capstract:2611.00040v1 [stat.ML] 20 Nov 2026

Transfer of Restart Policies under Ambiguous Failure Feedback

Mina Darven-Kim*
Department of Statistical Learning
Haldenmere Institute of Technology
m.darven-kim@haldenmere.example

Lucien Ortega-Vey
Laboratory for Geometric Inference
Valdorin Polytechnic

Adeola T. Mbadi-Lorne
Centre for Measurement Systems
Ternwell College of Engineering

Abstract

A restart after a failed local search can preserve the current hypothesis or discard it. The failure alone often provides weak evidence about which operation to perform. We study transfer of a sequential restart policy between tasks that share a binary latent state but have different observation units. A two-symbol rejection channel, estimated from 960 controlled insertion trials on keyed interfaces, initializes the policy for a held-out point-cloud registration benchmark. The transferred policy recovers 2,514 of 2,880 poses within four attempts, compared with 2,408 for strict alternation and 2,546 for a policy fitted on labeled registrations. Its paired improvement over alternation is 3.7 percentage points (95% object-bootstrap interval 2.2 to 5.2). Collapsing the rejection symbols recovers the strict-alternation decision sequence. The near-miss symbol reverses the direction of the posterior update after a failed attempt in both domains. Neither scale matching nor target pose labels are needed for this reversal.

1 Introduction

When two initializations are available, there are two ways to rerun a failed search. One can start somewhere else within the current initialization basin, or one can change basins. The distinction matters when the first initialization chooses a persistent discrete hypothesis. A poor local offset is refreshed on the next run. A wrong discrete hypothesis is refreshed only when explicitly replaced. Alternation is the standard inexpensive response when the caller cannot tell the two failures apart.

The information for a better decision is commonly present in the unsuccessful run. A fit that gets close to verification can be evidence *for* its initialization, even though a fit that stops early is evidence against it. An aggregate failure count discards the distinction. Section 3 gives a policy that uses one additional bit per rejected run and a single probability carried across runs. We ask whether the probabilities for this model have to be learned from labeled outputs of the solver. Such labels ordinarily require obtaining the pose by another means, which can be more costly than running the solver again.

Our source of labeled failures is the USB Type-A receptacle. In an occluded horizontal mount, plug orientation is binary and stationary, and insertion can fail in the admissible orientation because of a small lateral offset. Some plugs have entered the shell before they stop. Others have reached the face of the receptacle and stopped there. We learn the distinction from measured travel and apply it, without fitting to target pose labels, to a registration solver offered two headings separated by a

∗Corresponding author: m.darven-kim@haldenmere.example

Preprint.

half-turn. On both tasks a retry can freshen the local offset without changing the heading. The transfer concerns the controller's decision at a rejected attempt. The plugs supply no observations about the point sets.

The two rejection codes yield opposite likelihood ratios in the source. In the target they do so again, with different magnitudes. Section 5 tests whether this is sufficient at a fixed attempt budget and how much is lost by forgetting earlier failures. We also measure wrong certificates. On partially observed surfaces it is possible to produce a plausible fit to the back of an object even when the front heading was supplied to the benchmark.

2 Related work

An independent restart of a Las Vegas procedure samples another runtime from the same distribution, and the cutoff sequence controls the expected total runtime (Luby et al., 1993). Context-sensitive restart policies can instead use evidence collected during a search to estimate its remaining runtime (Kautz et al., 2002). We fix the cutoff and choose a candidate. The belief update is a sequential likelihood ratio of the form used by Wald (1945), with actions determining which hypothesis has which emission distribution. The controller uses a finite-horizon value function (Sutton and Barto, 2018). Here the initial hypothesis has no learned prior, and the only learned parameters are conditional frequencies of coarse outcomes.

Rigid registration with known correspondences admits a least-squares solution (Horn, 1987; Umeyama, 1991). The uncertain correspondence problem is substantially different. Consensus methods sample correspondences (Fischler and Bolles, 1981). Local methods update correspondences and motion in alternation (Besl and McKay, 1992; Chen and Medioni, 1991). Point-to-plane residuals and subsampling can improve the behavior of individual iterations (Rusinkiewicz and Levoy, 2001), and trimming can reduce the influence of nonoverlap (Chetverikov et al., 2002). Feature-based global registration (Zhou et al., 2016) and learned registration systems (Aoki et al., 2019; Wang and Solomon, 2019) change the initialization or optimization itself. Our principal comparisons use the same trimmed solver; the controller only decides which of two coarse headings to pass to it.

