

Fast, Robust, Permutation-and-Sign Invariant $SO(3)$ Pattern Alignment

Anik Sarker¹ and Alan T. Asbeck¹

Abstract—We address the correspondence-free alignment of two rotation sets on $SO(3)$, a core task in calibration and registration that is often impeded by missing time alignment, outliers, and unknown axis conventions. Our key idea is to decompose each rotation into its *Transformed Basis Vectors* (TBVs)—three unit vectors on S^2 —and align the resulting spherical point sets per axis using fast, robust matchers (SPMC, FRS, and a hybrid). To handle axis relabels and sign flips, we introduce a *Permutation-and-Sign Invariant* (PASI) wrapper that enumerates the 24 proper signed permutations, scores them via summed correlations, and fuses the per-axis estimates into a single rotation by projection/Karcher mean. The overall complexity remains linear in the number of rotations ($\mathcal{O}(n)$), contrasting with $\mathcal{O}(N_r^3 \log N_r)$ for spherical/ $SO(3)$ correlation. Experiments on EuRoC Machine Hall simulations (axis-consistent) and the ETH Hand-Eye benchmark (`robot_arm_real`) (axis-ambiguous) show that our methods are accurate, 6-60x faster than traditional methods, and robust under extreme outlier ratios (up to 90%), all without correspondence search.

I. INTRODUCTION

Estimating the 3D rotation that aligns one sensor or object frame to another is a fundamental problem in robotics and computer vision. Classically posed as *Wahba’s problem* [1], it underpins many applications: *point-cloud registration* [2], [3], [4]; *Robot Hand-Eye (RHE) / World-Hand-Eye (RWHE)* pipelines that first estimate rotations, then recover translations, often with time-offset handling or unsynchronized logs [5], [6], [7], [8], [9], [10]; *IMU-based wearables and gloves* for motion based sensor-to-segment (S2S) calibration, joint-axis identification, and cross-session re-calibration, where orientation distributions alone can recover mounting offsets and axes without time pairing [11], [12], [13], [14], [15]; *IMU-magnetometer heading calibration*, where yaw alignment is obtained by matching rotation-only statistics while robustly handling hard/soft-iron distortions via ellipsoid-fitting models [16], [17], [18], [19]; *VIO/SLAM initialization and relocation*, where batches of orientations are aligned to refine gravity/yaw and extrinsics before full state estimation [20], [21], [22], [23]; *molecular alignment* [24], and more.

Closed-form or least-squares solutions (e.g., Davenport/QUEST, SVD/Procrustes, and modern quaternion solvers) are mature [25], [26], [27], [28], [29], but they typically assume *paired* measurements (known correspondences) and degrade under heavy outliers or axis-convention mismatches.

A well known strategy treats an $SO(3)$ distribution as a spherical signal and aligns two such signals via FFT-based spherical/ $SO(3)$ cross-correlation [30], a technique with roots

in 3D shape analysis [31], [32] and used in modern spherical CNNs. However, harmonic/correlation pipelines scale as $\mathcal{O}(N_r^3 \log N_r)$ in the rotation-space resolution/bandwidth [33], [34], which can limit practical resolution and scalability [35].

In many real scenarios (e.g., multi-sensor rigs, asynchronous logging, or heterogeneous vendors), *pairwise correspondences are unavailable*, sampling rates differ, and axis conventions may be permuted and/or sign-flipped. A particularly important use case is Robot Hand-Eye / World-Hand-Eye calibration [36], [37], [10], [38], [39], [40], [41], [42], [43], which seeks X (hand→sensor) and Y (world→robot) satisfying

$$A_i X = Y B_i, \quad i = 1, \dots, N, \quad A_i, B_i \in SO(3).$$

Splitting into rotation and translation yields the rotational constraints $R_{A_i} X_R = Y_R R_{B_i}$, typically solved first on $SO(3)$, followed by translations. A common two-stage practice is to ignore nonlinear constraints, solve the overdetermined linear system in least squares, then project the resulting matrices onto $SO(3)$ [37]; translations are then recovered trivially given the rotations. In practice, however, most RHE/RWHE pipelines *assume synchronized, paired poses*; when sampling rates differ, authors time-align to build pseudo correspondences [10], [38], [39], [40], [41], [42], [43].

In this work we target the rotational subproblem only. We estimate, from two *unpaired* rotation sets $\mathcal{A} = \{R_{A_i}\}$ and $\mathcal{B} = \{R_{B_i}\}$, a single $R^* \in SO(3)$ that best aligns $\mathcal{B}$ to $\mathcal{A}$ without time alignment or correspondences. This correspondence-free $SO(3)$ alignment serves as a robust initializer or intermediate relative-frame estimate for modern (possibly certifiable) RWHE solvers [36], while leaving translation estimation outside the scope of this paper.

