Applied Artificial Intelligence and Computing Systems
https://www.scigate.org
Science Gate Academic
Vol. 5, No. 6, 2026

Article

Rotational Symmetry and Labeled Evidence in a Digital Heritage Registration Model

Alva Kestrel*, Nerin Vos, Ione Marwick, Sabel Orr

School of Computer Science, University of Adelaide, Adelaide, Australia

*Corresponding author: alvakestrel243@outlook.com

Abstract: Geometric alignment can be precise while absolute orientation remains unidentified. We construct an eightfold-symmetric, urn-like surface to separate these two issues in a yaw-only digital-heritage registration model. An original 1,920-point surface and a radial matching objective exhibit eight exactly equivalent orientation branches. An executed 400-trial branch-selection emulator compares a local branch chosen near the initialization with selection aided by a noisy labeled marker angle. Mean absolute yaw error is 90.22 degrees without branch evidence and 0.0985 degrees with marker selection; within-one-degree frequencies are 13.25% and 100%. These numbers arise from an explicit branch-and-noise model, not from running iterative closest point or a trained registration network. A three-dimensional surface rendering, polar symmetry orbit, repeated objective minima, and circular error distribution make the ambiguity visible. The analysis explains why geometric equivariance does not by itself identify an absolute branch and why extra labels change the information available to an estimator. No archaeological scan, artifact provenance, or full six-degree-of-freedom reconstruction is claimed.

Keywords: digital heritage; rotational symmetry; point-cloud registration; pose ambiguity; labeled correspondence

1 Geometry can align what it cannot name

A repeated ornament can be matched to several equally plausible locations on the same object. When the entire geometry repeats under a rotation, minimizing a geometric residual may identify the shape's alignment only up to that rotation. The numerical fit can be excellent while the reported absolute pose is wrong relative to a chosen reference. This distinction matters for digital heritage, where an orientation can carry cataloging, restoration, or comparison meaning beyond surface overlap. The present paper studies the distinction in an exact generated shape.

We use an urn-like surface with eight identical radial sectors. Only rotation about its vertical axis is unknown. Translation, scale, tilt, partial overlap, missing fragments, and perspective are excluded. These restrictions make the symmetry group and equivalent poses explicit. The shape is not a scan of an archaeological object and has no historical attribution. Its visual resemblance to a vessel supplies an application context, while its mathematical construction supplies the evidence. It should not be described as a reconstruction of a real artifact.

The paper makes two related claims. First, a geometry-only objective has eight equivalent yaw branches by construction. Second, an additional labeled marker angle can select among those branches in a specified noise model. The first is an exact structural result. The second is an executed branch-selection experiment that assumes a local continuous estimate has already been obtained. It is not a comparison between registration algorithms, and its favorable marker result should not be credited to a newly trained descriptor or a faster optimizer.

2 Registration methods and an information boundary

Iterative closest point [1] and range-image registration [2] provide foundational geometric alignment methods. Least-squares transformation estimation [3] and quaternion absolute orientation [4] solve related problems when correspondences or point pairings are given. Those assumptions matter: a labeled pairing can remove ambiguities that an unlabeled surface cannot resolve. We do not run these algorithms here. Their role is to distinguish numerical estimation of a transformation from the prior question of whether the observed geometry determines a unique transformation.

Fast global registration [5], learned local descriptors [6], point-set learning [7], and learned closest-point representations [8] broaden the available tools. Robust learned matching [9], low-overlap registration [10], coarse-to-fine correspondences [11], and geometric transformers [12] address difficult practical settings. Our generated experiment cannot rank itself against them. A descriptor or attention mechanism may use extra structure, training priors, or appearance, while our geometry-only model intentionally contains none of that asymmetry. A method can make a consistent branch choice without proving that the branch is observable from shape alone.

Spin-image recognition [13], heritage recording [14], fractured-object matching [15], and integrated heritage reconstruction [16] motivate the application context. The survey of geometric symmetry [17] provides the structural perspective. A 2026 study asking whether color resolves registration ambiguity [18] highlights the relevance of additional appearance evidence. We use a simpler, explicit labeled angle rather than a color-registration pipeline. That extra observation changes the statistical experiment, so the resulting improvement must be described as the value of added identity information rather than an equal-input algorithmic victory.

3 A surface with a known symmetry group

The surface has height from zero to 1.5 at twenty sampled levels and ninety-six equally spaced angles around the axis. Its radius combines a smooth height profile with a small eightfold angular modulation. The resulting 1,920 points form the displayed surface. The vertical profile gives the object a nontrivial silhouette, while the angular term preserves exact repetition after every forty-five-degree rotation. Model dimensions are arbitrary. No texture, material response, scan noise model, or archaeological damage is implied by the shaded rendering.

$$r(z,\theta)=0.7+0.22\cos(\pi z)+0.07\cos(4\pi z)+0.04\cos(8\theta) \tag{1}$$

Let the symmetry group contain rotations by integer multiples of pi over four. Applying any member of that group leaves the radius function unchanged. Therefore a rotated observation generated at yaw delta is geometrically equivalent to one generated at delta plus any group angle. The set of equivalent poses is an orbit. A geometry-only estimator can identify that orbit under favorable noise and coverage, but cannot identify which of its members has the chosen absolute label without additional information or a prior that privileges one member.

$$r(z,\theta+\pi/4)=r(z,\theta) \tag{2}$$

This observation is not a failure of optimization. Even an exhaustive global search cannot select a uniquely correct absolute branch from an objective that assigns equal values to all branches. A deterministic tie-breaking rule may return the smallest angle or the pose closest to initialization, but that is a convention. It can be useful in a pipeline if documented, yet it should not be confused with recovering a physically distinguished orientation. A benchmark that assigns one canonical branch must explain why its labels are observable or should use a symmetry-aware pose metric.