Bench studies of keyed connectors distinguish shell contact from contact with the internal insert (Salen-Kuroda and Essel-Neri, 2020). Meier-Haddou and Irete (2023) found that visually masking the receptacle changed the choice of first heading in a hand insertion task without measurably changing the characteristic travel at either contact. We use similar masking and a motorized stage because a person told to try again may change both the lateral offset and the angle in ways that depend on the preceding contact. On the registration side our verifier follows the separation of fitting points from witness points used by Velado-Fenn and Waseem-Dorne (2025).

3 Model

Let $H \in \{0, 1\}$ be the persistent correct heading and $A_t \in \{0, 1\}$ the heading chosen for attempt t. On each attempt we draw a new local perturbation. The observed symbol Y_t is an acceptance a, a shallow rejection s, or a deep rejection d. The terms shallow and deep refer to extent of progress toward the verification threshold. We suppress the question of a wrong acceptance until Section 4.3. The working emission law $p_m(y)$ depends on whether $m = \mathbf{1}\{A_t = H\}$. Conditional on A_t, H, it is taken to be independent of previous attempts. We fit this law under controlled offsets in Section 4.1 rather than estimating it from the failures produced by a previously deployed policy.

Let ℓ_t be the log posterior odds for $H = 0$ before attempt t. The first heading is fixed at zero by convention and H is balanced in evaluation. A rejection changes the odds by

$$\ell_{t+1} = \ell_t + (1 - 2A_t) \log \frac{p_1(Y_t)}{p_0(Y_t)}. \tag{1}$$

The probability of a correct heading, $q_A(\ell)$, is $\operatorname{logit}^{-1}((1 - 2A)\ell)$. Write $w_y(\ell, A) = q_A(\ell) p_1(y) + (1 - q_A(\ell)) p_0(y)$. For b attempts left, Eq. (2) gives the probability of eventual acceptance under the fitted model, starting from $V_0 = 0$.

$$V_b(\ell) = \max_{A \in \{0,1\}} \left\{ w_a(\ell, A) + \sum_{y \in \{s,d\}} w_y(\ell, A) V_{b-1}\left(\ell + (1 - 2A) \log \frac{p_1(y)}{p_0(y)} \right) \right\}. \tag{2}$$

We resolve ties in favor of changing the previous heading after rejection. At the first attempt the two headings are exchangeable; every local method receives the same initial candidate. The recursion uses the actual number of attempts remaining. In the success-by-attempt curves, decisions are made for a four-attempt horizon throughout; the value at two attempts is a prefix of that run and need not equal the value under a two-attempt budget.

On the final attempt with $p_1(a) > p_0(a)$, Eq. (2) selects the more probable heading. On an earlier attempt the informative outcome can also have value. In particular, accumulation matters when two weak pieces of contrary evidence fail to erase a strong observation. With the source estimates, one shallow rejection and two deep rejections at the same heading still favor the other heading: the combined log likelihood ratio is approximately $-2.34 + 2(0.79) = -0.76$. The policy that retains on the last deep rejection would disregard that history. This is a statement about hypothetical evidence; the policy need not visit all three measurements at a single heading.

3.1 A transfer condition

We distinguish disagreement on the code probabilities from disagreement on actions. For the last attempt and a uniform prior after one rejection, the policy stays if $p_1(y) > p_0(y)$. The magnitudes can differ between source and target without changing this decision. Later decisions depend on the sum in Eq. (1) and on the remaining horizon. Even identical signs no longer suffice.

