Continuous Pose Preservation in Substrate Extraction

Aino K. Vuorensalmi
Division of Manipulation and Measurement
Alveranta Institute
Turku
Finland
aino.vuorensalmi@alveranta.example

Rafael Quimbayo Sarta
Department of Mechanical Systems
Instituto Palmariego de Ingeniería
Cali
Colombia

Thảo Nguyên Lương-Vũ
Division of Manipulation and Measurement
Alveranta Institute
Turku
Finland

capstract:2611.00019v1 [cs.RO] 10 Nov 2026

***Abstract*—Extracting a deformable substrate from beneath rigid objects is often evaluated by comparing the final arrangement with the initial one. During extraction, limited friction and incidental contacts can displace objects that settle close to their original poses. We specify an extraction task by a bound on the complete pose history of each supported object and implement a controller that transfers the supporting load continuously between the substrate, an external fixture, and the receiving surface. In a matched, two-site evaluation of 144 layouts with one to four objects, a direct withdrawal meets the terminal specification in 129 layouts but meets the trajectory specification at 240 Hz in 61. Closed-loop support passes the trajectory specification in 125 layouts, a difference of 44.4 percentage points. Explicit pickup and replacement gives the highest terminal pass count, 139, and violates the trajectory bound in all 144 layouts. Subsampling the same recordings at 30 Hz increases the apparent trajectory pass count for direct withdrawal from 61 to 86. We report the effect of imposing contact restrictions in addition to pose tolerances and identify the residual ambiguity in calling an object stationary during a tabletop linen exchange.**

***Index Terms*—Manipulation, Deformable materials, Force control, Pose tracking**

I. Introduction

A task specification for manipulation usually names a set of objects to be moved and a target arrangement. Preserving the arrangement of the other objects can be formalized as an avoidance constraint. In some tasks the two sets are in direct contact at the initial instant and must remain so for part of the motion. To replace a surface liner, for example, the robot must move the material under the set of objects against which pose preservation is assessed. Common intermediate states include an object partly resting on each surface and an object making contact with a folded portion of the liner while fully supported elsewhere. These states are relevant even when the final arrangement agrees with the target.

We study the case in which the removable substrate is a tablecloth and the payload is a place setting. In a specification collected from tabletop demonstrations, the directive is to remove the cloth without moving any of the articles on it [1]. A photograph of the setting after the cloth is gone will often satisfy the disposition criterion. For a fixed camera looking down, two photographs are also insensitive to temporary lifting and replacement of the same set of objects. The photographs establish that the setting is in the same place twice. The motion requirement concerns the interval in between.

A robot can clear the table and reconstruct the arrangement if the contract specifies only final poses. For a path bound we study a fixture that supplies support during extraction. It allows a change of load path while limiting article excursions in the table frame.

We report a controlled comparison of three interpretations of the task, with a sampling-rate experiment on the same motion records. Section III defines the contract and Section IV gives the controllers and the observation procedure. The experiments in Section V concern complete cloth extraction on an enclosed benchtop rig with acrylic articles representing the outer geometry of place settings. The interpretation of “without touching” requires an additional predicate on the interaction history. We record it separately from the trajectory criterion.

II. Contact and Task Representations

Contact friction and the support geometry constrain induced motion [2]–[4]. Pushing measurements show variation across materials and travel directions [5]. Earlier work on liner handling studied withdrawal from beneath a moving clamp on an unloaded fabric region [6]. Daniotti and Almarz [7] measured shifts in stacked rigid articles during support exchange. A spread place setting puts independent normal loads and friction states on a shared sheet.

A finite-dimensional state and force constraints are standard in rigid-body trajectory optimization with contact [8], [9]. The compliance at the substrate matters here because its thickness and local wrinkles have the same order of magnitude as the displacement acceptance threshold. The support channels use position reference and normal-force feedback, following the

Preprint.

distinction between hybrid position and force control [10] and mechanical impedance [11]. Their purpose is to regulate different outputs during the transition. The model of rigid-body equilibrium over a contact region [12] explains why a single measured net upward force would be insufficient for a set of objects with different contact patches. One channel near saturation may be concealed in a measurement of force summed over the array.

