trigger holds: frame gap, inlier ratio, bidirectional error, hull coverage, or Top-2 projective divergence. Method output and reference correctness are never risk features.

3.3 Fixed-capacity mode retention

A shared generator returns candidates $h\hat{\mathbb{I}}\mathbb{P}_{st}$ with support $S(h)$ and a deterministic quality rank. For a fixed capacity K, a bank $\mathbb{B}\acute{\mathbb{I}}\mathbb{P}_{st}$ must satisfy $|\mathbb{B}|£K$. Let $\Pi_\tau(\mathbb{B})$ be its

complete-link partition under the homography distance in Sec. 4.2. We expose two complementary, measurable surrogates:

$$D_\tau(\mathrm{B}) = |\Pi_\tau(\mathrm{B})|, \mathrm{U}(\mathrm{B}) = \left|\bigcup_{h \in \mathrm{B}} \mathrm{S}(h)\right|. \quad (2)$$

D_τ counts retained geometric modes and U counts covered correspondences. The method is a deterministic construction targeting these surrogates; it does not claim a probabilistic optimum or a globally optimal subset. The central empirical question is whether the proposed retained bank preserves an evaluable reference-consistent mode that score-only Top-K truncation loses.

3.4 Controlled quantities and denominators

The primary controlled quantities are reference-consistent-mode presence (CMP), evaluable homography-release coverage (EHRC), and wrong-release counts. Let R be the registered set, E its evaluable subset, A the pre-adapter homography releases, and $W = A \setminus E$ the wrong releases. We define the following:

$$\begin{aligned} \mathrm{EHRC} &= \frac{|A \cap E|}{|E|}, \\ \mathrm{wrong}_{\mathrm{direct}} &= \frac{|W|}{|E|}, \\ \mathrm{wrong}_{\mathrm{conditional}} &= \frac{|W|}{|A \cap E|}. \end{aligned} \quad (3)$$

EHRC and direct wrong incidence use all evaluable registered pairs as denominator; conditional wrong-release rate uses evaluable releases. All 1,620 LinTrack references are evaluable. Homography release is measured before adaptation and is distinct from robot edge-add coverage.

References-consistent-mode presence (CMP) means that at least one retained homography satisfies the pair's reference rule; an empty bank counts as absence. We use "proposed M-V8 retained bank" for bank-level CMP and "MHAR release" for the final pre-adapter output.

4 Multi-Hypothesis Ambiguity Rejection

4.1 Candidate generation and ranking

We use SIFT features and descriptor-ratio filtering [14], followed by a reverse-mutual check. The frozen generator evaluates at most 4,096 homography hypotheses using normalized DLT [15]. Candidates must satisfy support, spatial-coverage, and refit-residual gates before ranking and deduplication; Table I lists numerical settings.

Candidates are ordered ascending by

$$q(h) = \left(-n_h, -r_h, -\min(c_h^s, c_h^t), e_{90,h}, \mathrm{id}_h\right). \quad (4)$$

Here n, r, c, and e90 denote inlier count, ratio, spatial coverage, and residual; candidate ID breaks ties. Homographies within 1 px projective distance are deduplicated before the score-only B-R0 bank retains the first eight.

4.2 Geometric cells and complementary banks

For the fixed 64×48 grid G of pixel centers, the operational distance is

$$\delta(H_a, H_b) = \mathrm{P95}_{x \in G} \| \pi(H_a \bar{x}) - \pi(H_b \bar{x}) \|_2 . \quad (5)$$

Here, $\bar{x}$ is the homogeneous lift of a two-dimensional pixel coordinate; π divides the first two homogeneous components by the third. The operator P95 denotes the 95th percentile over grid points. The operational cell threshold is $\tau = \tau_c = 10\mathrm{px}$.

Candidates are assigned deterministically to complete-link geometric cells in ascending ID order. A candidate joins the first cell whose members all satisfy the distance threshold; otherwise it starts a new cell. This construction is deterministic but not claimed optimal.

From the same pre-bank pool, R0 is score-only, while complementary banks favor geometric-cell representatives or uncovered support. R3 alternates these criteria and stops at K = 8; original support is preserved for adaptation.

4.3 Bank choice, observed ambiguity, and release

Let R_0 be the rank-only bank and r_0 its Top-1 inlier ratio. The frozen binary64 ratio threshold is $\tau_r \gg 0.5436337625178826$.