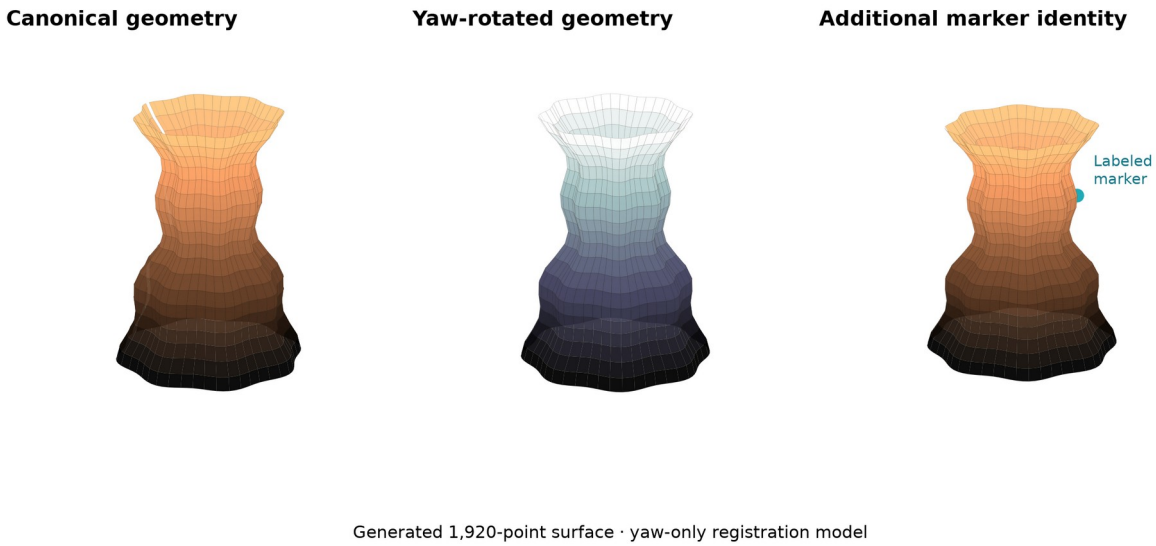


Figure 1. Original generated surface shown in a canonical frame, after a yaw rotation, and with an illustrative labeled marker. Shading is a visualization of the mathematical mesh, not a photograph, material measurement, or heritage scan.

4 The radial objective and its repeated minima

For the illustrated observation, the true yaw is thirty-seven degrees. Independent radial noise with standard deviation 0.002 is added to the generated surface profile. Candidate yaws cover negative pi to pi at 1,440 equally spaced points, corresponding to a quarter-degree spacing. At each candidate, the mean squared difference between its radial profile and the observation is computed. This is a yaw-only radial matching objective, not a nearest-neighbor point-cloud objective and not a full rigid-transformation fit.

$$J(\alpha)=\frac{1}{M}\sum_{m=1}^{M}\left[r(z_m,\theta_m-\alpha)-r_m^{obs}\right]^2 \tag{3}$$

The added radial noise does not remove the objective's periodicity. Every equivalent candidate predicts exactly the same radial array, so it has exactly the same residual against any fixed observed array. Noise can move the best continuous phase within an orbit, but cannot distinguish members of the orbit when the candidate family retains exact symmetry. This is a useful distinction from approximate symmetry in real objects. Slightly asymmetric geometry might in principle identify a branch, although the identifying signal could be weaker than acquisition noise or occlusion.

The forty-five-degree repetition also explains why local precision and absolute accuracy should be reported separately. A fit could estimate yaw modulo the symmetry group to a tenth of a degree and still differ from a canonical label by ninety or one hundred thirty-five degrees. Averaging such errors without describing symmetry would make a locally precise method appear simply inaccurate. Conversely, reporting only residual overlap would hide an absolute-label problem relevant to a restoration task. The metric should follow the intended use of orientation.

5 An explicit branch-selection experiment

The repeated-trial calculation is separate from the radial-grid illustration. In each of 400 trials, a true yaw is uniform over the full circle. Equivalent branches are constructed analytically. The geometry-only local estimate selects the branch closest to zero and adds a small continuous angular error with standard deviation 0.002 radians. This emulates a local initialization convention; it does not execute an

iterative closest-point solver. A noisy labeled marker angle equals the true yaw plus Gaussian noise with standard deviation 0.006 radians.