Ramal and Nordviken [13] reported endpoint agreement in single-article rail experiments. Mikkonen and Urrál [14] used two synchronized views for monitoring brief excursions in moving fixtures and found that camera placement influenced which excursions were missed. We measure pose relative to the table frame, assign completion using a separate cloth trace, and evaluate both the endpoint and the sampled trajectory against the same initial object identities. The optical tracking gives an operational answer at its designated sampling rate. A mathematical guarantee on the path between frames requires additional information about intersample motion.

III. Task Specification

The scene contains a textile sheet and k distinguishable rigid objects initially on its upper surface. Let $x_i(t)$ and $R_i(t)$ denote translation and rotation of object i in a rigid table coordinate frame. The reference (x_i^0, R_i^0) comes from the stationary pretrial window. We normalize translation by $d = 1.0$ mm and rotational geodesic angle by $\theta = 1.0$ degree. Eq. (1) gives the worst deviation among objects. The axis of rotation is unrestricted. We use the body frame of the tracking geometry to define the origin, so the pose comparison includes tipping, and also includes yaw of a near-circular object.

$$e(t) = \max_{1 \le i \le k} \max \left(\frac{\| x_i(t) - x_i^0 \|_2}{d}, \frac{\angle((R_i^0)^\top R_i(t))}{\theta} \right). \tag{1}$$

The cloth must clear the table as a single connected piece by time T, with all original objects present. The endpoint test asks whether $e_{\mathrm{end}} \le 1$, where e_{end} applies Eq. (1) to the pose estimated over the final stationary window. The trajectory test requires endpoint success and $e(t_j) \le 1$ at every recorded time t_j from the reference window through the final one. Including the separately estimated endpoint makes the two tests nested by construction. Robot withdrawal is part of the observation interval. The clock starts before any support hardware engages and stops after the support fixture has disengaged. A support plan that holds the third object in place while leaving the fixture attached at release does not complete the task.

The article identities matter. Replacing a displaced article with a second article of the same shape could meet an unlabeled registration score. Our acceptance procedure indexes each pose by its original fiducial identity. It also requires the cloth to exit whole. Cutting away the visible border and leaving an island under each object would violate that condition even if the perimeter camera gave an unobstructed view of the table edge. Each sheet has preprinted fiducials distributed through the interior and was checked in the receptacle after the trial. The fabric is a single piece before and after the run.

There are two relevant notions of contact. The narrow rule disallows deliberate direct contact between any robot tool and an article during the interval. The broad rule disallows interventions that change the article’s support state. Under the broad rule the task is infeasible by definition, because extracting the original load-bearing sheet entails replacing its support. We can measure the narrow rule with contact sensors and video review. A direct withdrawal can satisfy it while exerting appreciable tangential load on an article through the cloth. A fixture can satisfy the position criterion while touching every article.

IV. Apparatus and Control

A. Load transfer

In the proposed controller an overhead fixture has a compliant support channel for each occupied article site. There is a low-profile capture boss at a known coordinate on each experimental article, and the overhead channel engages that boss. The extraction actuator handles the textile. The fixture reacts against the table frame on a metrology substructure shared by the cameras. This lets the overhead channel compensate for the net weight that would otherwise load the substrate. A small residual load is useful during initial engagement because it limits the tendency to lift the article from its datum.

