capstract:2611.00018v1 [cs.CV] 10 Nov 2026

Recovery of prewash azimuth in undecorated white plates

Silvia D. Nevrani
Department of Imaging Metrology
Istituto Tecnico di Valcerro
Parma
Italy
s.nevrani@valcerro.example

Emiel K. Verbrugsen
Surface Imaging Group
Meerveld Institute of Applied Optics
Breda
Netherlands

Nari H. Belcourt
Department of Imaging Metrology
Istituto Tecnico di Valcerro
Parma
Italy

***Abstract*—Replacing a circular object in its previous position leaves one degree of freedom undetermined by its silhouette. We examine whether a record of the upper surface supplies enough information to restore the azimuth of an undecorated white plate after washing. Relative rotation was measured independently from below on 128 plates from 16 production lots. We reserved four lots for method development. On the remaining 96 plates, a median of normalized correlations over radial strips recovered the prewash heading to within five degrees on 462 of 576 dry placements following one wash (80.2%). The ordinary annular correlation recovered 381 headings (66.1%), a paired difference of 14.1 percentage points (95% lot-level interval 9.5 to 18.7). A threshold on separation between competing correlation peaks accepted 420 placements, of which 399 met tolerance. After ten washes, 397 of 576 placements met tolerance against the original reference. The method requires a reference image of the actual plate in its original setting. Substituting a reference from the same carton reduced one-wash success to 18 of 576 placements.**

***Index Terms*—Image registration, rotation estimation, surface inspection, object symmetry**

I. Introduction

The position of a soup plate at a table can be recorded by the center of its rim. This record suffices to put a replacement plate at the same place. It leaves open which part of the rim had faced the diner. The ambiguity also arises when the *same* plate is returned after a meal. A plate with a border motif can be turned until the motif agrees with the record. For plain white service there is seldom an agreed point on the circumference from which to take the bearing.

We consider the following specification for returning the plate. Register its center to the old center, bring its plane parallel to the old plane, and return its upper surface to the azimuth in a photograph made before the table was cleared. At a radius of 110 mm a five-degree error puts a point on the rim about 9.6 mm from its former position along the arc. We use five degrees as the acceptance tolerance. We audit this tolerance from below. Whether a diner observes an error of that size is a separate question.

This last degree of freedom is usually taken out of an object pose benchmark. For a solid of revolution, a rotation about its axis can leave a CAD rendering identical. Distance to the nearest visually equivalent rendering is then a sensible loss [1], [2]. The dishes considered here have that symmetry in the catalog drawing. A particular dish, at sufficient magnification, may have iron specks under the glaze or a small deviation in the shape of its lip. If one of these features survives a wash and appears in both photographs, it defines a material point. A plate need not have a front in order to have a previous heading.

The difficulty is occasional gross error. A camera can resolve an abrasion to a fraction of a degree while the abrasion is there. On a plain piece there may be little to distinguish its trace from a drying streak in the next image. Averaging these cases into one angular precision would obscure the distinction between localizing a match and finding one. We count all placements within an angular tolerance, including those for which a method finds no convincing match. We then give coverage separately for a method that can decline to provide a return angle.

Our estimates concern a known item in a controlled overhead view. Among other things, this choice removes the problem of determining which plate was cleared from which place. A return procedure involving a mixed stack would also need to solve that problem. Even perfect plate identification supplies no angle by itself.

II. Image correspondence and pose equivalence

Handling symmetry by predicting several admissible poses [3] retains information that can be lost when a regressor interpolates between them. Sampling a conditional pose distribution offers another way of representing this uncertainty [4]. Image-specific pose sets also allow an evaluator to distinguish a symmetry of the object from the temporary loss of an informative feature [5]. In our setting, several distinct maxima of a matching function are often defensible from the photograph. Only one is correct in material coordinates. The underside measurement in Section III makes it possible to score this quantity on images where the face alone leaves the answer undecided. In particular, we can keep unmarked pieces in the test sample without asking an annotator to choose between indistinguishable views.

For a similarity transformation in the image plane, resampling on a polar or log-polar grid changes rotation into translation along one axis [6]. Cartesian registration lends itself to fast correlation and subpixel refinement [7]. We use the outline

to remove translation and size first. A nearly circular outline is quite helpful at this stage, although its registration score as a function of rotation can be almost flat. Once the center is known we can look for agreement around a ring without estimating a projective transform for every candidate angle. The treatment of the center error matters when the signal is a hairline scratch.