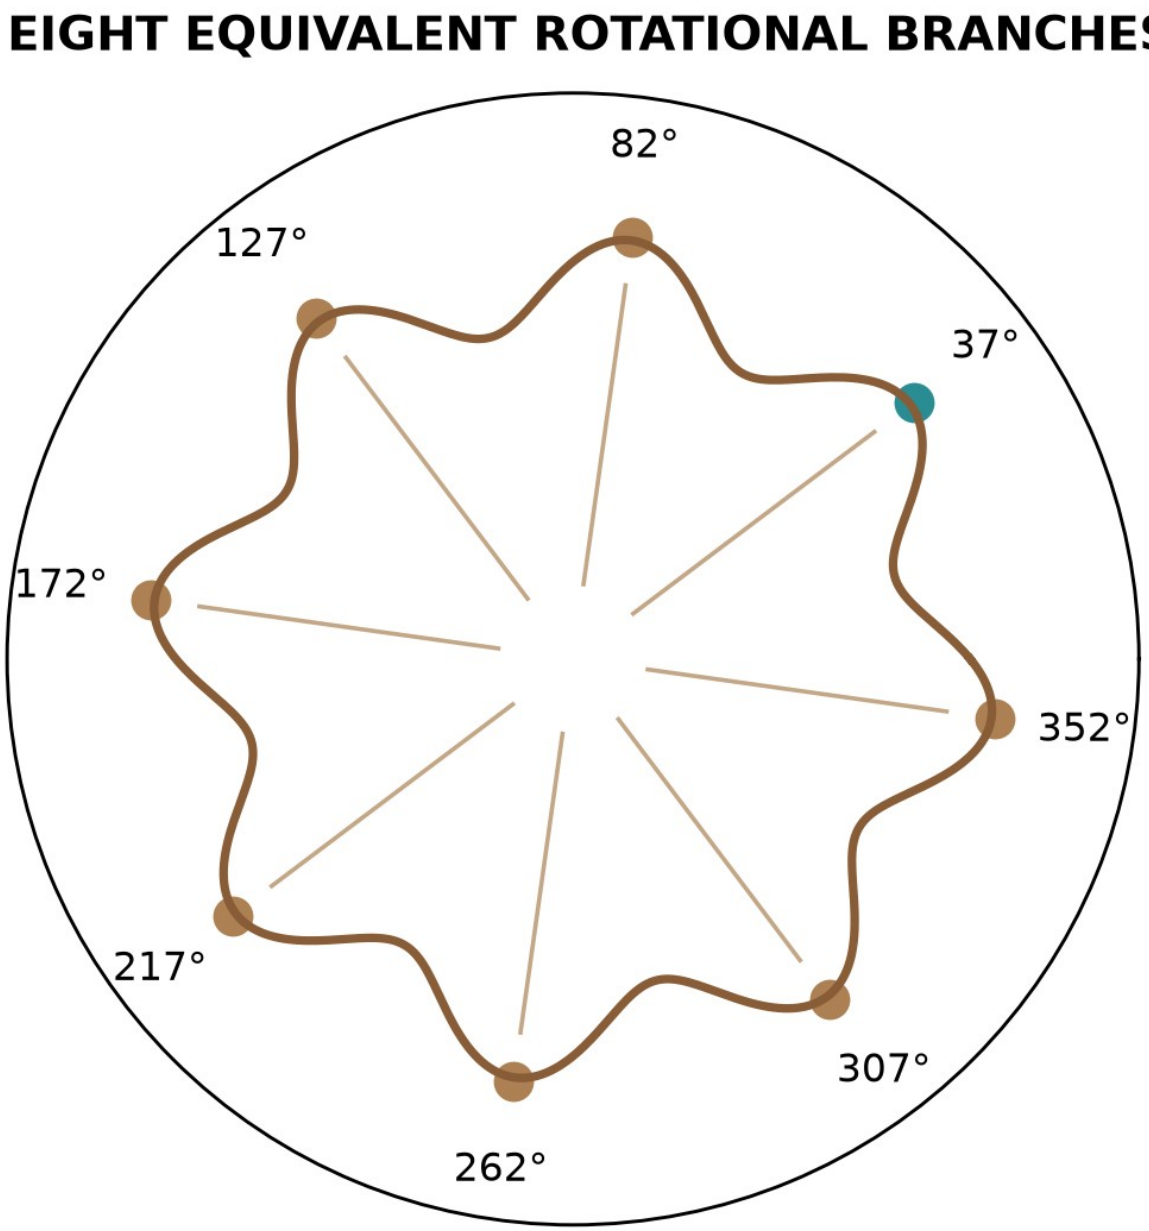


Figure 2. Polar diagram of the eight equivalent yaw branches for the illustrated rotation. Colored positions identify equivalent geometric orientations; the highlighted branch denotes an externally labeled reference, not a branch preferred by the shape itself.

The marker-assisted estimate enumerates eight branches around the same local continuous estimate and selects the one closest to the marker angle under circular distance. The marker chooses a discrete branch; it does not replace the continuous local angle. Consequently, the final fine-scale error comes mainly from the 0.002-radian local error when branch selection is correct. This distinction explains why the reported mean pose error can be smaller than the marker's own noise scale. The method uses the marker to answer a coarser identity question.

$$d_{S^1}(a,b)=|atan2(\sin(a-b),\cos(a-b))| \tag{4}$$

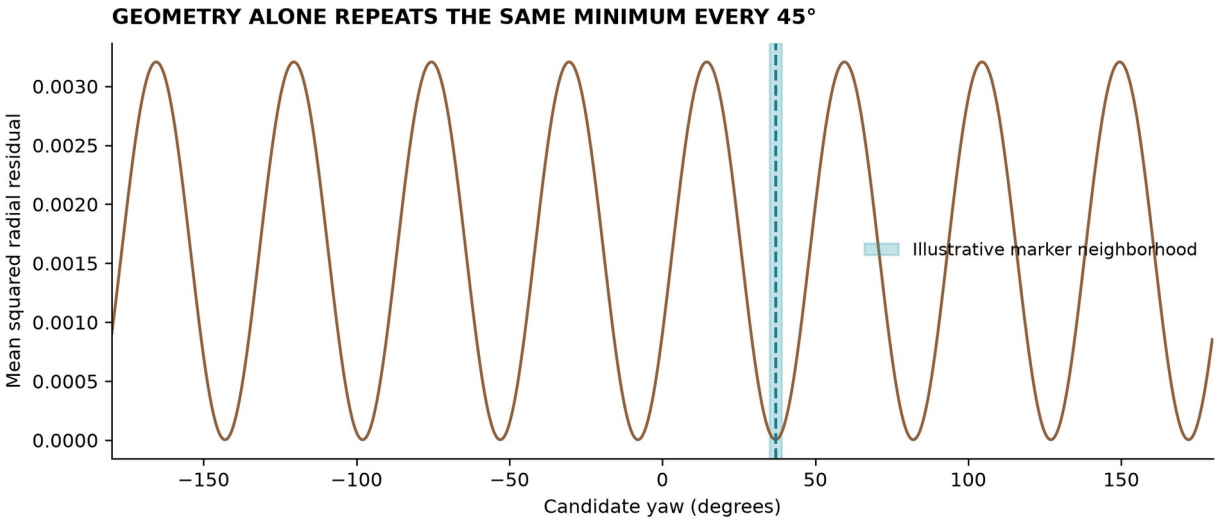


Figure 3. Actual radial matching objective for the generated noisy observation. Its period is forty-five degrees because the candidate model is exactly symmetric. The highlighted marker neighborhood illustrates additional branch evidence and is not part of the geometry-only residual.