Its exact hexadecimal floating-point representation is 0x1.16572a1e080eap-1. Define

$$J_\Delta = \begin{cases} 1 - \dfrac{|R_0 \cap R_3|}{|R_0 \cup R_3|}, & |R_0 \cup R_3| > 0, \\ \perp(\mathrm{missing}), & |R_0 \cup R_3| = 0 \end{cases}. \quad (6)$$

The missing value $J_\Delta = \perp$ (BANK_UNION_EMPTY) triggers the first branch of (7) before any threshold comparison.

The frozen method rule selects the first applicable branch, evaluated from top to bottom:

$$(\mathcal{B}^*, d) = \begin{cases} (R_3, \mathrm{WITHHOLD}), & \text{missing/unsupported}, \\ (R_3, \mathrm{EMIT}), & J_\Delta \geqslant 14/15, \\ (R_0, \mathrm{WITHHOLD}), & r_0 < \tau_r, \\ (R_0, \mathrm{EMIT}), & \text{otherwise} \end{cases}. \quad (7)$$

The frozen thresholds and comparison operators are applied exactly as specified. JΔ chooses a branch, while geometric cells define observed ambiguity. Only EMIT+SINGLETON proceeds to adaptation.

4.4 Why homography before Sim(2)

Homography is used only as an intermediate hypothesis language. Approximately planar wall views may contain perspective changes that separate competing modes better than a direct similarity fit, while the deployed back end accepts a single Sim(2) constraint. The selected candidate's original support is therefore refitted rather than truncating homography coefficients.

ISSN 3106-7018 (Print) ISSN 3106-5791 (Online)

The adapter minimizes a Huber objective

$$\min_{s,R,t} \sum_{i \in S} \rho_4 \left(\| x_i^t - (sRx_i^s + t) \|_2 \right). \quad (8)$$

Here, S indexes the selected candidate's original inliers; ρ_4 has a 4-pixel Huber transition in canonical coordinates. The fit estimates a positive scale s, a planar rotation R, and a 2D translation t.

Using the Huber loss [16], our adapter performs iteratively reweighted support refits with similarity updates based on Umeyama's least-squares formulation [17]. The adapter applies frozen support, scale, rotation, and residual gates. These gates test Sim(2) compatibility rather than physical correctness.

Support coordinates and residual thresholds use the canonical 640×480 pixel-center frame; Section 7.3 reports the implementation audit.

4.6 Algorithm

Algorithm 1. Deterministic risk-edge release

1. Materialize the risk-pair and ordered-correspondence registry.
2. Generate, validate, refit, rank, and deduplicate the candidate pool.
3. Construct R0, R1, R2, and R3 with K = 8.
4. Apply the frozen (r0, JΔ) bank-choice and release rule.
5. Classify the selected bank as EMPTY, SINGLETON, or AMBIGUOUS.
6. Unless the result is EMIT + SINGLETON, return SKIP_EDGE; otherwise select the minimum (rank, ID) member.
7. Refit Sim(2) on its original support; return ADD_EDGE only if every gate passes, otherwise ADAPTER_FAILURE.

Table 1. Frozen primary parameters of the MHAR pipeline.

Component	Value
Canonical frame / hypothesis evaluations	640´480 / at most 4096
SIFT / matching	4096; BF-L2; Lowe 0.80; reverse best
homography support	3 px; n^3 12 ; coverage ≥ 0.10
Candidate quality gate	inlier residual P90£4px
Dedup / mode distance	1 px / τ_c = 10px
Bank capacity	K = 8
Release features	$r_0 < \tau_r$; $J_\Delta{}^3$ 14 / 15
Sim(2) Huber / iterations	δ = 4px / 10
Sim(2) validity	$s \in [0.75, 1.25], \lvert\theta\rvert \leqslant 25^\circ$
Sim(2) residual	median ⩽ 2.5px ; P95 ⩽ 6px
Tie-break / registered runtime	rank then SHA-256 ID / one thread, seed 2026080701