Local descriptors allow most of the image to change between exposures if a few points can be matched [8], [9]. Regional tests on otherwise similar surfaces face the opposite problem from wide-baseline matching of richly textured patches. There can be several hundred plausible points in a stack of photographs and fewer than four on the plate actually being returned. A learned comparison of local patches could change how those points are weighted [10]. The limited number of independent lots in this study restricts what can be learned about rare glaze finishes. We include a sparse feature baseline and two dense methods that require no training images from the test lots.

In earlier measurements of fired test coupons, Sorel-Davenne and Ivasko-Renn [11] obtained repeatable bearings from glaze pits at fixed illumination, using a coupon holder keyed to one edge. Tersen and Rivelo [12] measured registration through a rinse on circular glazing samples and allowed the receiving fixture to fix the initial quadrant. An inventory of refectory tableware [13] records the backstamp as an aid to identifying pieces from mixed services. These are useful reference conditions. Clamping the plate through a wash, providing its quadrant, or reading a backstamp on the underside would change the problem studied here. In particular a backstamp can be placed at any angle relative to a reference photograph of the upper face.

III. Plates and image acquisition

A. Sampling and allocation

We obtained sixteen cartons, each from a separately documented production lot, through three ceramic stockrooms. We selected eight round, white glazed soup plates per carton in packing order. The order was recorded before examining individual faces. Eligible lots had a plain circular rim, an unprinted upper face and a foot capable of resting on the open support. A color band, a relief motif or a scalloped edge rendered the entire lot ineligible. Small inclusions, scratches and mold seams on individual pieces were admissible. Excluding them would have conditioned the sample on precisely the evidence available for registration.

Four lots, chosen by a draw before imaging, supplied the development set. Their 32 plates were used to select the annulus, the score aggregation rule and the margin for abstention. The other twelve lots supplied 96 test plates. All pieces from a carton stayed together, including those with consecutive kiln numbers. The plates were unused when enrolled. The outside diameters covered 204 to 248 mm. We recorded the original finish as glazed white in the inventory regardless of slight differences in cream tone. We did not attempt to label a faint mark on a single image as a defect.

Each plate supplied one upper reference frame at enrollment. We made six new placements before washing, six after one wash, and six after the tenth wash. The counts in Table I include the first six placements as actual removals from the support. Digital rotations of the reference would leave dust on the sensor and the reflected light field rigidly attached to the plate. Each placement required the operator to lift the piece clear, turn it and reseat it. All decisions on camera gains and score settings were finished with the development data before comparison on the test cartons.

TABLE I
Acquisition counts. Each query is paired with the original reference of its own plate. The lower-camera exposures used for scoring are additional to these counts.

Partition	Lots	Plates	Upper references	Dry reseats	After one wash	After ten washes
Development	4	32	32	192	192	192
Test	12	96	96	576	576	576
Total	16	128	128	768	768	768

B. Camera arrangement

A fixed overhead monochrome camera viewed a 310 mm square field at 2048 by 2048 pixels. We used a diffuse white dome with a polarizer across its exit and a crossed analyzer at the lens. A flat-field frame of matte white ceramic and a dark frame were taken at the start of each session. Gain and exposure stayed fixed during that session. The surface image is the ratio of the dark-corrected signal to the corresponding flat field. Low-frequency illumination gradients remained appreciable on the steep part of the well. The polarizers also left a narrow reflection on some rolled lips. The treatment of both appears in Section IV.

A separate camera looked up through the center of an annular support. Two distinguishable black enamel dots on the interior base of the foot supplied a direction vector. This part of the foot was opaque in transmitted light and lay below the opaque upper glaze. We varied the direction of the dot pair relative to the stamped manufacturer’s name. The operator saw the plate from above and could not use the dots to align it. The two exposures shared a trigger. Before each session we imaged a double-sided plane target in the support, estimated distortion from tilted views [14], and measured the change of handedness and angle between the upper and lower view of the table. We removed the target for all plate exposures.