Table 1. Branch-selection emulator, 400 trials, seed 3901.

Information supplied	Mean absolute yaw error	Within 1 degree	Continuous local-noise SD
Geometry branch near zero	90.2198°	13.25%	0.002 rad
Geometry plus labeled angle	0.0985°	100.00%	0.002 rad

The geometry-only within-one-degree frequency is close to one eighth, as expected when a uniform true orientation is assigned to one of eight equivalent branches near a fixed initialization. Its large mean absolute error primarily reflects branch selection, not poor local fitting. The marker-assisted policy selects the intended branch in all realized trials, leaving mean error 0.0985 degrees. A zero observed branch-failure count in 400 favorable trials is not a proof of zero failure probability. Marker mislabeling, occlusion, or much larger angular noise is not included.

6 Three distinctions for learned registration

Rotation equivariance concerns how a representation changes when its input rotates. It does not necessarily identify a canonical orientation for a symmetric input. An equivariant representation of an exactly symmetric object must remain consistent with the object's stabilizing transformations. A network can learn a dataset convention or exploit imperceptible asymmetries, but its success then depends on information beyond the ideal symmetric geometry. This distinction should be considered when interpreting learned registration results on repeated manufactured or heritage shapes.

Correspondence identity is also different from geometric proximity. A point at one repeated ornament may be geometrically compatible with several sectors. A label that says which ornament is which changes the correspondence problem. The marker in our experiment supplies exactly that kind of identity at the branch level. It is deliberately privileged information, not a feature discovered by a model. A practical study would need to show how such an identity is acquired reliably, whether it survives scanning and lighting changes, and what happens when it is wrong.

Finally, canonical pose is a task definition. For visualization, any symmetry-equivalent branch may be acceptable. For aligning a repaired fragment to a cataloged orientation, one branch may matter. A

symmetry-aware error would judge all equivalent branches as correct, while an absolute error would not. Neither metric is universally superior. Reporting both can separate geometric registration quality from semantic orientation recovery. This paper uses absolute error to expose the information gained from the marker and derives the equivalent-orbit structure explicitly.

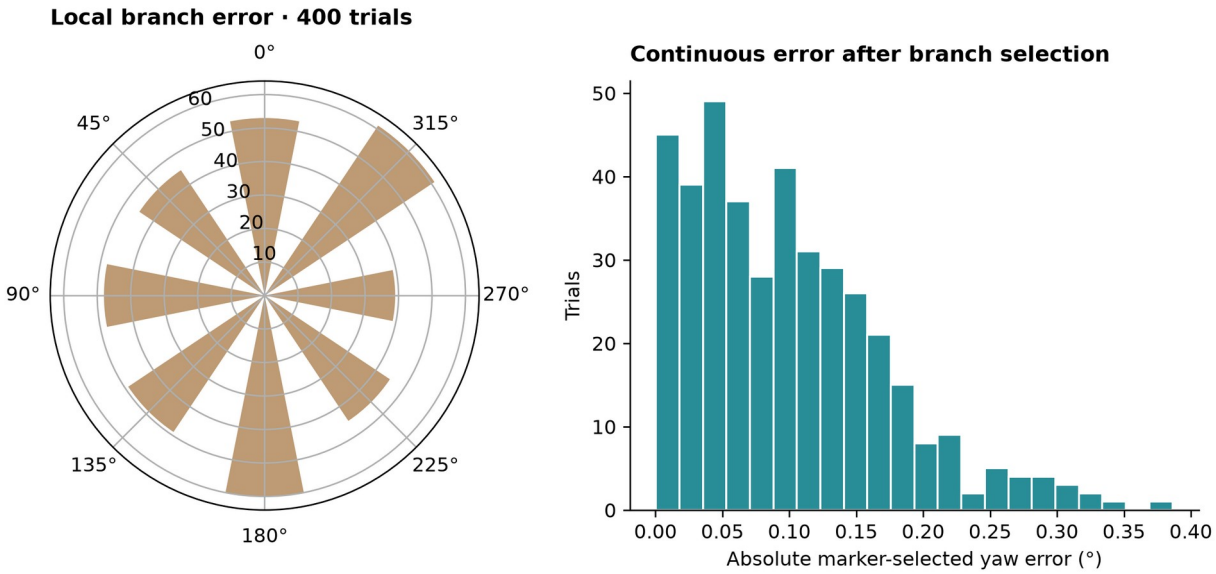


Figure 4. Circular distribution of signed local-branch error and histogram of absolute error after marker-based branch selection. The left panel reveals discrete symmetry branches; the right resolves continuous residual error on a much smaller angular scale.

7 Limits for heritage practice

Real heritage objects rarely have exact analytic symmetry, complete overlap, or a fixed vertical axis. Surface wear, inscriptions, broken handles, and acquisition artifacts may all break symmetry. Some asymmetries are useful evidence; others are noise or reconstruction errors. A real registration study would need to distinguish them and evaluate repeated scans with known reference transformations. The present model omits those complexities, so it cannot establish how many markers are needed for a museum object or whether a particular scanning method preserves identifying detail.