We recast $SO(3)$ set alignment as *three* correspondence-free alignments on S^2 : each rotation is decomposed into *Transformed Basis Vectors* (TBVs), yielding three spherical point sets per distribution. We then apply fast, robust spherical matchers from prior work [44]—SPMC, FRS, and their hybrid—to each axis and *reconstruct* a single 3×3 rotation by projection onto $SO(3)$. To handle unknown axis conventions, we make the pipeline *Permutation-and-Sign Invariant (PASI)* by exhaustively scoring signed axis permutations and selecting the maximizer (restricting to right-handed cases when appropriate). Fig. 1 provides an overview of one variant of our method.

A. Contributions

Correspondence-free $SO(3)$ formulation with PASI. We align rotation sets via TBVs on S^2 and enforce invariance to axis relabels/sign flips by enumerating signed permutations $L = PS$ and selecting the hypothesis that maximizes a correlation objective [44].

*This work was supported by the National Science Foundation, Grant # 2014499

¹The authors are with the Dept. of Mechanical Engineering, Virginia Tech, Blacksburg, VA, USA. `aniks`, `aasbeck@vt.edu`




Fig. 1. The figure shows how to align two sets of (SO_3) distribution using SO_3 -SPMC. The major steps of the SO_3 -SPMC alignment algorithm in axis consistent setting are depicted through subfigures (a-h). Sub-figures show the progression from one step to another.

(a) Two sets of orientations (SO_3) are presented in two separate blocks. (b) S_2 -distributions of two sets and corresponding mean directions are presented. The darker colors represent the S_2^A, S_2^B, S_2^C , while the lighter colors correspond to S_2^B, S_2^C, S_2^A . (c) In the three-sphere visualization, pairs of $(S_2^{AN}, S_2^{BN}), (S_2^{AN}, S_2^{CN}),$ and (S_2^{AN}, S_2^{BN}) are depicted. (d) The projected S_2 -distributions are visualized using a 2D histogram. The color legends in the histogram maintain the same notation as mentioned earlier, where darker colors represent the relevant 2D histogram of set A , while lighter colors correspond to set B . (e) After performing the 1D cross-correlation, the respective 2D histograms are aligned. (f) The 2D histogram data, represented in geographic coordinates, is converted to 3D Cartesian coordinates. (g) All S_2 -distributions are rotated by corresponding inverse of the $R_{xN}^A, R_{yN}^A, R_{zN}^A$. (h) Block diagram illustrates the process of extracting an set of SO_3 from the aligned $S_2x, S_2y,$ and S_2z -distributions of set B .

Linear-time algorithms with strong robustness. Each spherical matcher (SPMC, FRS, hybrid) runs in $\mathcal{O}(n)$ time [44]; our $SO(3)$ variants require a constant number of calls (3 for axis-consistent; 18 precomputed for PASI), thus remaining $\mathcal{O}(n)$, versus $\mathcal{O}(N_r^3 \log N_r)$ for spherical/ $SO(3)$ correlation [33], [34], [35]. Empirically, the pipeline remains accurate under very high outlier rates (up to 90%) [44].

Extensive evaluation on synthetic and real data. We validate axis-consistent alignment on EuRoC MAV Machine Hall sequences [45] and axis-ambiguous alignment on the ETH Hand-Eye benchmark (robot_arm_real) [10], solving the *rotational* RWHE component directly on $SO(3)$ and comparing against time-aligned RANSAC variants.

II. MATHEMATICAL PRELIMINARIES AND CONCEPTS

A. Rotating 3D Cartesian Points

Let the points in $\mathbf{P} \in \mathbb{R}^{n \times 3}$ and let $\mathbf{R} \in SO(3)$. Applying the same rotation to all points uses right-multiplication:

$$\mathbf{P}_{\text{rotated}} = \mathbf{P} \mathbf{R}. \quad (1)$$

B. Rows vs. Columns of a Rotation Matrix

For the active, column-vector convention $\mathbf{v}' = \mathbf{R} \mathbf{v}$ with $\mathbf{R} \in SO(3)$:

- **Columns:** column j equals $\mathbf{R} \hat{\mathbf{e}}_j$; i.e., the rotated canonical axes expressed in the original frame.
- **Rows:** row i equals $\hat{\mathbf{e}}_i^T \mathbf{R}$, giving components via $(\mathbf{v}')_i = \text{row}_i(\mathbf{R}) \cdot \mathbf{v}$; geometrically, original-frame axes expressed in the rotated frame (since $\mathbf{R}^{-1} = \mathbf{R}^T$).

Example: a z -axis rotation by θ ,

$$R_z(\theta) = \begin{bmatrix} \cos \theta & -\sin \theta & 0 \\ \sin \theta & \cos \theta & 0 \\ 0 & 0 & 1 \end{bmatrix},$$

has first/second columns as the rotated x/y axes, while the first row $[\cos \theta, -\sin \theta, 0]$ extracts the x -component after rotation.