The reference axis on the table points toward the nominal diner’s position. It defines how a clockwise return instruction would be read on the upper camera. It need not pass through a visible feature on the plate. For any one reference-query pair the difference in lower-camera bearings gives the signed rotation of that material about the table normal. A constant clocking error in the lower camera cancels within a session. To compare across sessions we apply the daily target transform before differencing. Appendix A gives an independent check of this step.

We assigned six sectors of 60 degrees in random order to the six queries of each plate at each stage. A target offset within each sector was drawn uniformly. A removable pointer on the lifting fixture allowed coarse settings to these targets. The pointer left the camera field before exposure. We used the continuous measured bearing, including the operator’s setting error, in scoring. The estimate for a given pair cannot use the other five trials, their order or their sectors. Sectors serve to cover the circle even when a plate is so smooth that the predicted heading stays close to zero.

C. Wash history and exposure control

One complete rack pass with detergent and a rinse constituted a wash. We distributed lots across rack positions and changed their positions between passes. Pieces cooled on edge and air-dried before top and bottom imaging. A visible pool of water in the well delayed capture until the pool was gone. We did not polish the surface after drying. No plate was used for serving food between cycles. The one-wash comparison therefore tests changes in an initially clean service, including the loss of loose marks present at enrollment. It gives little information on the persistence of evidence introduced by a meal.

The reference and the first six queries were made on the same day. The last two stages were photographed on separate days using the same camera and support. We changed the center position within the available play in the fixture at every placement, and alternated between two matte backing inserts according to a schedule independent of angle. We kept the identifiers used by the stockroom, the absolute lower angles and all sequence counters out of the scoring inputs. A query comprised an upper reference array and an upper target array. The algorithms were permitted to assume that these showed the same plate.

IV. Angular registration

A. Preprocessing

We locate the outer edge against the dark insert with a gradient edge detector [15], then fit a circle to the external contour. The plate may be slightly oval. The fitted radius is a unit of sampling distance, rather than a claim that every radius is equal. After undistortion the optical axis is within 0.3 degrees of the support normal. At the radii used below the dominant residual error is horizontal decentering. We reject contour points more than two pixels from the initial fit, refit once, and retain the resulting circle for all methods using the upper face. Every test photograph admitted a circle fit with at least 80% of the circumference represented.

Let r denote radial distance divided by this fitted radius, and let the angular bin index j increase clockwise in the upper view. We bilinearly sample the corrected image at 720 bins and 96 equally spaced radii between 0.54 and 0.94 inclusive. This leaves the outer rim available to locate the center but excludes the boundary between table and plate from the texture score. In one family the bottom of the soup well ends inside this interval. We retained these rows. Their radial mean can be much darker than that of the lip without supplying information about angle.

We subtract a circular running mean of width 31 angular bins from each row and average neighboring groups of four radial rows. The result is a 24 by 720 array for each view. Clipped pixels at either sensor limit are invalid. A strip entry requires at least two valid rows in its group, and the running mean uses the available samples in its window. Concentrating on circumferential contrast attenuates broad changes in shading. It also removes some large, stable variations in glaze tone. Narrower and wider windows were compared on the development lots only.

B. Estimators

For strip k and positive shift d, $C_k(d)$ is the sample Pearson correlation of the reference strip $u_k(j)$ with the query strip $v_k(j+d)$. Indices wrap at 720 bins. Means and sums of squares are evaluated on $J_k(d)$, the bins valid in both views at this shift. Writing a_j and b_j for the respective signals with their overlap means subtracted,

$$C_k(d) = \frac{\sum_{j} a_j b_j}{\sqrt{\sum_{j} a_j^2}\sqrt{\sum_{j} b_j^2}}, \qquad j \in J_k(d). \tag{1}$$

The means in Eq. (1) are specific to the strip and shift. We omit strip scores supported on fewer than 540 angular bins or having a zero denominator. A shift with fewer than twelve supported strip scores receives no score. The two views do not generally clip in the same bins. Merely replacing clipped bins with black produced a peak at the angle required to align two reflections. The computations can be expressed as circular correlations of the signals, squared signals and masks, using the correlation form of the convolution theorem [16]. We evaluate the complete circular score before returning a hypothesis.