A single angular marker is sufficient here because only yaw remains unknown and the candidate orbit is already known. Full six-degree-of-freedom alignment generally requires more geometric and correspondence information. One point does not uniquely fix a rigid transformation in three dimensions, and a marker angle assumes a reference axis has been established. We state these assumptions to prevent the elegant one-marker result from being generalized beyond its dimensionality. Translation, tilt, and scale are fixed by the experiment rather than estimated successfully.

The marker's noise is independent and correctly centered, and its identity is always correct. A mislabeled marker could confidently select the wrong branch while the geometric residual remains excellent. This is a particularly important boundary because adding semantic evidence can create a new failure mode alongside its benefit. A realistic system should check consistency among multiple cues and allow unresolved orientation when identity is uncertain. No such reliability mechanism is implemented in the branch emulator, so its favorable result should remain conditional on trustworthy labels.

The visual design makes those conditions visible. The three-dimensional rendering identifies the generated geometry. The polar orbit states the set of equivalent poses. The objective plot shows why a global search alone cannot choose among them. The two error panels separate branch ambiguity from continuous precision. A photorealistic artifact image would imply provenance and measured detail that

the experiment does not have. The mathematical surface is therefore both a more honest and a more informative scientific figure.

8 The geometry of an unresolved orientation

8.1 The object is an equivalence class under rotation

Let a shape be represented by an unlabeled set of points and let a rotation act on every point. If a nonidentity rotation maps the set to itself, that rotation belongs to the shape's stabilizer. For the generated vessel, the stabilizer contains eight yaw rotations. An observation of the rotated set therefore identifies a pose only modulo that stabilizer. The natural geometric target is an equivalence class of rotations rather than one absolute angle. This statement concerns the information in the input, independent of a particular optimizer or descriptor.

A benchmark can still assign a single angle as ground truth, but that label may encode a convention unavailable from the shape. An algorithm trained on such labels may learn the convention from sampling artifacts, texture, or asymmetric details. It can be useful operationally, yet the source of information should be identified. Otherwise, a reported pose error can mix geometric reasoning with exploitation of a dataset-specific orientation cue. The exact symmetry model removes such cues so that the distinction becomes unavoidable.

The generated sampling itself preserves the eightfold structure. Ninety-six angular positions contain twelve samples per sector, so a forty-five-degree yaw permutes the angular columns. The height profile is unchanged by yaw. Thus the discrete point set, not only the continuous radius equation, has the intended symmetry under the sector rotations. This matters because careless sampling can break an exact analytic symmetry and give a model a numerical cue that was not intended. Our construction aligns the discrete sampling with the symmetry order.

8.2 A global optimum can remain ambiguous

Optimization quality and identifiability are different properties. A global optimizer can return a minimum of the radial objective, but every symmetry-equivalent branch has the same objective value. Improving search precision can locate each minimum more accurately without ranking them by absolute identity. The objective plot makes this visible through repeated equal minima. A low residual therefore supports geometric consistency with an orbit, not uniqueness of a canonical pose. This is a structural limitation of the measurement model rather than evidence that the optimizer is poorly implemented.

The same issue can appear in a nearest-neighbor registration objective, although its discrete sampling and correspondence rules may introduce additional effects. We do not assert numerical equivalence between our radial objective and every point-cloud loss. The purpose of the radial construction is to isolate an exact symmetry argument with a one-dimensional search. Extending the study to iterative closest point would require specifying correspondences, initialization, stopping criteria, surface sampling, and noise. None of those unexecuted details should be inferred from the urn rendering.

A tie-breaking convention can be documented and useful. Choosing the branch closest to a previous pose can provide temporal continuity, while choosing a catalog-defined branch can support consistent display. Those rules introduce history or metadata. They do not derive an absolute orientation from geometry alone. The branch emulator uses the closest-to-zero convention precisely so that this extra choice is explicit. Its large absolute error reflects disagreement between that convention and a uniformly random canonical label, not an inability to align the repeated shape.

8.3 Expected branch error can be derived

Ignoring the small continuous local noise, the absolute circular errors associated with eight equally likely branch choices are zero, forty-five, ninety, one hundred thirty-five, one hundred eighty, one hundred thirty-five, ninety, and forty-five degrees. Their mean is ninety degrees. A uniformly distributed

true yaw and a fixed near-zero branch convention produce this pattern. The executed mean of 90.22 degrees is close to that analytical reference. This agreement checks the branch construction without requiring an unreported numerical registration algorithm.

The probability of selecting the canonical branch under that convention is approximately one eighth, except for measure-zero boundary ties. With local noise much smaller than one degree, the within-one-degree frequency should therefore be near twelve and a half percent. The observed 13.25% is consistent with ordinary finite-trial variation. It should not be presented as a learned classifier's success rate or a measure of artifact recognizability. It measures how often a convention happens to agree with an independently labeled orientation in the specified emulator.

Once the correct branch is selected, the remaining absolute error is approximately the absolute value of the local Gaussian angular error. Its expected magnitude is the local standard deviation times the square root of two over pi. With standard deviation 0.002 radians, this is about 0.0914 degrees. The observed marker-assisted mean of 0.0985 degrees is near that scale. The calculation explains why a coarse identity cue can yield a final error below its own noise standard deviation: the cue selects a branch while a different estimate supplies continuous precision.

8.4 Marker selection is a classification problem on the orbit

The marker does not need to estimate yaw to a tenth of a degree to distinguish branches forty-five degrees apart. It needs to identify which branch is closest under circular distance. When marker noise is small relative to half the branch spacing, that classification is easy. Our standard deviation of 0.006 radians is deliberately favorable. The experiment therefore tests the principle of adding identity evidence, not the difficult regime in which marker uncertainty approaches the separation between equivalent branches.