For a single article near equilibrium, let $N_{c,i}$ be normal load on the cloth, $F_{h,i}$ the upward component from the fixture, and $N_{r,i}$ the normal reaction from the receiving surface. Neglecting vertical acceleration yields Eq. (2). We track the relative magnitudes from contact probes. During extraction, $N_{r,i}$ is initially zero and becomes observable once the cloth clears the base. The tangential reaction from the cloth is bounded by $\mu_{c,i} N_{c,i}$ in a Coulomb approximation. This bound is a local model and ignores snagging and capillary adhesion.

$$N_{c,i} + F_{h,i} + N_{r,i} \simeq m_i g. \tag{2}$$

Load transfer lasts through two events detectable with the rig. During engagement the compliant channel establishes its spring reference in the initial pose, then the fixture takes the load. After the trailing fabric boundary clears the article, the receiving surface takes the load as the fixture releases. This changes article height by approximately the compressed local thickness of the fabric. Thick folded hems can transiently load a channel even when its mean support force is near nominal. Channel travel and end-effector deflection then matter. The pose current at recontact is the one evaluated in the same table frame as the reference pose. A level array with all articles supported by the tool can pass the height tolerance if the descent through the removed material occupies less than the remaining vertical error budget.

Horizontal position feedback holds each capture point at its registered datum. Vertical force feedback maintains the assigned load share until probes indicate cloth clearance, then

transfers load to the receiving surface. The controller pauses extraction when any channel exceeds its displacement margin and aborts when a channel runs out of travel or the cloth loses tension. A separate multiview tracker with no command link supplies the scored poses.

B. Layouts and comparison methods

Two sites used independently assembled rigs. Each fabricated three textiles from the same three lot specifications, identified as light plain weave, medium twill, and a heavy weave with a bound edge. At each site and for each fabric, we drew six layouts at each of four article counts. This gives 72 layouts per site and 144 overall. A draw selects article shape and orientation from a fixed inventory and places objects in independently drawn disjoint sites within a marked region. All articles are inert. A decorated face makes yaw trackable without relying on the outline. The test inventory has no liquid, breakable serviceware, or weighted cutlery. A set of independently prepared calibration layouts was reserved at each site.

A removable jig reset the layout before each trial. Every controller used a separately conditioned specimen from the same cloth lot. Controller order was randomized within layout. We measured the reference positions after every reset and kept all assigned trials.

Rapid withdrawal follows a fixed site-specific command chosen on calibration layouts, with the support fixture stowed. Replacement lifts the articles beyond the trajectory tolerance, extracts the cloth, and returns them to the recorded poses. All methods run on the same rigs with the same extraction actuator. The 90 s timeout includes support engagement and release where applicable.

C. Observation and analysis

Three views record poses at 240 Hz. We use independent camera estimation with fixed fiducials on the table for hand-eye and multiview checks [15], [16], and article tags for point association [17]. An event signal shared with the cloth sensor gives the acquisition timestamp. The reference and final pose estimates use medians for center coordinates and the least-squares rotation in the corresponding 200 ms dwell window. Within the observation interval we score each native frame, with no temporal smoothing applied to the sequence used for maxima. The point spread, pose residuals and jitter of the calibration records were evaluated before measuring outcomes. Appendix A describes the checks.

Every retained frame must have a valid pose for every article to pass the trajectory check. An aborted run fails both tests. A run with an unresolved intermediate pose can pass the endpoint test if the cloth clears and the terminal pose is available. For direct contact, reviewers inspected pad sensor records and videos from the overhead and oblique cameras. They saw the state of the support fixture, with the numerical pass decisions withheld.

The primary contrast is the paired difference in trajectory pass probabilities between closed-loop support and rapid withdrawal. With $Y_{m\ell}$ the binary trajectory result for method m and layout ℓ, the contrast is Eq. (3). We report a percentile interval from 20,000 paired resamples of layouts stratified by site. Resampling a layout carries every treatment and outcome recorded for it. This conditions inference on the two test sites and their proportions in the design [18]. The sample includes six draws per fabric and article count at each site. Fabric-specific comparisons and the sampling-rate analysis are descriptive.

$$\widehat{\Delta} = \frac{1}{144} \sum_{\ell=1}^{144} (Y_{\mathrm{support},\ell} - Y_{\mathrm{rapid},\ell}). \tag{3}$$