There is a simple, more stringent condition for full policies. Suppose a source and a target start with the same law of H and, conditional on any common history and its next action, their outcome laws differ by at most δ in total variation. Here the outcome alphabet can distinguish correct and incorrect acceptance for the evaluator, even though both terminate play for the policy. Under the same reward of one for correct acceptance by step B, the loss from running an optimal source policy in the target obeys

$$0 \le V_T^\star - V_T(\pi_S^\star) \le 2\{1 - (1 - \delta)^B\}. \tag{3}$$

The proof in Appendix A requires a bound on whole conditional laws, including the law in the target after a run of failures. We use Eq. (3) to state what would suffice for transfer. The measured first-rejection ratios cannot establish such a bound. Repeated offsets on a mesh can give conditionally dependent failures even when the perturbations themselves are independent.

4 Experimental design

4.1 Connector trials

We used 48 plug and receptacle assemblies, each consisting of a plug with its own unpowered receptacle. Twenty-four assemblies supplied training data and 24 were reserved as an audit set. The split was by assembly and was blocked by one of six measured shell geometries. One assembly contributed 40 insertions, 20 at each heading in shuffled order. The correct heading was read from a fiducial on the back of the receptacle and joined to the trace only after processing. Front masking hid the keyway during positioning. It had no contact with the plug.

The insertion stage began each trial fully withdrawn. A lateral offset and small pitch perturbation were drawn from a fixed zero-centered distribution independently of the heading. The stage halted at the displacement limit or at its instrument force limit, whichever came first. No electrical connection was measured. Mechanical closure, read by a separate displacement switch, yielded acceptance. For other trials we divided greatest forward travel by the geometric stroke and used 0.35 as the code boundary. This was the upper nominal extent of contact with the mouth of the shell in the fixture specification. The same cutoff was applied to every assembly. No cutoffs were optimized on the held-out records. A trace was kept when the stage met its prescribed start and reset positions, and all 1,920 attempted insertions met these conditions.

The channel is estimated from labeled counts in Table 1. In the model we retain the structural zero $p_0(a) = 0$ and add one-half to each other cell before normalizing by row. The structural zero is an approximation for the target, where verification is imperfect. We chose to transfer the mechanical constraint unmodified and report evidence against it in Section 5.

Table 1: Source emission counts by orientation and outcome. Counts within every row sum to 480. The orientation label describes the commanded heading, not the calculated insertion depth.

Assemblies	Heading	Accepted	Shallow	Deep
Fit, 24	Correct	286	39	155
Fit, 24	Wrong	0	410	70
Audit, 24	Correct	282	46	152
Audit, 24	Wrong	0	399	81

4.2 Registration episodes

We constructed episodes from four ModelNet40 categories (Wu et al., 2015). We drew 30 meshes per category from the official training split for the target-supervised comparison and 60 per category from the test split for evaluation, using chairs, tables, dressers, and bookshelves. Meshes with zero surface area after normalizing to unit bounding-box diagonal were ineligible. Twelve episodes per mesh yielded 1,440 labeled development episodes and 2,880 test episodes. Six of each set of 12 had their correct heading at candidate zero and six at candidate one. This balancing fixes the initial posterior without learning a prior and prevents a shape category from being informative about the index.

Two clipped surface regions were sampled per mesh. The extent of common visible area, relative to the smaller of the two regions, lay in $[0.35, 0.65]$. Source and target points were sampled independently from triangular faces by area. Each input cloud had 1,024 points. A third set of 384 witness points per view was independently sampled from the respective visible regions and withheld from the local fit. The source view was rotated about the supplied vertical axis and translated. The two candidates used on input were a coarse heading within eight degrees of the planted angle and the heading obtained by adding 180 degrees, presented in the balanced order described above. The controller received integer candidate identifiers and rejection codes. It did not receive point coordinates, the witness residual, or the coarse heading itself.

An attempt used at most 20 point-to-plane ICP iterations. We discarded the 30% largest absolute correspondence residuals at each iteration and constrained yaw to remain within 45 degrees of the candidate heading. The initial translation aligned the two observed centroids and added an isotropic Gaussian perturbation with standard deviation 0.04 bounding-box diagonals per coordinate. Each trial had a preassigned perturbation for each attempt index, shared across all policies even when they chose different headings. This prevents a rule from gaining systematic access to favorable restart seeds. Trimming the residual is an instance of the procedure described by Chetverikov et al. (2002). The nominal overlap was used to stratify the report, not to select a trimming fraction per example.