If marker noise became large enough, the nearest-branch rule could select an adjacent branch and create a large absolute error while preserving excellent geometric overlap. The error distribution would then combine a narrow continuous component with discrete branch failures. Reporting only mean error could obscure that mixture. A practical evaluation should separate branch accuracy from local orientation precision and should show tail behavior. The present trial histogram resolves the narrow component because no branch failures occurred in the realized favorable sample.

A mislabeled marker is more dangerous than a noisy correct marker in a different way. It can systematically point to another sector rather than fluctuate around the true one. Increasing the number of repeated measurements of that same wrong identity may strengthen the wrong decision. The emulator assumes identity is always correct, so it does not test this failure. This is an important example of how additional semantic information can both resolve a geometric ambiguity and introduce a new dependence on annotation reliability.

8.5 Circular distance avoids an artificial discontinuity

Angles near negative pi and positive pi are close on a circle even though their ordinary numerical difference is almost two pi. The evaluation therefore wraps differences through sine and cosine before taking an absolute value. Without this step, a small boundary-crossing error could be counted as nearly a full rotation. The same circular distance is used for branch selection. Consistent treatment of the circle is necessary for the reported mean and within-one-degree frequency to have their intended geometric meaning.

A circular mean would be appropriate for some orientation summaries, but it can be misleading for a multimodal symmetry orbit. Averaging several equally plausible branches can yield an undefined or arbitrary resultant direction. A set-valued or mixture representation is more faithful in that case. The paper does not collapse the eight branches into a mean pose. It shows them explicitly in the polar orbit and evaluates absolute error only after specifying a branch convention or adding a labeled cue.

The rose plot bins signed local-branch errors, so its angular coordinate carries orientation meaning and its radial coordinate carries trial count. The histogram beside it uses a much smaller linear angular range for continuous residual error. These panels intentionally use different scales because they answer different questions. Comparing bar heights across them would be inappropriate, but comparing the structure of discrete modes with a narrow residual distribution clarifies the two sources of error. Captions state the distinction rather than relying on visual intuition alone.

8.6 Exact and approximate symmetry should not be conflated

A real vessel may have a small inscription, uneven wear, or a missing fragment that breaks exact symmetry. In principle, such details can identify an absolute branch. In practice, their signal may be weaker than scan noise, incomplete coverage, or reconstruction artifacts. The problem then becomes one of statistical distinguishability rather than exact nonidentifiability. Our analytic shape is exactly symmetric, so it establishes the limiting structural case. It does not claim that every visually repetitive heritage object is mathematically ambiguous under all measurements.

Approximate symmetry also changes the role of training priors. A learned descriptor may attend to weak asymmetric details more effectively than a generic geometric loss. That could be a real advantage if the details are stable and physically meaningful. It could also exploit scanner artifacts or canonical sampling patterns that disappear on a new collection. A careful benchmark should perturb acquisition conditions and document the asymmetries available to the model. The current generated study supplies no such empirical comparison, so it should not be used to dismiss learned features.

One useful extension would add a controlled asymmetric term with varying amplitude and compare it with radial noise. That would create a transition from exact orbit ambiguity to weakly identifiable absolute pose. Another would remove sectors to model partial overlap, which can either introduce or remove distinguishing information depending on the visible region. These are proposed experiments only. Their purpose here is to show how the exact analysis can guide a richer design without pretending that the richer results have already been obtained.

8.7 Equivariance preserves structure rather than inventing labels

An equivariant model transforms its representation predictably when the input transforms. This can improve data efficiency and geometric consistency. For a symmetric input, however, the representation must remain compatible with the transformations that leave the input unchanged. It cannot assign a uniquely observable canonical branch unless some additional cue or convention breaks the symmetry. This is a statement about information and group action, not a claim that equivariant architectures are incapable of useful registration or that invariant descriptors are inherently inferior.

An invariant descriptor can help match local geometry while deliberately discarding orientation. An equivariant feature can retain orientation relationships. A canonicalization network can choose a pose convention. These functions are related but distinct. The generated experiment does not implement any of them, so it makes no architecture ranking. It provides a test that their outputs can be interpreted against: what is the observable target, what transformations are equivalent, and which extra information permits an absolute label to be recovered?

The distinction is especially relevant when a benchmark counts only one branch as correct. A method that returns a different symmetry-equivalent branch may be penalized despite perfect overlap. Another method may score well by learning an arbitrary canonical orientation from the training distribution. Reporting symmetry-aware and absolute errors together can reveal this difference. The appropriate primary metric depends on whether the application needs geometric merging, semantic orientation, or both. Our absolute metric is chosen to measure the added labeled cue, not to redefine every registration task.

8.8 One marker here is not one point in full rigid space

The model assumes a known vertical axis, fixed translation, fixed scale, and no tilt. Under those restrictions, an angular cue can select among yaw branches. A single three-dimensional point correspondence in an unconstrained rigid problem generally leaves rotational freedom, and an angular observation already presupposes a reference frame. The one-marker result should therefore be understood as a low-dimensional identity demonstration. It does not establish a universal minimum number of fiducials for scanning, robotics, or artifact assembly.

The marker is also not used to estimate a full continuous transformation. It chooses among branches of a local estimate that is assumed to have a known noise distribution. That local precision is supplied by the emulator rather than obtained from an executed optimizer. This is the strongest modeling boundary in the numerical result. The radial-grid plot demonstrates the objective's repeated structure, while the repeated trials demonstrate branch selection under explicit local-error assumptions. Keeping those two calculations separate prevents a false claim of end-to-end registration performance.