V. Results

A. Terminal and trajectory assessment

For endpoint assessment, replacement gave the highest pass count in Table I at 139 of 144. It passed no trajectory assessment because every run contained deliberate lift. We observed the lift on the fixed-view metrology records even when the same object identities returned within 1 mm and 1 degree. Rapid withdrawal completed with terminal poses inside tolerance in 129 layouts, but failed the trajectory test in 68 of those. Closed-loop support met terminal tolerance in 137 and the trajectory criterion in 125. Four layouts passed the trajectory criterion by rapid withdrawal alone, 68 passed by support alone, and 57 passed by both. The paired estimate was 44.4 percentage points (95% interval 35.4 to 53.5). Each site gave the same direction of effect. At the Turku site the pass counts were 33 rapid and 64 support. In Cali they were 28 and 61. Within-site counts of layouts passed by both methods were 31 and 26, respectively.

TABLE I
NUMBER OF COMPLETE EXTRACTIONS PASSING EACH CRITERION. EACH FABRIC SUPPLIES 48 MATCHED LAYOUTS ACROSS THE TWO SITES. T DENOTES TERMINAL AND P DENOTES TRAJECTORY. TRAJECTORY EVALUATION USES 240 HZ RECORDS.

Fabric	Layouts	Rapid T	Rapid P	Replace T	Replace P	Support T	Support P
Plain weave	48	46	31	47	0	47	45
Twill	48	43	22	46	0	46	42
Bound weave	48	40	8	46	0	44	38
All	144	129	61	139	0	137	125

In Fig. 1 we plot the earlier middle-ranked rapid trajectory by peak error among its 68 endpoint-only passes, with the corresponding support trial. The rapid peak occurs before the fabric tail clears the second article. The smaller support peak follows clearance, during pad release. The traces align on each trial's own last clearance event.

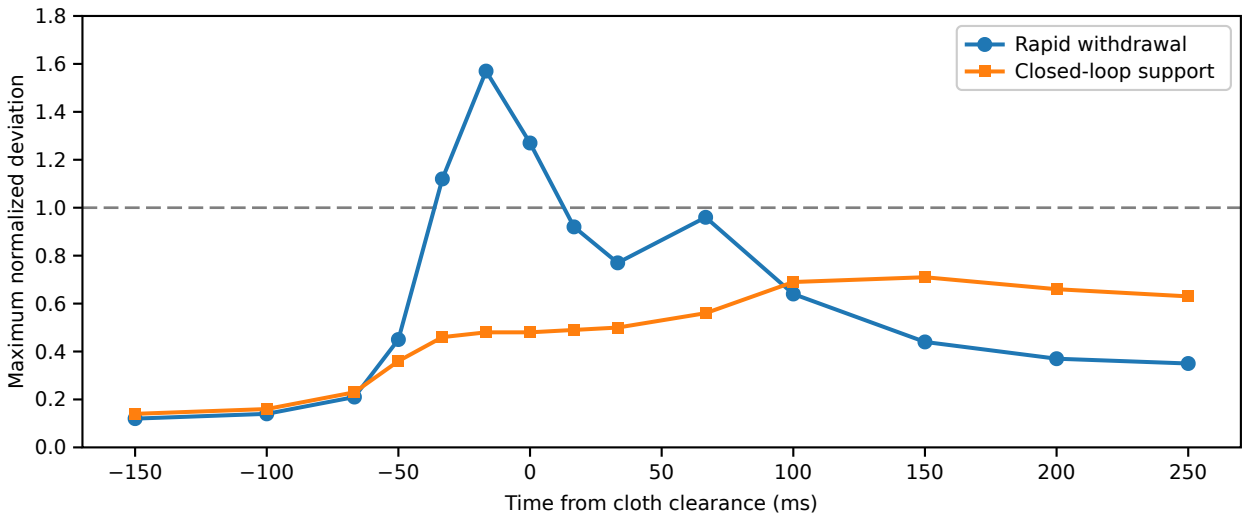


Fig. 1. Pose deviation for two controllers on the same layout. Time zero is clearance of the fabric tail from the last article's projected footprint. The dashed line gives the fixed trajectory threshold. Data shown near clearance only.