4.3 Certificate and outcome

The verifier used the trimmed symmetric nearest-neighbor root-mean-square distance r, in units of the bounding-box diagonal, between the two witness sets after applying the proposed motion. It kept the smallest 60% of distances in each direction, averaged the pooled squares, and took their square root. No witness distances contributed to the optimizer’s correspondence updates. The coarse output passed to the controller was

$$Y = \begin{cases} a & r \le 0.025, \\ d & 0.025 < r \le 0.050, \\ s & r > 0.050. \end{cases} \tag{4}$$

The acceptance threshold in Eq. (4) was inherited from the witness protocol (Velado-Fenn and Waseem-Dorne, 2025). We set the deep cutoff at twice that tolerance before inspecting target traces. The factor of two amounts to asking whether the rejected fit was inside the next larger error band. No conversion from millimeters of plug travel to distance in a point set was used.

A success required a certificate and a final pose error at most five degrees in rotation and 0.05 diagonals in translation, computed from the withheld transform. Any certificate stopped search, including a wrong certificate. We count a stopped run with an incorrect pose as a failure. We also count exhaustion of the budget as a failure, even if some earlier rejected transform met the pose tolerances. The unit of analysis is a full episode; the unit of resampling for target intervals is the mesh, with all episodes and methods from a sampled mesh kept together.

4.4 Comparisons and tests

We compare a fixed heading, independent redraw after each rejection, and strict alternation. The latest-code rule stays after a deep rejection and switches after a shallow one. It uses the direction of the observed effect without learning its size or tracking multiple outcomes. The collapsed-code model sums each row’s two rejection frequencies and solves Eq. (2) with only that symbol. For the target-fitted policy we precomputed both candidate outcomes at each of the four attempts on the 1,440 development episodes. We then fit a decision tree by backward induction on these complete replay records. At each attainable history and heading we chose the heading with the larger empirical count of eventual correct acceptances, using descendant decisions already fit. The root heading was fixed at zero. At nodes visited by fewer than 20 development episodes we used the latest-code rule for the remaining attempts. This comparison has the same inputs at evaluation and can learn from the exact history of attempts that led to a rejection.

Separately we run Fast Global Registration (Zhou et al., 2016) once per episode using the fit clouds, with witness checking and the same pose criteria. This diagnostic receives no candidate set. Its success rate is included to put the restricted local task in context. It is omitted from the local attempt curves since its correspondence features, failure behavior, and per-attempt time are different.

The main estimand is the paired difference in four-attempt success from source transfer versus strict alternation. We take 2,000 stratified percentile bootstrap resamples (Efron and Tibshirani, 1993), sampling meshes with replacement within each of the four categories. We keep the plug fit fixed. The resulting interval describes target sampling conditional on the trained policy. For the bench audit we resample assemblies within shell geometry. The first-rejection plot shows pointwise intervals of 1.96 bootstrap standard errors on either side of each test and audit estimate.

5 Results

5.1 First-rejection evidence

Under the training channel, the unsmoothed probability of being in the correct heading, given a shallow first rejection, is $39/449 = 0.087$. For a deep first rejection it is $155/225 = 0.689$. Those fractions use both headings in the denominators. Conditioning on rejection inside each heading before comparing them would erase the difference in acceptance rates. The policy uses the joint probability of each code and rejection, that is, the corresponding cell in a three-outcome row of Table 1.

The audit results in Figure 1 show the reversal in both source and target. On the 960 audit insertions the two probabilities are 0.103 and 0.652. On registration there were 875 first-attempt certificates, of which 870 met the pose-error limits. Among the 1,328 shallow and 677 deep first rejections, the candidate was correct in 138 and 432 cases, respectively. The agreement is more pronounced in direction than in magnitude. Treating 0.689 as a calibrated target probability for the deep code would overstate the frequency by about five points. The natural reference here is a conditional empirical frequency in the sense of Dawid (1982). As a proper-score check (Gneiting and Raftery, 2007), the two smoothed source probabilities give a Brier score of 0.141 on these 2,005 rejected first attempts. Using the marginal correct-heading frequency among rejected attempts, estimated on that same set and therefore given an optimistic reference, gives 0.203. These scores compare information in the codes. They do not demonstrate calibration of the evolving posterior.