At each shift we take the median of supported strip scores. The clockwise motion in the query relative to the reference and the opposite instruction for returning the dish are

$$\begin{aligned} S(d) &= \operatorname{median}_k C_k(d), \\ \hat{\theta} &= \tfrac{1}{2}\,\arg\max_{0 \le d < 720} S(d), \\ \theta_{\mathrm{return}} &= -\hat{\theta} \pmod{360^\circ}. \end{aligned} \tag{2}$$

Here the factor of one half has units of degrees per bin. We break exact ties at the smaller nonnegative bin index. With no supported shift the default estimate is zero degrees and the

confidence margin is zero. There is no refinement beyond the bin grid. Its worst nearest-bin quantization error is a quarter of a degree, smaller than the inspection tolerances for which we report performance. Sub-bin localization would leave a 140-degree selection of the wrong lobe essentially unchanged.

The ordinary annular baseline stacks all supported strip pairs into one vector at each shift and computes a single centered, normalized correlation. It uses the same mask, rows, and input resolution as Eq. (2). It can benefit from a clear mark confined to one strip, but a broad band of scratches confined to one radius can also dominate the entire result. Taking a median requires a reasonable fraction of radial strips to favor the same angle. Since adjacent strips can share a physical scratch, the number of strips cannot be read as a sample size.

The sparse baseline detects scale invariant feature transform keypoints on the corrected face [8], limited to the same annulus. We use mutual descriptor matches with a nearest to second-nearest distance ratio below 0.8. A two-pair similarity fit with random sample consensus [17] uses a three-pixel inlier tolerance and at most 500 hypotheses. We refit on the largest consensus set with at least four inliers and read rotation from the fitted similarity. This baseline can recover a mark displaced across two polar strips. Insufficient matches produce a zero estimate, included in the success rate.

A final baseline correlates the radius of the external contour as a function of angle. It retains the faint ovality discarded by an ideal circle fit. The two lobes of an ellipse predict bearings separated by half a turn. We score the chosen lobe exactly as returned. For reference, an estimate independent of a uniform true bearing has probability $10/360$ of being within five degrees. The measured offsets in this study are stratified and subject to placement error, so this number is a reference expectation, not a fitted null distribution.

V. Evaluation

The primary endpoint is success after a single wash at the five-degree tolerance. We specified the comparison of strip median with ordinary annular correlation on this endpoint when we closed the development partition. All ablations and other wash counts are descriptive. Using the angular distance on the circle [18], success for method m in lot l, plate i, placement j is

$$\begin{aligned} e &= \hat{\theta}^{(m)}_{lij} - \theta_{lij}, \\ D &= \left|((e + 180) \bmod 360) - 180\right|, \\ H^{(m)}_{lij}(\delta) &= \mathbf{1}\{D \le \delta\}. \end{aligned} \tag{3}$$

Angles in Eq. (3) are in degrees. We average over six placements per piece and eight pieces per lot. Each of the twelve test lots thus contributes one rate (48 placements), and each contributes equally to the reported estimate. To compare methods we compute the twelve within-lot differences and report their mean plus or minus the two-sided Student multiplier for eleven degrees of freedom times their standard error. This interval describes lot-to-lot variation under an approximation of independent, similarly sampled lots. It conditions on the acquisition apparatus. It cannot quantify variation among stockrooms, of which we used only three. In particular, applying an independent binomial error to 576 trials would count repeated use of the same photograph as fresh enrollment each time.

The confidence margin is the winning score minus the largest score more than five degrees away in either direction on the circle. Development images set a margin threshold of 0.040, the smallest value on a grid spaced by 0.005 to give at least 95% success among accepted one-wash placements. We used no test images to adjust the threshold. Coverage is the fraction with margin at least the chosen value. Conditional success counts only accepted angles, and unconditional success of the return procedure counts accepted correct angles against the full number of placements. A method that declines every case can achieve zero angular mistakes at zero coverage. These quantities keep the cost of that choice visible.

VI. Results

A. Rotation after one wash

The strip median placed 462 of 576 one-wash queries within tolerance. Table II gives unrounded counts. The corresponding mean lot rate was 80.2% (95% interval 73.4 to 87.0), compared with 66.1% (60.2 to 72.1) for ordinary correlation. Their paired difference was 81 placements or 14.1 percentage points (9.5 to 18.7).

TABLE II
Test placements within five degrees of the measured rotation, out of 576 at each wash stage. The same prewash reference is used in all three columns.