Among the 19 support trajectory failures, 3 aborted before full cloth withdrawal and 1 had an unresolved intermediate pose. By first threshold crossing, 9 exceeded the translation limit near the cloth tail, 4 exceeded it at release, and 2 exceeded the rotation limit during engagement. Height dominated the error in 10 of the 13 first translation crossings. Four of these 15 numerical failures also finished outside terminal tolerance. One of the four first crossed the rotation limit on one article and finished with a translation error on another. Rapid withdrawal had 2 incomplete extractions and 3 intermediate tracking failures. Replacement cleared all 144 cloths. Its five terminal failures came from landing offsets.

We show sensitivity to translation tolerance in Fig. 2. For each tolerance we applied the specified endpoint and trajectory rules, leaving the angular tolerance and the observations fixed. Support gains most between 0.5 and 1.0 mm, an interval that contains the compressed thickness of the twill and of most bound specimens. The cutoff at 1 mm was set during the calibration stage and shared across sites. At 5 mm and above, the six residual support failures consist of three aborts, one unresolved pose, and two angular excursions. The departure of a single object remains sufficient to fail a layout.

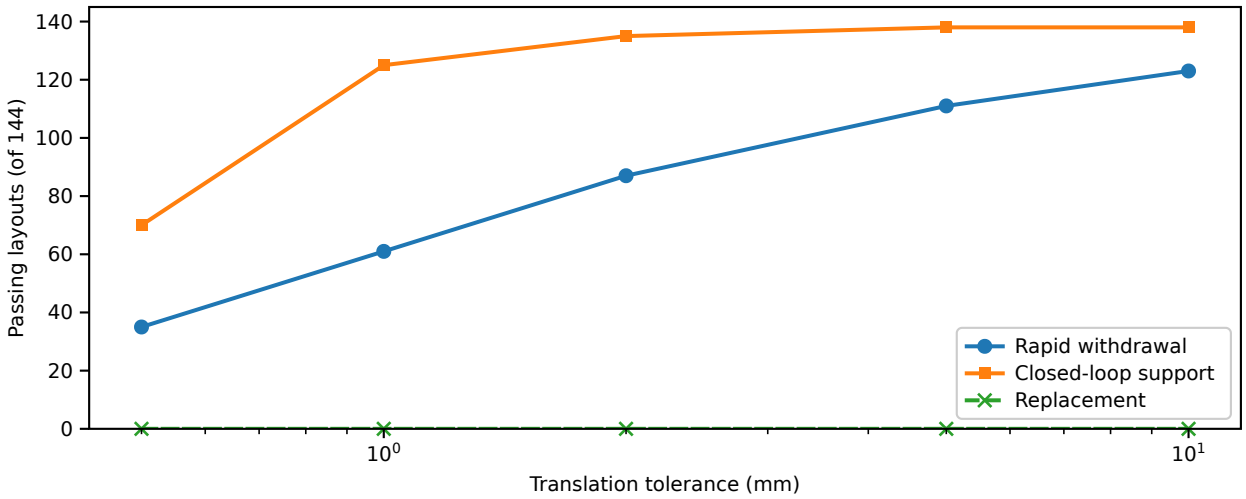


Fig. 2. Trajectory pass counts as a function of the translation tolerance. The rotational limit remains 1 degree. Complete extraction, visibility, and endpoint criteria are applied at each tolerance.

B. Evidence available to an observer

Fig. 3 compares the same runs at nested temporal resolutions. We started decimation at the first complete frame of the trajectory interval and retained the independently estimated endpoint in every condition. Reduced sampling alters both the chance of registering an excursion and the chance of landing on an unresolved pose. The lack of smoothing is deliberate. Sampling includes fewer observations of the same process, with identical measurements at common timestamps. Since the assessment is a conjunction over observations, decimation can only increase the apparent pass count at the fixed phase. On visually uneventful video the first motion may be visible as a slight departure near fabric clearance. The test at 30 Hz admits 25 rapid trajectories rejected at 240 Hz. A comparison of terminal images admits 68.