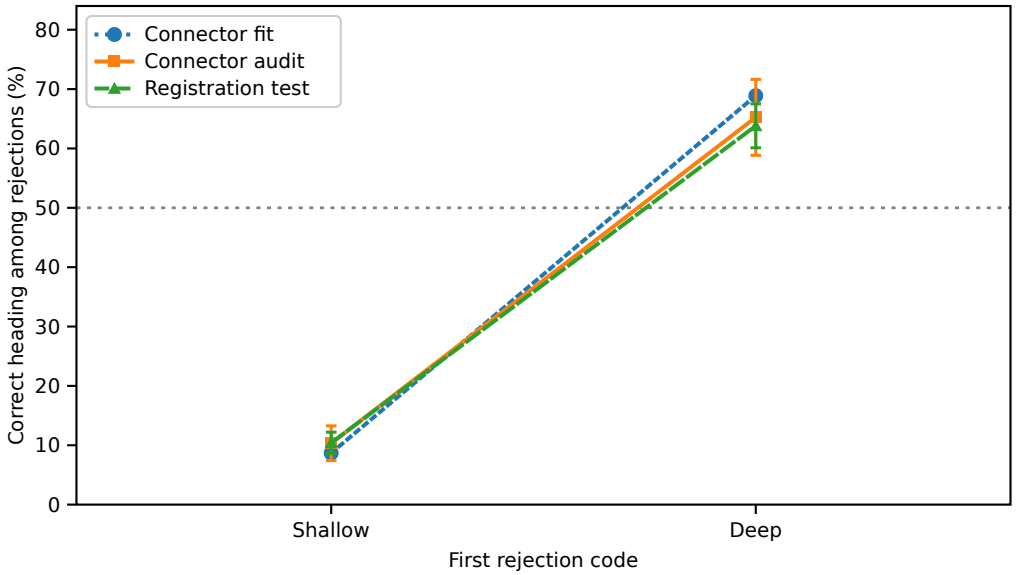


Figure 1: Correct-heading frequencies conditional on the first rejection. Bars on the audit and registration points are pointwise 95% bootstrap intervals. Training values use counts without smoothing.

5.2 Recovery at a fixed budget

The four-attempt policy fit to the plug records recovers 2,514 poses. Strict alternation recovers 2,408; the resulting difference is 106 poses, or 3.7 percentage points (95% interval 2.2 to 5.2). The target-fitted policy recovers 2,546. Fast Global Registration succeeds for 2,480 episodes (86.1%). Its median wall time is 221 ms per episode on the same CPU, compared with 164 ms for the transferred policy. Those measurements include witness checking. A comparison at equal wall time would require varying the local and global solver implementations and is outside this benchmark.

Table 2 reports cumulative successful episodes. Figure 2 shows the timing of the differences. The target-fitted policy and the source policy both gain by the second attempt. The latest-code rule agrees with the connector fit through attempt three. At the fourth attempt there is one rejection history where the rules differ. A shallow rejection at the first heading, followed by a deep then a shallow rejection at the second heading, leads the latest-code rule back to the first heading. The source policy stays with the second. The two shallow rejections oppose and cancel one another in Eq. (1). The earlier partial entry at the second heading determines the last attempt.

Table 2: Cumulative certified poses within tolerance by attempt, out of 2,880. All rows use the same initial candidate, fitter, and verifier.

Restart rule	1	2	3	4	4 (%)
Fixed heading	870	1,205	1,330	1,376	47.8
Independent redraw	870	1,475	1,893	2,191	76.1
Strict alternation	870	1,737	2,078	2,408	83.6
Latest code	870	1,826	2,235	2,497	86.7
Collapsed code	870	1,737	2,078	2,408	83.6
Connector fit	870	1,826	2,235	2,514	87.3
Target fit	870	1,826	2,251	2,546	**88.4**