A full empirical pipeline would need to estimate the axis, acquire marker correspondences, solve the transformation, and evaluate repeated scans against a credible reference. Each stage could fail independently or interact with the others. A marker that is easy to see in one scan may be occluded in another, and surface damage can change correspondence quality. The current model does not include those processes. Its contribution is to clarify which ambiguity an extra label resolves before a more complicated pipeline is evaluated.

8.9 The rendering should not imply provenance

The vessel is shaded to reveal its surface profile and eightfold sectors. Color and lighting are visualization choices, not measured material properties. The highlighted marker is an illustrative cue corresponding to the branch-selection concept. No museum, excavation, date, or restoration history is attached to the shape. This restraint is scientifically useful because a photorealistic artifact image could imply an empirical case study that does not exist. The original mathematical rendering communicates the modeled geometry directly and avoids borrowing authority from an invented object.

The surface plot also contains a finite sampling mesh. Its purpose is to make the generated point structure inspectable, not to claim a watertight manufacturing model or a measured scan resolution. A different triangulation could change visual appearance without changing the underlying symmetry argument. The numerical objective uses the stated radial arrays, and the branch trials use analytic equivalent rotations. The figure is linked to those data, but it should not be treated as a hidden source of additional geometric measurements.

The polar orbit and objective curve are therefore essential companions to the three-dimensional view. A perspective rendering can make some sectors appear different because of lighting and viewpoint even when the geometry is exactly repeated. The orbit states the equivalence explicitly, and the objective verifies its consequence numerically. This combination prevents visual asymmetry in the drawing from being mistaken for physical asymmetry in the model. The evidence comes from the equation and calculated residual, not from whether a shaded object looks perfectly symmetric to a viewer.

8.10 A reporting standard for ambiguous pose

A pose report should identify the transformation group, the shape's symmetry, the correspondence information, and the metric's treatment of equivalent solutions. It should also distinguish initialization conventions from observed evidence. These details can change the meaning of a registration score as much as the optimizer does. The generated example is small enough to state them completely. A more complex benchmark may need an explicit symmetry annotation or a task-specific equivalence definition rather than assuming that a single canonical transform is always observable.

When absolute orientation remains unresolved, returning several plausible branches can be more honest than returning one confident pose. A downstream task may accept any branch, request another observation, or use metadata to choose. That decision layer should be evaluated separately from geometric fitting. Our emulator forces a single choice to make absolute-error consequences visible, but it does not claim that forced selection is the best practical interface. The exact orbit provides a natural representation of uncertainty that a future system could preserve.

The study's strongest result is thus not the dramatic numerical improvement from ninety degrees to a fraction of a degree. It is the explanation of why that improvement occurs: a labeled cue answers a question the symmetric shape cannot answer alone. The executed branch trials quantify the effect under favorable, explicit noise assumptions. Real heritage validation must then test cue reliability, incomplete geometry, approximate symmetry, and full pose estimation. That sequence preserves the distinction between structural reasoning and empirical performance.

8.11 The marker supplies identity rather than more geometry

An ideal label distinguishing eight equally likely branches carries three binary digits of branch information. This is an information-counting observation, not a statement that three arbitrary binary sensor readings will reliably solve the problem. Noise, label errors, and prior probabilities change the effective information. The marker angle in the emulator supplies a highly reliable branch cue because its uncertainty is small relative to branch spacing. The dramatic improvement should therefore be credited to an informative additional observation, not to an unexplained gain in geometric precision.

The cue can be coarse and still useful. A broad angular sector narrower than the separation between competing branches could identify the correct member while leaving local fitting to determine the fine angle. Conversely, a very precise measurement of an unlabeled repeated feature could remain ambiguous because it does not say which sector was observed. Precision and identity are separate dimensions of evidence. The experiment chooses a noisy but correctly labeled angle to isolate that distinction and keeps the continuous local error unchanged across policies.

If prior orientation were concentrated near zero, the local convention could perform much better without any marker. That improvement would come from the prior distribution of true yaw, not from a change in the geometry. Our uniform prior prevents such a preference and makes all branches equally relevant. A benchmark should state whether initial poses are uniformly distributed, near aligned, or drawn from an application-specific range. Initialization difficulty and absolute identifiability can otherwise become mixed in a single registration score.

8.12 A symmetry-aware metric would answer another question

A symmetry-aware yaw error minimizes circular distance over all equivalent branch offsets. Under that metric, the local estimate in our emulator would have only its small continuous error, even when its absolute error is ninety degrees. The marker-assisted estimate would offer little additional geometric benefit because branch identity is deliberately ignored. This is not a contradiction. It shows that the value of a marker depends on whether the downstream task needs a canonical orientation or merely a geometrically equivalent alignment.

An assembly task may need both notions. Surface overlap could be correct up to symmetry, while an inscription or functional feature must face a particular direction. A visualization task might accept any branch. A cataloging task might preserve a metadata convention even when geometry cannot verify it. These examples motivate reporting the task definition explicitly, but they are not empirical applications tested by the generated vessel. The model supplies a language for distinguishing their requirements without pretending to resolve the practical correspondence problems each would introduce.

A combined score could penalize geometric misalignment and semantic branch error separately. Choosing their weights would be a decision about application costs. The present paper avoids such a

composite because it would obscure the mechanism and require values not supplied by the experiment. It reports absolute yaw error, within-one-degree frequency, and the exact orbit structure. A reader can then see which part of the improvement concerns identity and which concerns continuous fitting, rather than relying on an arbitrary blended performance number.

8.13 Reliability cannot be inferred from a zero count alone

All four hundred marker-assisted trials select the correct branch under the favorable noise model. This observed zero failure count does not establish immunity to rare Gaussian tails, marker occlusion, or annotation mistakes. The last two are absent from the generator entirely. A reliability claim would require an explicit failure model and a sample design targeted at its relevant rare events. The current result is a demonstration of a mechanism at one noise scale, with finite-trial performance stated exactly.