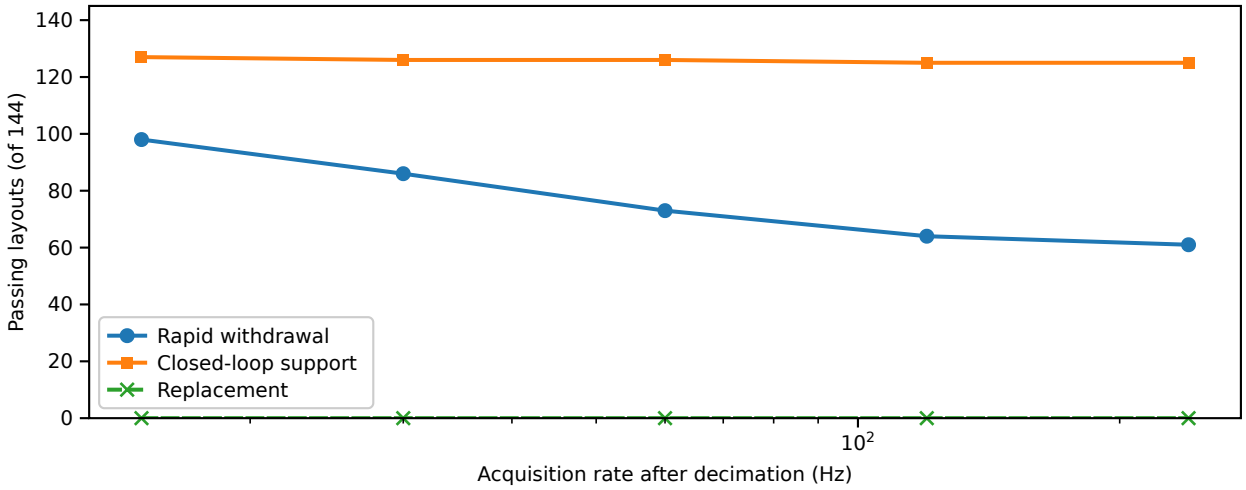


Fig. 3. Apparent trajectory passes under nested decimation of the 240 Hz motion records. Both accuracy tolerances and the terminal observations are fixed. Each series uses the same 144 layouts at every rate.

We also swept the eight start phases of 30 Hz sampling over a common window excluding camera-transition frames. Rapid passes varied from 82 to 90 by phase. The union was 108 layouts and the intersection was 64, equal to the native-rate result on that restricted window. Extending the window to the full primary interval adds three tracking failures at native rate.

C. Contact restrictions

With the additional direct-contact criterion in Table II, all 61 rapid passes survive. Every support pass relies on a fixture touching each article. The cloth alone couples the rapid actuator to the payload.

TABLE II
PASS COUNTS WHEN ADDING DIRECT TOOL CONTACT AS AN ADDITIONAL DISQUALIFIER. CONTACT-FREE TRAJECTORIES MUST ALSO SATISFY ALL PRIMARY CRITERIA.

Controller	Trajectory	Contact-free trajectory
Rapid withdrawal	61	61
Replacement	0	0
Closed-loop support	125	0

We logged an auxiliary contact history for accepted support trials. Every article in those trials left the fabric and was carried by the fixture alone for at least one registered time interval before returning to surface support. An article can thus satisfy the pose history, fail the direct-contact criterion, and be unsupported by the tabletop while the cloth is removed. No determination based only on the apparent placement of the article can exclude that sequence.