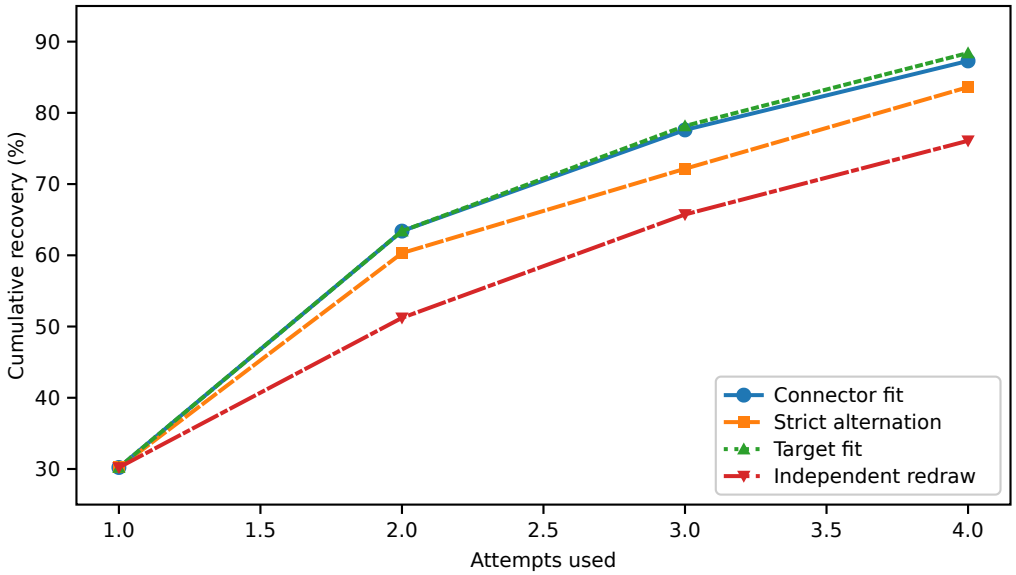


Figure 2: Certified pose recovery on the same 2,880 episodes. Curves use a fixed four-attempt horizon. The first attempt is shared.

Incorrect certificates account for 27 of the 366 unrecovered cases under source transfer. The other 339 ended with four rejections. For strict alternation, 22 ended with an incorrect certificate and 450 with four rejections. The transferred policy is more willing to remain in a well-fitting basin. If the basin is wrong, doing so provides a new opportunity for the same surface to give a wrong certificate.

5.3 Sensitivity to shared surface area

The overlap breakdown is in Table 3 and Figure 3. We ranked the nominal fractions of common area, with ties broken by episode identifier. The largest absolute difference from alternation is in the lowest-overlap quarter. On the failure scale the policy removes 42 of 190 failures in the lowest quartile and 17 of 57 in the highest. The latter is the larger proportional reduction.

Table 3: Successful recoveries at four attempts by overlap quartile. Each quartile has 720 episodes.

Overlap quartile	Alternation	Connector fit	Target fit
Lowest	530	572	585
Second	585	612	621
Third	630	650	655
Highest	663	680	685