Estimator	Dry reseat	One wash	Ten washes
External contour	106	102	97
Local features	447	329	279
Annular correlation	507	381	314
Strip median	532	462	397

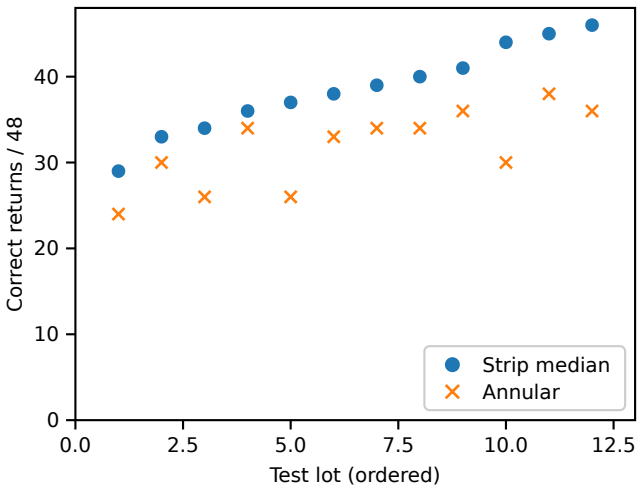


Fig. 1. Correct one-wash returns per lot, out of 48 placements. Lots are ordered by the strip median result. Both series at a given position refer to the same lot.

The difference between the two dense methods varied across lots (Fig. 1). In the smoothest-looking cartons the image often carried little evidence at any radius. In two of the more heavily striated cartons, ordinary correlation tended to align a bright arc near the well even when the neighboring strips disagreed. One particularly good carton yielded 46 correct returns out of 48 with the strip median. Two failures in that carton involved the same plate. A carton-level success rate near one therefore offers limited assurance about the plate remaining on the table.

The tolerance curve in Fig. 2 shows the position of the five-degree result relative to smaller and larger errors. It rises rapidly on the first few degrees and slowly beyond ten. Moving the tolerance from five to twenty degrees recovered another 28 strip-median cases. Even twenty degrees permits an arc displacement of 38 mm at 110 mm radius. Raising the limit is an expensive way to admit additional pieces if the placement specification was written in millimeters.

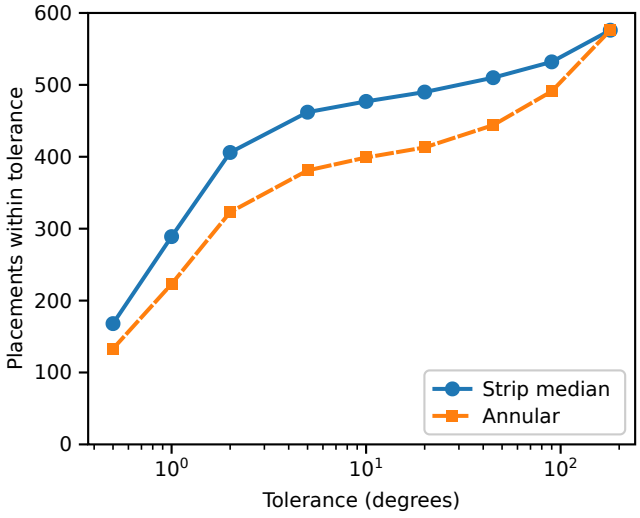


Fig. 2. Cumulative correct placements after one wash (counts out of 576). Lines join the measured tolerances. The horizontal scale is logarithmic.

The feature baseline occasionally gave a better bearing than either dense method. Several pieces had a small inclusion near the rim, far enough from the center for its angular position to be well determined. Two or three sharp points on nearby scratches can then determine a similarity transformation. When the well had a homogeneous glaze, the median over strips often had too few responding strips for the inclusion to influence the maximum. Across all pieces the sparse method paid a substantial cost for those on which it could fit no consensus. Its overall score includes the resulting default estimates as specified in Section IV.

B. Retaining the original reference

From dry reseating to the first postwash series the strip median lost 70 correct placements, and it lost a further 65 by the tenth wash (Table II). Camera setup and elapsed time changed along with wash history. We therefore treat the columns as operating conditions, and do not estimate a per-cycle rate of surface erosion. Individual inspections found that several apparent pits on first viewing had actually been loose spots. Conversely, a mottled area visible through the glaze often survived ten passes with little change in azimuthal profile.