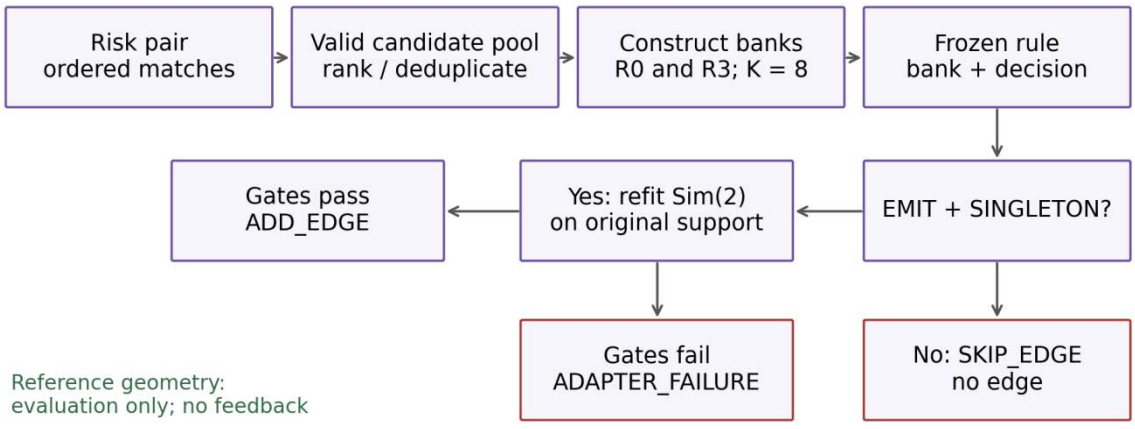


Fig. 2. MHAR release interface. Only EMIT+SINGLETON proceeds to support-preserving Sim(2) adaptation; all other cases add no edge.

5 Robot System Integration

5.1 Existing single-edge boundary

For each registered robot risk pair, ADD_EDGE replaces the production transform, while SKIP_EDGE and ADAPTER_FAILURE remove it. Non-risk temporal edges and downstream optimization, rendering, exposure, and crack processing remain unchanged.

5.2 Risk-pair injection

The method acts only on the pre-materialized risk registry. Pair identity, edge type, trigger evidence, and exclusions are shared across methods; output and evaluator truth cannot change registry membership.

5.3 Unchanged downstream and global boundary

The optimizer accepts at most one Sim(2) constraint per pair. Withholding may fragment the graph, so connectivity and no-output metrics are reported. This local interface is not claimed to outperform global multi-hypothesis inference.

6 Experimental Setup

6.1 Research questions and evidence roles

The LinTrack evaluation asks whether M-V8 improves CMP over B-R0 (Q1) and how MHAR compares with the exact-coverage diagnostic (Q2). Robot evidence addresses integration/connectivity (Q3), while adaptation/resource diagnostics address costs and limitations (Q4).

6.2 Diagnostic questions

6.3 Data roles and controlled evaluator

LinTrack provides planar-tracking sequences and reference geometry [18]. We use a predeclared 1,620-pair registry from the mouse_pad sequence and report finite-registry counts because all pairs come from one sequence.

Primary thresholds were calibrated on a separate 540-pair role before evaluation-reference access. CAL and TEST share no pairs or frame IDs, but this does not establish cross-sequence generalization. Four robot video families were sealed as ROBOT_UNSEEN.

6.4 Retrospective robot diagnostic and evaluator

The retrospective robot study is secondary diagnostic evidence, not a prospective holdout. It contains nine ready scans and four technical failures; a 305-pair reference subset was frozen before annotation, with disagreements adjudicated. The role-isolated analysis uses six ready ROBOT_UNSEEN scans.

Robot edges use a frozen transfer-error rule with no-output and uncertain-reference states retained. Historical outputs are separated from the later unit-consistent adapter refit; corrected downstream optimization and rendering were not replayed.

6.5 Baselines and implementation status

Table II summarizes executable comparison identities and the frozen robot reference protocol.

Table 2(a). Comparison identities and implementation status.

ID	Decision	Role/status
B-PROD	Frozen direct-Sim(2) production estimator	End-to-end system baseline
B-TOP1	Highest-ranked homography, then shared adapter	Single-hypothesis baseline
B-CONF-Fixed	B-TOP1 iff finite $r_0^3 0.4806201550387597$; else withhold	Deployable; 1,105 releases
B-R0	Score-only $K = 8$ bank; singleton-only release	Bank-retention comparator
M-V8-NR	Proposed bank; force one representative if nonempty	Withholding ablation
MHAR (M-V8)	Proposed bank choice and singleton-only release	Full method
B-CYCLE	Gate B-TOP1 against a non-risk alternate path	Required robot-only global comparator
Oracle	Reference-consistent mode available in the pre-bank pool	Evaluator-only ceiling, not a method

Table 2(b). Sealed robot input and frozen reference protocol.