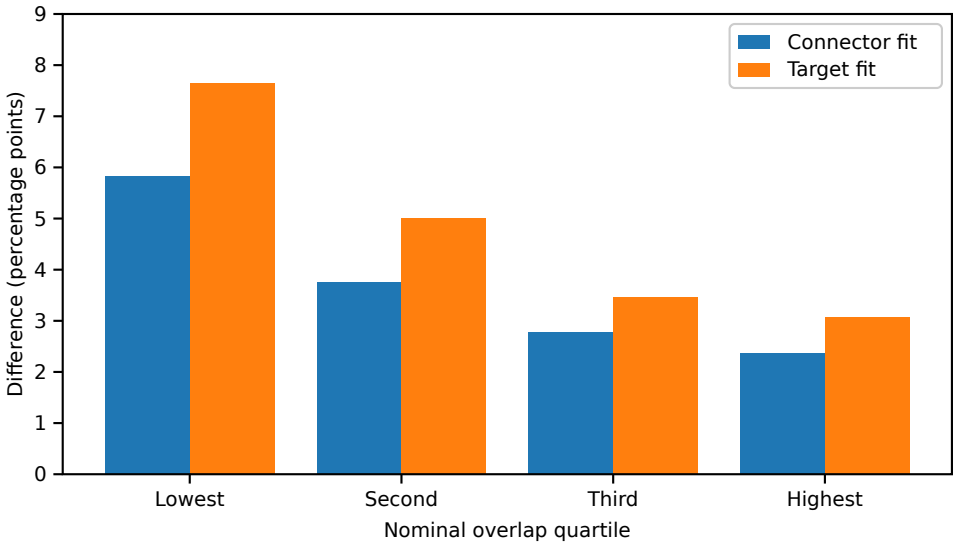


Figure 3: Pairwise differences in successful recovery by nominal overlap quartile. Differences are in percentage points over strict alternation.

Because the surface is shared across attempts, offsets within an episode can repeatedly lead to essentially the same correspondence set. On development traces with both headings run at all four offsets, we examined pairs of attempts at the correct heading. The deep-rejection fraction on the second attempt was 0.39 after a deep first rejection and 0.28 after a shallow first rejection. Both figures use all second outcomes as the denominator. Conditioning on the first symbol changes the population of surfaces that the second attempt sees, even with full reset of the optimizer.

6 Discussion

The source device is a binary keyed object with a physical endpoint. It is easy to assign the correct heading in that setting. The analogous label in partial registration is expensive: two motions can bring the shared patch into correspondence and differ by a half-turn on the rest of the object. Our policy asks the witness to distinguish two kinds of rejection and uses a separately counted source for the relative weight of each. The gain over the latest-code rule is 17 poses. The 106 gained over strict alternation mainly reflect access to the extra bit.

The benchmark relies on a procedure for obtaining two useful candidates. An exactly half-turn-symmetric object can make the task underdetermined: the benchmark’s planted heading then expresses a labeling convention. The pose metrics score that convention. All success rates are conditional on the chosen categories and moderate visible overlap. The verifier reuses the same witness samples on every attempt, so adapting to its rejection codes can select a fit that benefits from chance properties of those samples. The wall times omit the cost of obtaining the coarse headings.

Each mechanical assembly supplied 40 traces in a single session. The audit tests whether the channel retains its sign on other assemblies from the same six geometries. A comparison with natural insertions would need to separate fixture effects from a user’s tendency to rotate and translate the plug together. The structural zero for wrong-heading acceptance makes the connector useful as a source of labels while making it an incomplete model for self-verification.

7 Conclusion

A patch of coincident surface can produce the same observable response as the first few millimeters of entry into a keyed connector. The two responses warrant a second attempt at the same heading, even though neither offers verification. The downstream consequence depends on how much surface is available to verify that second attempt. When a caller only receives an undifferentiated failure code, it loses the evidence needed to make that choice.

Compute statement

We ran the registration replays and global diagnostic on one 16-core workstation, for 31 CPU core-hours including geometry sampling and verification. Policy fitting, resampling, and plotting took less than one CPU core-hour. Deployment requires two rejection likelihood ratios or the corresponding table of decisions for histories of length one to three. The bench stage logged 1,920 traces in 18 instrument-hours including assembly changes.

A Derivation of the transfer inequality

Fix a policy π. Couple the initial hidden headings and use the same private randomization for π in both domains. While histories agree, the actions agree. At each step use a maximal coupling of the two conditional outcome laws, given that common history and its chosen action. The augmented outcomes record whether an acceptance is correct. With probability at least $(1-\delta)^{B}$, all outcomes through termination or exhaustion agree. After termination pad with a fixed terminal symbol. Since the reward is binary, $|V_T(\pi) - V_S(\pi)|$ is bounded by $1-(1-\delta)^{B}$. Apply the bound on each side of the source optimality inequality $V_S(\pi_T^\star) \le V_S(\pi_S^\star)$ to obtain Eq. (3).

B Replay details

The episode generator used root seed 615901 and separate streams for the mesh draw, visible-region geometry, face samples, planted motion, and offsets. Direction signs were balanced within mesh before assigning the geometry seed. The two bootstrap analyses used seed 362714 with separate stream identifiers. Each resample of the first-rejection plot recomputes a ratio of pooled counts, so an assembly with fewer rejections of that code contributes fewer observations to the conditional frequency.