We also compared each tenth-wash query against the *first* one-wash view of the same plate. The correct rotation for this auxiliary pair was recomputed from the two corresponding lower views. This pairing recovered 430 of the 576 relative bearings with the strip median. It is a different target. A new reference on its own measures displacement since the first wash. Recovering the original table heading from it also requires carrying forward the estimated displacement during the first wash. A bad estimate on the intervening view would become part of the stored heading and could persist through otherwise successful registrations.

C. Radial support and abstention

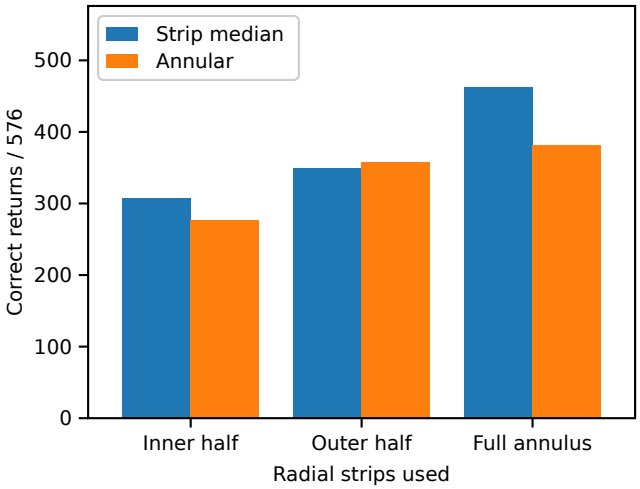


Fig. 3. Successful one-wash placements at five degrees for three radial supports. Each bar counts all 576 trials. Inner and outer halves each contain 12 of the original 24 strips.

We reran both dense scores using the inner twelve or outer twelve strips alone (Fig. 3). The minimum support was half the number of retained strips, as in the full method. On the outer half the ordinary score yielded 358 successes against 349 for the strip median. The improvement of the median on the complete annulus thus did not extend to every band in isolation. The ledge carried isolated high-contrast marks on several pieces, a setting in which retaining the weight of a strong strip is reasonable. The half-band results reuse the test images and were examined after the primary analysis.

At the fixed margin of 0.040 the estimate was returned on 420 placements, of which 399 were within tolerance (Table III). This gives 72.9% coverage and 95.0% conditional success. In a sequence requiring an instruction at every place the method supplies an accepted correct return on 69.3% of placements. The 21 accepted mistakes include matches with a single prominent lobe far from the true heading. Their separa-

tion from the next-best lobe is evidence for a preference *within the available images*. It says little about whether the prominent feature was present on the dish before washing.

TABLE III
One-wash decisions at the margin fixed on development data. Correctness uses the forced estimate even on queries where the instruction is withheld.

Margin decision	Within 5 degrees	Outside 5 degrees	Total
Return angle	399	21	420
Withhold angle	63	93	156
All placements	462	114	576

VII. Controls and error inspection

For a check on fixtures and carton-level cues we assigned each query the reference of the next plate in packing order within its own lot, wrapping at eight. The target for scoring this control remained the query’s displacement from its *own* enrollment. There is no corresponding point between unrelated dot pairs. The question here is whether a method given the wrong enrollment can still guess the required return angle, for example from the way the support was seated. At one wash the strip median yielded 18 apparent successes in 576 comparisons (3.1%). The expected count under a uniform guess is 16. A particular pair of molds may share a direction, but this control offers little support for using a spare plate as a substitute reference.

In a second control we replaced every angular sample in each radial row by that row’s mean, with a complete valid mask. All shifts then have zero within-row contrast after filtering. The documented default of zero was within tolerance on 15 one-wash queries. This control specifically checks the fallback, the sign of wrapping at zero and use of a complete denominator. It does not test for information leakage in contour centering, which has already taken place by the time we apply the replacement. The wrong-reference control uses the original pixel arrays and does include centering.

A half-turn error has several possible causes. After a narrow feature has been lost, an oval lip can still recover a major axis. This is a line, with two admissible senses. In other cases the strongest pair of residual dark streaks fell on opposite sides of the dish. Folding the error at 180 degrees would count such pairs as successful replacements. That treatment would be appropriate for a rim specification expressed as the direction of an axis. It would place the material point opposite its prescribed seat in ours. We found no justification for changing the loss on the basis of a second maximum seen after the wash.