Candidates (ready/fail)	Pairs (all/risk)	Reference subset	Primary/secondary/adjud.	Required unique landmark-visible frames	Current lifecycle
13 (9/4)	2,839/2,838	305	305/305/305	≥180 total; ≥20/scan	Final audit PASS

B-CYCLE is a robot-only path-consistency comparator. Graph-Cut RANSAC failed the determinism gate, while Progressive-X and CONSAC were not integrated; they are related-work context only.

B-CONF-Match@338 is an evaluator-only exact-coverage diagnostic, not a deployable threshold; it ranks the same

ISSN 3106-7018 (Print) ISSN 3106-5791 (Online)

frozen confidence score to match 338 releases.

6.6 Metrics, comparison scope, and shared controls

Q1 compares bank-level CMP for M-V8 and B-R0. Q2 compares MHAR with M-V8-NR for withholding and B-CONF-Match@338 for matched-coverage release quality. Methods share frames, correspondences, registry membership, solver budget, evaluator, and downstream initialization.

7 Results

7.1 Q1: Bank-level CMP gain; mode selection remains untested

On LinTrack, M-V8 contains a reference-consistent homography in 1,183/1,620 pairs (73.02%), versus 881/1,620 (54.38%) for B-R0 and 849/1,620 (52.41%) for Top-1. The census is 881 retained by both, 302 by M-V8 only, none by B-R0 only, and 437 by neither. This is a retention gain, not successful mode selection.

Periodic-texture diagnostic

In the 100-pair synthetic diagnostic, all 80 alias-positive periodic pairs have empty valid pools. Among 20 nonperiodic controls, reference-consistent modes reach the valid pool in 5 PRIMARY and 13 SUCCESSOR cases. Thus periodic failure occurs before bank selection.

7.2 Q2: More wrong releases at matched coverage

Figure 3 summarizes the controlled release counts and wrong releases; Table 3 provides the complete census.

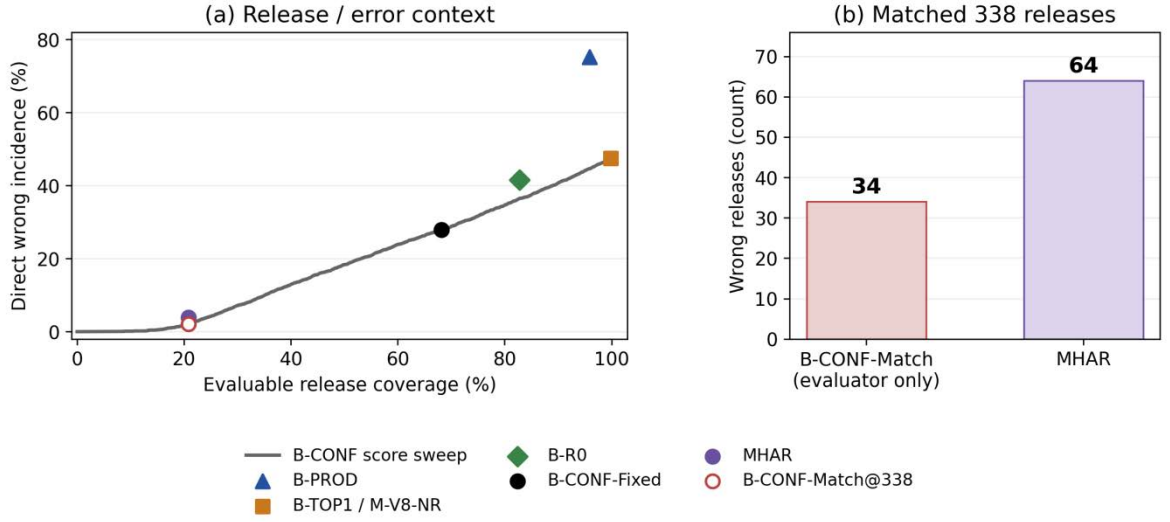


Fig. 3. Controlled release/error results. At 338 releases, B-CONF-Match@338 has 34 wrong homographies and MHAR has 64; the matched diagnostic is evaluator-only.

Table 3. Complete 1,620-pair controlled census. CMP is retained-bank reference-consistent-mode presence; A-C/A-W/W-E are accepted-correct, accepted-wrong, and withheld-evaluable counts; EHRC is pre-adapter release coverage.