We store the signed counts n_s, n_d for rejections under each candidate and compute the odds in Eq. (1) as $n_s \lambda_s + n_d \lambda_d$, where $\lambda_y = \log(p_1(y)/p_0(y))$. An observation at heading zero increments the corresponding count and one at heading one decrements it. This representation cancels opposing observations before floating-point arithmetic. For the collapsed channel the stated tie rule breaks even-numbered rejection histories toward the first candidate and yields identical traces to strict alternation.

Acknowledgments

We thank the mechanical standards group at Ternwell for the fixture loan and Inez Tavera-Keil for independent coding of the contact traces.

References

Yasuhiro Aoki, Hunter Goforth, Rangaprasad Arun Srivatsan, and Simon Lucey. PointNetLK: Robust & efficient point cloud registration using PointNet. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7163–7172, 2019.

Paul J. Besl and Neil D. McKay. A method for registration of 3-D shapes. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 14(2):239–256, 1992.

Yang Chen and Gérard Medioni. Object modeling by registration of multiple range images. In *IEEE International Conference on Robotics and Automation*, pages 2724–2729, 1991.

D. Chetverikov, D. Svirko, D. Stepanov, and P. Krsek. The trimmed iterative closest point algorithm. In *International Conference on Pattern Recognition*, volume 3, pages 545–548, 2002.

A. Philip Dawid. The well-calibrated Bayesian. *Journal of the American Statistical Association*, 77(379):605–610, 1982.

Bradley Efron and Robert J. Tibshirani. *An Introduction to the Bootstrap*. Chapman and Hall, 1993.

Martin A. Fischler and Robert C. Bolles. Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM*, 24(6):381–395, 1981.

Tilmann Gneiting and Adrian E. Raftery. Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association*, 102(477):359–378, 2007.

Berthold K. P. Horn. Closed-form solution of absolute orientation using unit quaternions. *Journal of the Optical Society of America A*, 4(4):629–642, 1987.

Henry Kautz, Eric Horvitz, Yongshao Ruan, Carla Gomes, and Bart Selman. Dynamic restart policies. In *Proceedings of the Eighteenth National Conference on Artificial Intelligence*, pages 674–681, 2002.

Michael Luby, Alistair Sinclair, and David Zuckerman. Optimal speedup of Las Vegas algorithms. *Information Processing Letters*, 47(4):173–180, 1993.

Sabine Meier-Haddou and Conall P. Irete. Heading choice and terminal travel in visually occluded plug insertion. *Transactions on Manual Assembly*, 19(3):177–191, 2023.

Szymon Rusinkiewicz and Marc Levoy. Efficient variants of the ICP algorithm. In *Third International Conference on 3-D Digital Imaging and Modeling*, pages 145–152, 2001.

Mira Salen-Kuroda and Paavo Essel-Neri. Axial travel at shell and insert contact in keyed interfaces. *Journal of Interface Metrology*, 12(2):73–86, 2020.

Richard S. Sutton and Andrew G. Barto. *Reinforcement Learning: An Introduction*. MIT Press, 2018.

Shinji Umeyama. Least-squares estimation of transformation parameters between two point patterns. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 13(4):376–380, 1991.

Jana Velado-Fenn and Ridwan Waseem-Dorne. Witness-point tolerances for two-candidate registration. Technical Report LGI-25-06, Valdorin Polytechnic, 2025.

Abraham Wald. Sequential tests of statistical hypotheses. *The Annals of Mathematical Statistics*, 16(2):117–186, 1945.

Yue Wang and Justin M. Solomon. Deep closest point: Learning representations for point cloud registration. In *IEEE/CVF International Conference on Computer Vision*, pages 3523–3532, 2019.

Zhirong Wu, Shuran Song, Aditya Khosla, Fisher Yu, Linguang Zhang, Xiaoou Tang, and Jianxiong Xiao. 3D ShapeNets: A deep representation for volumetric shapes. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 1912–1920, 2015.

Qian-Yi Zhou, Jaesik Park, and Vladlen Koltun. Fast global registration. In *European Conference on Computer Vision*, pages 766–782, 2016.