We reviewed signed errors greater than 45 degrees alongside their score curves, revealing lower-camera labels only after selecting the images. Fixed reflections partly survived the polarization filter near rolled edges. When a circular highlight remained stationary in the camera view, the dominant apparent angle could be close to zero despite a large material rotation. Some narrow, well-separated peaks instead followed a water mark that crossed most radial strips. The score can support the wrong lobe in both situations. The examples were chosen by their errors and cannot estimate the fraction attributable to each optical cause.

VIII. Discussion

Restoring a plate relative to its enrollment frame is possible on a substantial fraction of this stock, including pieces that a catalog photographer would reasonably present in a single view. The information resides in differences among nominally identical surfaces. An estimator trained to ignore instance differences while identifying the shape would lose exactly the evidence required for a return. Conversely, an exact mark on the underside supplies an excellent audit direction even on a plate whose upper face carries very little angular information.

The results depend on the scale and organization of the clearing operation. Identity was provided. We fixed the camera close to the support normal and moved a single clean plate into its field. A table photograph taken across several place settings would provide fewer plate pixels and an oblique view of the far rim. The upper lip can hide part of the well at such angles. A homography of the table plane cannot recover that missing portion. In service there may be utensils across the plate at enrollment and residue after use. Neither intervention was represented in our test. Likewise the bottom stamp was essential to the audit but inaccessible to the matching method. Permitting an operator to turn the plate over for both photographs would give a different and much easier task.

The production lots came from stock that the three stores could still supply by the carton. Pieces retained by a household over many years can acquire stable knife marks, while a replacement in the same service can be almost featureless. Our wash series lasted ten cycles, which limits what it tells us about either state. The decline between columns in Table II is compatible with removal of surface marks and with small differences in illumination on the later days. The calibration checks constrain an error in the angular reference. They do not measure how a light field interacts with a slightly tilted glossy surface.

A system that records only the outline of each place cannot later report a previous material bearing. Where the exact bearing is part of the clearing instruction, an enrollment view and a policy for unassigned angles therefore belong in the initial record. For a plate on which the evidence has washed away, the old place remains well defined in translation. The outstanding entry on its return sheet is the angle.

Appendix A. Angular audit

The lower image contains two circular dots of unequal diameters inside the recessed foot. We locate their intensity centroids after local background subtraction. The bearing is the angle of the line from the smaller dot to the larger one. A change in camera distance alters the length of the line without reversing it. We kept the dots more than 35 mm apart on every plate. Measuring the pair by a principal axis without recording which end is which would introduce a 180-degree reference ambiguity of its own.

A separate lapped disk carried a line visible from both sides through two narrow windows. Before and after each acquisition day we placed this disk at twelve encoder readings spaced by 30 degrees with an independently chosen origin. We converted each side to table coordinates and subtracted the measured bearings. The largest absolute discrepancy was 0.43 degrees. Repeating the same upper measurement on the disk after recalibration gave a median absolute discrepancy of 0.16 degrees across readings. These checks apply to the calibration disk on its support. A chipped foot could introduce a tilt specific to a production piece. Our qualification measurement on each test plate was simply a stable seat under its own weight.

Because the two reference dots remain on the same opaque body, their absolute azimuth relative to the unmarked upper surface never enters Eq. (3). There is no occasion to define a common zero across pieces in a carton. Subtracting the bearing of the lower view of plate A from the lower view of plate B, even with exact camera calibration, would supply an arbitrary number. This is the reason for retaining the original target when constructing the wrong-reference test in Section VII.

Appendix B. Numerical and acquisition details

The circular score was evaluated in double precision. At 720 bins the indices 718, 719, 0, 1 and 2 form consecutive samples across the chosen origin. Padding the row at its ends as for linear correlation splits a scratch that crosses the nominal diner’s direction. Each pair of real-valued images and each pair of valid masks must undergo the same wrap convention. As a check on sign we moved the upper-camera calibration disk clockwise by 30 degrees. The score maximum moved by 60 bins with the conventions of Eq. (2). A return instruction from the test image therefore called for a counterclockwise turn of 30 degrees.