VI. Discussion

The endpoint criterion and the path criterion define different tasks even when the tolerance values are identical. For inventory arranged by a robot after the operation, obtaining a satisfactory final photograph may be the practical objective. During simultaneous service, a client may care whether the contents of a vessel experienced a transient jerk. The ordinary pose criterion approximates such a requirement only indirectly. Bounding the displacement of a vessel does not bound acceleration without temporal or dynamical assumptions. It cannot certify the absence of sloshing inside a vessel. An internal state constraint would require a different measurement model. The present results apply to inert payloads.

For a strict movement specification, at least one surface definition needs attention. If the top of the textile initially supports an article and the tabletop supports it at the end, height differs by a compressed thickness even after a perfect transfer. Under a tabletop datum the path must include that vertical change. An endpoint test that absorbs the sheet thickness into the nominal pose would make the terminal error smaller while treating the transition inconsistently. If the aim were to preserve only each article’s projected position as seen from above, the height correction and clearance behavior could both be ignored. Such a score could give identical readings for a full retreat in plane-normal direction and a trajectory with no detectable translation.

The supported trials benefit from engineered capture features, a flat reference and access from above. On a restaurant table, a tool able to reach several individual supports would encounter lighting, people and tall articles. We did not assess those conditions. The broad meaning of contact-free transfer also differs from withdrawing a sheet quickly as a demonstration. Any direct change to the support state makes the broad interpretation impossible for a load resting on the sheet. Most disagreement can be resolved by naming the forbidden contact and the coordinates in which preservation is measured. Both details should be fixed before judging the performance of an extraction.

VII. Conclusion

The acquisition record permits a reviewer to assess the pose and contact requirements independently. A recording that begins after pickup can conceal a whole stage of replacement. A completed support trace can satisfy the pose requirement even though for some of the interval the place setting is carried by neither the cloth nor the tabletop.

Appendix A. Measurement and Sampling Checks

We calibrated views using a rigid artifact with independent linear and angular settings spanning the acceptance neighborhood. After alignment, a static interchange between the two outer cameras had a median 3-D position discrepancy of 0.09 mm and a 95th percentile of 0.21 mm. The corresponding angular discrepancies were 0.08 and 0.19 degree. These dispersions characterize agreement at repeated poses, not absolute accuracy [19]. The fixture carries its own reference markers. A static load on the support array displaced its pointers with respect to the table fiducials. Unmodeled deflection of the camera support would enter measurement as an apparent common object shift. We measured the reference drift on unloaded and loaded calibration runs and set the table frame from its own fiducials every trial.

An interframe statement is possible when a valid bound exists on the instantaneous rate of the normalized pose deviation. Suppose $e(t)$ in Eq. (1) is L-Lipschitz over an interval with uniform sampling period h and both boundaries sampled. The sampled maximum M_h is then related to the continuous maximum M by Eq. (4). A maximizer is within $h/2$ of a sample. This is an amplitude bound. It does not require a band-limited trajectory and is separate from signal reconstruction guarantees [20]. Impacts and contact transitions make a reliable small global L difficult to establish. Peak estimates from finite differences cannot substitute for that bound.

$$M_h \le M \le M_h + \frac{L h}{2}. \tag{4}$$

We used 240 Hz as a defined observation rate and did not extrapolate its pass rate to arbitrarily fine time resolution. The shortest accepted motion was still assessed over the entire engagement and release interval. Drift estimates were computed before treatment comparison. If validation required a bound on continuous displacement, error from static calibration and interframe motion would together reduce the threshold available to the sample maximum. That stricter threshold would make fewer claims of stationarity for specimens near the lower thickness limit.

Appendix B. Paired Outcome Arithmetic

For the predeclared primary comparison each matched pair contributes $+1$, -1, or zero to the numerator of Eq. (3). In Turku the contributions were 33 positive, 2 negative, and 37 zero. In Cali they were 35 positive, 2 negative, and 35 zero. The observed difference is $(68-4)/144 = 0.4444$. Stratified resampling uses 72 draws with replacement at each site.