Method / bank	CMP ↑	D_τ ↑	A-C/A-W/W-E	EHRC ↑	Wrong / released ↓
B-PROD	–	–	335/1218/67	1553/1620	1218/1553
B-TOP1	849/1620	0.999/1	849/769/2	1618/1620	769/1618
B-CONF-Fixed	849/1620	0.999/1	653/452/515	1105/1620	452/1105
B-CONF-Match@338†	849/1620	0.999/1	304/34/1282	338/1620	34/338
B-R0	881/1620	1.278/3	669/672/279	1341/1620	672/1341
M-V8-NR	1183/1620	2.817/6	849/769/2	1618/1620	769/1618
MHAR (M-V8)	1183/1620	2.817/6	274/64/1282	338/1620	64/338
Oracle‡	1348/1620	–	–	–	–

Notes: Dτ reports mean/P95 retained-cell count per pair. Wrong/released is |W|/|A ∩ E| in (3). ↑/↓ indicate the preferred direction; higher CMP or Dτ does not imply better release quality. † Evaluator-only exact-coverage diag-

nostic. ‡ Evaluator-only ceiling from the shared pre-bank pool.

At the same 338 evaluable releases, MHAR emits 274 correct and 64 wrong homographies, versus 304 correct and 34 wrong for B-CONF-Match@338; this is not a comparison of two deployable rules. Sixty of MHAR's 64 errors occur in the 187-pair R0=SINGLETON, R3=AMBIGUOUS, EMIT_R0 subgroup. R3 exposes ambiguity, but the frozen rule selects R0 and releases Top-1. Removing this subgroup post hoc would also remove 127 correct releases, so it is used only for failure localization.

7.3 Q3: Historical robot outcome and unit-consistent adapter audit

As secondary retrospective evidence, the role-isolated ROBOT_UNSEEN reaggregation contains six ready scans and four technical failures. All six historical ready outputs fall below the frozen 90% coverage criterion and are fragmented. A coordinate audit identifies an approximately 3.6-fold unit mismatch: support points were mapped to raw-undistorted cropped coordinates while Huber and residual thresholds remained in canonical coordinates. Unit-consistent refitting changes adapter-passing edge availability, but corrected edge sets were not propagated through optimization or rendering; no corrected downstream improvement is claimed. Figure 4 summarizes the historical nine-ready-scan results in Table IV(b).

Table 4(a). Historical six-scan ROBOT_UNSEEN results under the primary weight policy. Ratios are computed from the displayed FC values.

Scan	B-PROD FC	MHAR FC	MHAR/B-PROD	Wrong edges T1/MHAR
Scan 1	0.2028	0.0420	0.2071	9/4
Scan 2	0.2484	0.0392	0.1578	13/6
Scan 3	0.1171	0.0341	0.2912	15/7
Scan 4	0.1513	0.0336	0.2221	7/3
Scan 5	0.1304	0.0117	0.0897	20/7
Scan 6	0.0690	0.0219	0.3174	22/9

All six ready scans fail the 90% coverage criterion and end with disconnected injected graphs. The historical wrong-edge-count check does not test edge identity.