The estimator has no fitted weights. Development consisted of evaluating strip widths, angular running means and the margin on images from the four reserved lots. We held the same upper references throughout this search. All production timings include circle fitting, feature extraction and the angular search for both images, even though the reference could be stored after enrollment. On a single 3.4 GHz CPU core the median pair time was 0.18 s for the strip median and 0.14 s for ordinary correlation over the 576 primary pairs. Timing excludes camera transfer and the wait for the ware to dry. No accelerator was used.

For reproduction of the lot interval, the one-wash counts for the strip median in Fig. 1 order are 29, 33, 34, 36, 37, 38, 39, 40, 41, 44, 45 and 46. The corresponding annular counts are 24, 30, 26, 34, 26, 33, 34, 34, 36, 30, 38 and 36. Subtract the counts within each position before dividing by 48. Sorting the two columns independently would preserve their means and alter the standard error of their difference. The totals in Table II and the one-wash margin table count a placement exactly once per method, including unsuccessful returns and every zero supplied by a fallback rule.

Acknowledgments

We thank Aila V. Senot for preparing the two-sided calibration disk and recording the stockroom lots before the cartons were opened.

References

[1] T. Hodan *et al.*, “BOP: Benchmark for 6D Object Pose Estimation,” in *Proc. European Conference on Computer Vision*, 2018, pp. 19–34.

[2] T. Hodan, D. Barath, and J. Matas, “EPOS: Estimating 6D Pose of Objects With Symmetries,” in *Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 11703–11712.

[3] F. Manhardt *et al.*, “Explaining the Ambiguity of Object Detection and 6D Pose From Visual Data,” in *Proc. IEEE/CVF International Conference on Computer Vision*, 2019.

[4] T.-C. Hsiao, H.-W. Chen, H.-K. Yang, and C.-Y. Lee, “Confronting Ambiguity in 6D Object Pose Estimation via Score-Based Diffusion on SE(3),” in *Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024, pp. 352–362.

[5] B. Meden, A. Brazi, F. Mayran de Chamisso, S. Bourgeois, and V. Lepetit, “BOP-Distrib: Revisiting 6D Pose Estimation Benchmarks for Better Evaluation under Visual Ambiguities,” in *Proc. IEEE/CVF Winter Conference on Applications of Computer Vision*, 2026.

[6] B. S. Reddy and B. N. Chatterji, “An FFT-based technique for translation, rotation, and scale-invariant image registration,” *IEEE Transactions on Image Processing*, vol. 5, pp. 1266–1271, 1996.

[7] M. Guizar-Sicairos, S. T. Thurman, and J. R. Fienup, “Efficient subpixel image registration algorithms,” *Optics Letters*, vol. 33, pp. 156–158, 2008.

[8] D. G. Lowe, “Distinctive Image Features from Scale-Invariant Keypoints,” *International Journal of Computer Vision*, vol. 60, pp. 91–110, 2004.

[9] V. Balntas, K. Lenc, A. Vedaldi, and K. Mikolajczyk, “HPatches: A Benchmark and Evaluation of Handcrafted and Learned Local Descriptors,” in *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, 2017.

[10] S. Zagoruyko and N. Komodakis, “Learning to Compare Image Patches via Convolutional Neural Networks,” in *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, 2015.

[11] M. Sorel-Davenne and T. Ivasko-Renn, “Bearing repeatability from pit constellations on fired glaze coupons,” *Annals of Surface Inspection*, vol. 14, pp. 63–79, 2021.

[12] A. I. Tersen and N. S. Rivelo, “The stability of radial signatures in circular glaze specimens after rinsing,” *Transactions on Applied Optical Measurement*, vol. 27, pp. 108–122, 2023.

[13] S. Varrin-Leduc, *Cataloguing Undecorated Refectory Ware*. Colline Technical Editions, 2022.

[14] Z. Zhang, “A flexible new technique for camera calibration,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 22, pp. 1330–1334, 2000.

[15] J. Canny, “A computational approach to edge detection,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. PAMI-8, pp. 679–698, 1986.

[16] A. V. Oppenheim, R. W. Schafer, and J. R. Buck, *Discrete-Time Signal Processing*. Prentice Hall, 1999.

[17] M. A. Fischler and R. C. Bolles, “Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography,” *Communications of the ACM*, vol. 24, pp. 381–395, 1981.

[18] K. V. Mardia and P. E. Jupp, *Directional Statistics*. Wiley, 2000.