Replacement is excluded from the primary contrast because its planned clearance moves all articles outside the support neighborhood. Its trajectory result instead confirms the distinction between the task formulations using the same object registration and reset protocol.

Acknowledgments

We thank the Alveranta and Palmariego instrument workshops for fabricating the matched fixture arrays.

References

[1] H. E. Shevlin and J. Vallerand, “Disposition Clauses in Tabletop Demonstrations,” *Journal of Experimental Material Practice*, vol. 7, pp. 53–71, 2011.

[2] M. T. Mason, *Mechanics of Robotic Manipulation*. MIT Press, 2001.

[3] R. M. Murray, Z. Li, and S. S. Sastry, *A Mathematical Introduction to Robotic Manipulation*. CRC Press, 1994.

[4] K. M. Lynch and M. T. Mason, “Dynamic Nonprehensile Manipulation: Controllability, Planning, and Experiments,” *The International Journal of Robotics Research*, vol. 18, pp. 64–92, 1999.

[5] K.-T. Yu, M. Bauza, N. Fazeli, and A. Rodriguez, “More than a Million Ways to Be Pushed: A High-Fidelity Experimental Dataset of Planar Pushing,” in *Proc. IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2016.

[6] P. A. Cordel and M. Paredes Leiva, “Sequential Isolation of Loaded Regions in Surface Liner Handling,” in *Proc. Symposium on Textile Handling and Automation*, 2018.

[7] L. Daniotti and H. Almarz, “Observations of Base Displacement During Quasistatic Support Exchange,” *Transactions on Applied Manipulation*, vol. 16, pp. 207–223, 2021.

[8] M. Posa, C. Cantu, and R. Tedrake, “A Direct Method for Trajectory Optimization of Rigid Bodies Through Contact,” *The International Journal of Robotics Research*, vol. 33, pp. 69–81, 2014.

[9] K. M. Lynch and F. C. Park, *Modern Robotics: Mechanics, Planning, and Control*. Cambridge University Press, 2017.

[10] M. H. Raibert and J. J. Craig, “Hybrid Position/Force Control of Manipulators,” *Journal of Dynamic Systems, Measurement, and Control*, vol. 103, pp. 126–133, 1981.

[11] N. Hogan, “Impedance Control: An Approach to Manipulation. Part I: Theory,” *Journal of Dynamic Systems, Measurement, and Control*, vol. 107, pp. 1–7, 1985.

[12] T. Bretl and S. Lall, “Testing Static Equilibrium for Legged Robots,” *IEEE Transactions on Robotics*, vol. 24, pp. 794–807, 2008.

[13] N. Ramal and S. Nordviken, “Terminal Disposition Following Guided Substrate Withdrawal,” in *Proc. Workshop on Contact Assessment in Manipulation*, 2022.

[14] S. E. Mikkonen and F. Urrál, “Occlusion and Transient Pose Errors in Dual-View Fixture Metrology,” *Industrial Motion Measurement*, vol. 12, pp. 307–322, 2024.

[15] R. Y. Tsai and R. K. Lenz, “A New Technique for Fully Autonomous and Efficient 3D Robotics Hand/Eye Calibration,” *IEEE Transactions on Robotics and Automation*, vol. 5, pp. 345–358, 1989.

[16] R. Hartley and A. Zisserman, *Multiple View Geometry in Computer Vision*. Cambridge University Press, 2004.

[17] E. Olson, “AprilTag: A Robust and Flexible Visual Fiducial System,” in *Proc. IEEE International Conference on Robotics and Automation*, 2011.

[18] B. Efron and R. J. Tibshirani, *An Introduction to the Bootstrap*. Chapman and Hall, 1993.

[19] J. M. Bland and D. G. Altman, “Statistical Methods for Assessing Agreement Between Two Methods of Clinical Measurement,” *The Lancet*, vol. 327, pp. 307–310, 1986.

[20] C. E. Shannon, “Communication in the Presence of Noise,” *Proceedings of the IRE*, vol. 37, pp. 10–21, 1949.