The noise scale is also known to the evaluator but not used to construct a calibrated confidence probability. The nearest-branch rule returns a choice without a stated posterior. A future method could combine marker uncertainty with geometric orbit hypotheses to return branch probabilities or abstain when evidence is weak. That extension would need to evaluate calibration and coverage. Simply observing a large score margin in this easy setting would not justify a probability of correctness, especially if identity errors are possible.

The branch experiment is therefore best read as a controlled contrast in available information. Geometry supplies an orbit with local precision; the labeled cue selects its identity. The experiment quantifies that contrast while deliberately omitting acquisition, correspondence discovery, and full rigid pose. Preserving those omissions is essential to its scientific value. They make the source of improvement understandable and identify the empirical work needed before the idea can support a heritage-recording or restoration claim.

9 Conclusion

Exact rotational symmetry creates multiple equally valid geometric alignments. In the generated yaw-only heritage model, a labeled angular cue resolves the discrete branch while preserving the fine precision of a local estimate. The executed numbers describe an explicit branch-and-noise emulator, not a new registration algorithm or a real artifact study. The transferable lesson is to report the symmetry group, the pose convention, the information in correspondences, and the metric's treatment of equivalent orientations before interpreting a low residual or a large absolute pose error.

References

[1] P. J. Besl, N. D. McKay, “A method for registration of 3-D shapes,” IEEE Transactions on Pattern Analysis and Machine Intelligence, vol. 14, pp. 239-256, 1992. doi: 10.1109/34.121791.

[2] Y. Chen, G. Medioni, “Object modelling by registration of multiple range images,” Image and Vision Computing, vol. 10, pp. 145-155, 1992. doi: 10.1016/0262-8856(92)90066-c.

[3] S. Umeyama, “Least-squares estimation of transformation parameters between two point patterns,” IEEE Transactions on Pattern Analysis and Machine Intelligence, vol. 13, pp. 376-380, 1991. doi: 10.1109/34.88573.

[4] B. K. P. Horn, “Closed-form solution of absolute orientation using unit quaternions,” Journal of the Optical Society of America A, vol. 4, pp. 629, 1987. doi: 10.1364/josaa.4.000629.

[5] Q. Y. Zhou, J. Park, V. Koltun, “Fast Global Registration,” Lecture Notes in Computer Science, pp. 766-782, 2016. doi: 10.1007/978-3-319-46475-6_47.

[6] A. Zeng, S. Song, M. Niessner et al., “3DMatch: Learning Local Geometric Descriptors from RGB-D Reconstructions,” 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 199-208, 2017. doi: 10.1109/cvpr.2017.29.

[7] C. R. Qi, H. Su, K. Mo, and L. J. Guibas, “PointNet: Deep Learning on Point Sets for 3D Classification and Segmentation,” 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 77-85, 2017. doi: 10.1109/cvpr.2017.16.

[8] Y. Wang, J. M. Solomon, “Deep Closest Point: Learning Representations for Point Cloud Registration,” arXiv preprint arXiv:1905.03304, 2019. [Online]. Available: https://arxiv.org/abs/1905.03304.

[9] Z. J. Yew, G. H. Lee, “RPM-Net: Robust Point Matching using Learned Features,” arXiv preprint arXiv:2003.13479, 2020. [Online]. Available: https://arxiv.org/abs/2003.13479.

[10] S. Huang, Z. Gojcic, M. Usvyatsov et al., “PREDATOR: Registration of 3D Point Clouds with Low Overlap,” arXiv preprint arXiv:2011.13005, 2020. [Online]. Available: https://arxiv.org/abs/2011.13005.

[11] H. Yu, F. Li, M. Saleh et al., “CoFiNet: Reliable Coarse-to-fine Correspondences for Robust Point Cloud Registration,” arXiv preprint arXiv:2110.14076, 2021. [Online]. Available: https://arxiv.org/abs/2110.14076.

[12] Z. Qin, H. Yu, C. Wang et al., “Geometric Transformer for Fast and Robust Point Cloud Registration,” arXiv preprint arXiv:2202.06688, 2022. [Online]. Available: https://arxiv.org/abs/2202.06688.

[13] A. E. Johnson, M. Hebert, “Using spin images for efficient object recognition in cluttered 3D scenes,” IEEE Transactions on Pattern Analysis and Machine Intelligence, vol. 21, pp. 433-449, 1999. doi: 10.1109/34.765655.

[14] F. Remondino, “Heritage Recording and 3D Modeling with Photogrammetry and 3D Scanning,” Remote Sensing, vol. 3, pp. 1104-1138, 2011. doi: 10.3390/rs3061104.

[15] Q. X. Huang, S. Flöry, N. Gelfand et al., “Reassembling fractured objects by geometric matching,” ACM Transactions on Graphics, vol. 25, pp. 569-578, 2006. doi: 10.1145/1141911.1141925.

[16] S. F. El-Hakim, J. A. Beraldin, M. Picard et al., “Detailed 3D reconstruction of large-scale heritage sites with integrated techniques,” IEEE Computer Graphics and Applications, vol. 24, pp. 21-29, 2004. doi: 10.1109/mcg.2004.1318815.

[17] N. J. Mitra, M. Pauly, M. Wand et al., “Symmetry in 3D Geometry: Extraction and Applications,” Computer Graphics Forum, vol. 32, pp. 1-23, 2013. doi: 10.1111/cgf.12010.

[18] M. Kresović, T. Theoharis, J. Y. Hardeberg, “Does color resolve symmetry ambiguities in 3D point cloud registration? Literature survey and adversarial testing,” Computers & Graphics, vol. 139, pp. 104696, 2026. doi: 10.1016/j.cag.2026.104696.