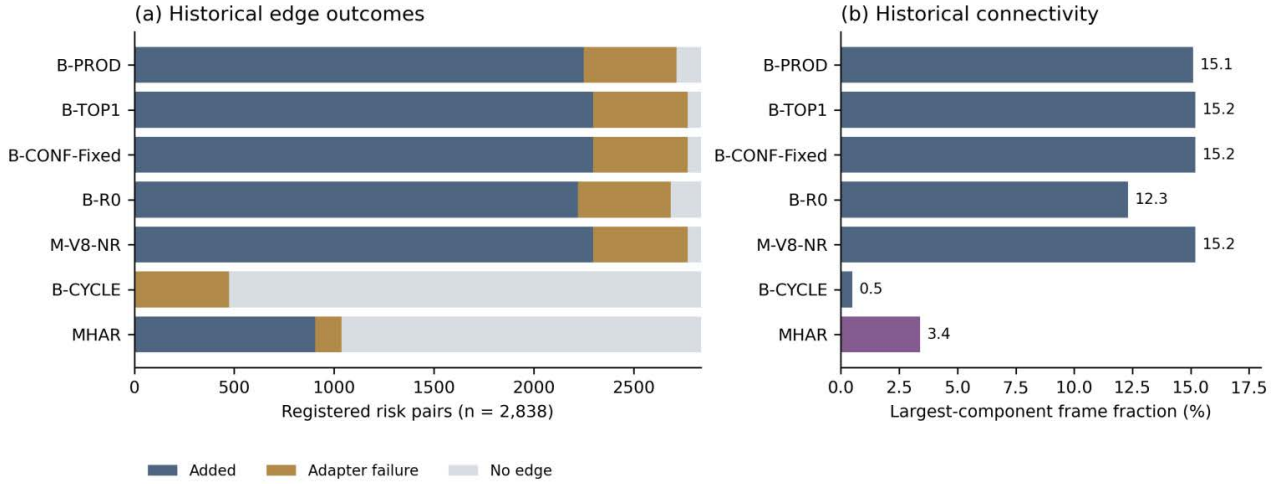


Fig. 4. Historical nine-ready-scan diagnostic from Table IV(b): (a) added, adapter-failed, and no-edge risk pairs; (b) largest-component frame fraction. Outputs are historical, not corrected downstream results.

Table 4(b). Historical nine-ready-scan diagnostic. A/AF/NE denote add/adapter-failure/no-edge counts; Ref. C/W denotes correct/wrong added edges on the reference subset.

Method	A/AF/NE	Ref. C/W	CC ↓	FC ↑	F/N	Time (h) ↓
B-PROD	2249/465/124	115/106	47	0.151	9/0	4.92

ISSN 3106-7018 (Print) ISSN 3106-5791 (Online)

Method	A/AF/NE	Ref. C/W	CC ↓	FC ↑	F/N	Time (h) ↓
B-TOP1	2297/474/67	117/109	50	0.152	9/0	4.90
B-CONF-Fixed	2297/473/68	117/109	50	0.152	9/0	4.90
B-R0	2222/463/153	114/103	59	0.123	9/0	5.00
M-V8-NR	2297/474/67	117/109	50	0.152	9/0	5.00
B-CYCLE	0/474/2364	0/0	208	0.005	1/8	4.90
MHAR	906/132/1800	53/43	152	0.034	9/0	5.00

7.4 Q4: Adaptation, support sensitivity, and resource costs

Of 338 MHAR homography releases, 288 pass the adapter. The 61/288 source-H wrong fraction does not measure the physical error of the refitted Sim(2) output. Alternative distance definitions change some bank states, confirming partition sensitivity. The maximum pair time is 427.415 s and peak RSS is 158.441 MiB, so K = 8 does not bound computation or memory.

7.5 Failure taxonomy and crack boundary

Failures include pool absence, cap loss, competing cells, withholding, adapter failure, fragmentation, and non-planarity/parallax. Crack projection remains qualitative because the real-mask coordinate and denominator audit has not passed.

8 Discussion and Limitations

Q1 — Bank-level CMP is not selection. M-V8 increases reference-consistent-mode availability over B-R0, but every released candidate remains Top-1. Periodic alias-positive pairs fail before a valid pool forms, so successful periodic mode selection is not established.

Q2 — Release quality is the main controlled failure. At matched coverage, MHAR has more wrong releases than B-CONF-Match@338, with most errors where R3 is ambiguous but the frozen rule selects singleton R0. This localizes failure but does not justify a post-hoc rule change.

Q3 — The robot comparison is implementation-confounded. Historical adapter thresholds used inconsistent coordinate units, and corrected edge sets were not replayed downstream. The robot evidence therefore diagnoses integration and connectivity costs rather than causal downstream effects of MHAR.

Q4 — Adapter compatibility is not correctness; fixed bank capacity is not fixed computational cost, and withholding trades release risk against connectivity. Retention, release, adaptation, and downstream utility therefore require separate evaluation.

9 Conclusion

Within the separated CAL/TEST LinTrack registry, the proposed M-V8 bank increases reference-consistent-mode presence relative to same-capacity B-R0, but the release rule does not use those additional modes to replace Top-1. At matched release coverage, MHAR produces more wrong homographies than the evaluator-only B-CONF-Match@338 diagnostic, with most errors concentrated in the branch that selects singleton R0 despite ambiguity in R3. Alias-positive periodic cases fail before valid candidate pools are formed, so periodic mode selection remains untested. Historical robot outputs are further limited by a coordinate-unit mismatch, and corrected edge sets were not replayed downstream. The evidence therefore separates retention, release, adaptation, and connectivity as distinct problems. Future work should test a preregistered selection rule on periodic cases with valid competing modes and repeat robot evaluation with a unit-consistent adapter and complete downstream replay.

References

[1] S. Sano, D. Takaki, A. Ishida, and T. Ishida, “Image mosaicking and localization using a camera mounted on a hanging-type wall climbing robot,” J. Robot. Mechatron., vol. 33, no. 6, pp. 1373–1383, 2021, doi: 10.20965/jrm.2021.p1373.

[2] L. Yang et al., “Automated wall-climbing robot for concrete construction inspection,” J. Field Robot., vol. 40, no. 1, pp. 110–129, 2023, doi: 10.1002/rob.22119.

[3] M. Pellikka and V. Lahtinen, “A robust method for image stitching,” Pattern Anal. Appl., vol. 24, no. 4, pp. 1847–1858, 2021, doi: 10.1007/s10044-021-01005-8.

[4] M. A. Fischler and R. C. Bolles, “Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography,” Commun. ACM, vol. 24, no. 6, pp. 381–395, 1981, doi: 10.1145/358669.358692.

[5] D. Barath and J. Matas, “Graph-Cut RANSAC,” in Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR), 2018, pp. 6733–6741.

[6] D. Barath and J. Matas, “Progressive-X: Efficient, anytime, multi-model fitting algorithm,” in Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV), 2019, pp. 3780–3788.

[7] F. Kluger, E. Brachmann, H. Ackermann, C. Rother, M. Y. Yang, and B. Rosenhahn, “CONSAC: Robust multi-model fitting by conditional sample consensus,” in Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR), 2020, pp. 4634–4643.

[8] C. Herrmann et al., “Robust image stitching with multiple registrations,” in Proc. Eur. Conf. Comput. Vis. (ECCV),

2018, pp. 53–69, doi: 10.1007/978-3-030-01216-8_4.

[9] N. Sünderhauf and P. Protzel, "Switchable constraints for robust pose graph SLAM," in Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst. (IROS), 2012, pp. 1879–1884, doi: 10.1109/IROS.2012.6385590.

[10] P. Agarwal, G. D. Tipaldi, L. Spinello, C. Stachniss, and W. Burgard, "Robust map optimization using dynamic covariance scaling," in Proc. IEEE Int. Conf. Robot. Autom. (ICRA), 2013, pp. 62–69, doi: 10.1109/ICRA.2013.6630557.

[11] J. G. Mangelson, D. Dominic, R. M. Eustice, and R. Vasudevan, "Pairwise consistent measurement set maximization for robust multi-robot map merging," in Proc. IEEE Int. Conf. Robot. Autom. (ICRA), 2018, pp. 2916–2923, doi: 10.1109/ICRA.2018.8460217.

[12] H. Yang, P. Antonante, V. Tzoumas, and L. Carlone, "Graduated non-convexity for robust spatial perception: From non-minimal solvers to global outlier rejection," IEEE Robot. Autom. Lett., vol. 5, no. 2, pp. 1127–1134, 2020, doi: 10.1109/LRA.2020.2965893.

[13] Y. Geifman and R. El-Yaniv, "SelectiveNet: A deep neural network with an integrated reject option," in Proc. 36th Int. Conf. Mach. Learn. (ICML), ser. Proc. Mach. Learn. Res., vol. 97, pp. 2151–2159, 2019.

[14] D. G. Lowe, "Distinctive image features from scale-invariant keypoints," Int. J. Comput. Vis., vol. 60, no. 2, pp. 91–110, 2004, doi: 10.1023/B:VISI.0000029664.99615.94.

[15] R. Hartley and A. Zisserman, Multiple View Geometry in Computer Vision, 2nd ed. Cambridge, U.K.: Cambridge Univ. Press, 2004.

[16] P. J. Huber, "Robust estimation of a location parameter," Ann. Math. Stat., vol. 35, no. 1, pp. 73–101, 1964, doi: 10.1214/aoms/1177703732.

[17] S. Umeyama, "Least-squares estimation of transformation parameters between two point patterns," IEEE Trans. Pattern Anal. Mach. Intell., vol. 13, no. 4, pp. 376–380, 1991, doi: 10.1109/34.88573.

[18] K. Zimmermann, J. Matas, and T. Svoboda, "Tracking by an optimal sequence of linear predictors," IEEE Trans. Pattern Anal. Mach. Intell., vol. 31, no. 4, pp. 677–692, 2009, doi: 10.1109/TPAMI.2008.119.