C. Transformed Basis Vectors (TBVs)

Let $R_{A \leftarrow B} \in SO(3)$ map coordinates from frame B to frame A , i.e., $v_A = R_{A \leftarrow B} v_B$. We define *Transformed Basis Vectors (TBVs)* using the rows of $R_{A \leftarrow B}$:

$$\text{TBV}_i \triangleq \text{row}_i(R_{A \leftarrow B})^T = (R_{A \leftarrow B})^T \hat{\mathbf{e}}_i^A, \quad i \in \{x, y, z\}.$$

Thus, $\text{TBV}_x, \text{TBV}_y,$ and TBV_z are the first, second, and third rows of R (transposed), respectively; equivalently, they are the columns of R^T . Geometrically, TBV_i is frame A 's i -axis expressed in frame B . Because $R \in SO(3)$, the TBVs form an orthonormal triad. For clarity, the *columns* of $R_{A \leftarrow B}$ are frame B 's axes expressed in frame A , and the entries satisfy $(R_{A \leftarrow B})_{ij} = \hat{\mathbf{e}}_i^A \cdot \hat{\mathbf{e}}_j^B$. In computation, the row view provides direct component extraction,

$$(v_A)_i = \text{row}_i(R_{A \leftarrow B}) v_B,$$

and we adopt this row-centric convention throughout unless stated otherwise.




Fig. 2. S^2 representation of a normal vector and S^2 representation of the sets of Transformed Basis Vectors (TBVs). (a) In S^2 representation, we view an arbitrary normal vector $\vec{OP}$ as a point P on the surface of the unit sphere; P is the S^2 -point of $\vec{OP}$. Black arrows indicate the reference-frame directions (equivalently, the TBVs of the identity rotation). The red (x), green (y), and blue (z) dots are the S^2 -points of the identity rotation. (b) The $S2_x$, $S2_y$, $S2_z$ distributions are colored red, green, and blue, respectively; their mean directions are shown as arrows.

D. Set of TBVs From Set of Orientations (SO_3)

Given a set of n orientation samples in SO_3 , we store them as an array of rotation matrices of shape $n \times 3 \times 3$. From this array, we extract the TBVs, yielding three TBV sets: the sets of x -TBVs, y -TBVs, and z -TBVs.

E. S^2 Representation and S^2 -point: Normalized 3D Vector Representation on Sphere Surface

The n -sphere framework (S^n), a family of compact manifolds embedded in $\mathbb{R}^{n+1}$, is widely used in computer vision, including viewpoint and surface-normal estimation [46], [47]. On the 2-sphere (S^2), a unit 3D vector is represented by its endpoint on the sphere surface, which we refer to as an “ S^2 -point” [48]. Thus, a normal vector $\vec{OP}$ of unit length corresponds to the point $P \in S^2$; the coordinates of P encode the direction of the original unit vector. The axes of a global reference frame intersect the sphere at three points (on x , y , z), which can likewise be viewed as the S^2 representation of that frame.

F. S^2 Representation of TBVs

Typically, in the n -sphere paradigm, a quaternion (in $\mathbb{R}^4$) is represented on S^3 [48]. Here, we instead use the S^2 representation of normal vectors for orientation-derived features. A single rotation matrix induces three S^2 -points via its x -, y -, and z -TBVs. In Fig. 2(a), the identity rotation yields three S^2 -points at the coordinate axes. For n SO_3 samples, we first extract three TBV sets (Section II-D) and then represent them as three S^2 -distributions on the sphere surface, denoted $S2_x$, $S2_y$, and $S2_z$; see Fig. 2(b).

G. S^2 Point-Pattern Registration

Spherical point-pattern registration seeks the optimal rotation that aligns two S^2 distributions (target and source). Prior work [44] introduced three algorithms— $SPMC$ (Spherical Pattern Matching by Correlation), FRS (Fast Rotation

Search), and a hybrid $SPMC_FRS$ —for robust alignment of spherical signals. These methods are correspondence-free and robust to moderate noise as well as extremely high outlier ratios. Unlike traditional solutions to the Wahba problem [49], [50], [51], [52], [53], [26], [54], algorithms such as $SPMC_FRS$ do not require pairwise correspondences between spherical signals. Moreover, the source may contain up to 90% outliers provided the underlying signal remains sufficiently complete relative to the target. The output is a rotation that aligns the source spherical distribution to the target, which can be applied to 3D points via Eq. 1. Pseudocode for $SPMC$, FRS , and $SPMC_FRS$ is provided in Appendix A.

H. Axis-Consistent vs. Axis-Ambiguous

Each $R \in SO(3)$ yields $TBV_x(R)$, $TBV_y(R)$, $TBV_z(R) \in S^2$. A TBV *sign flip* is the antipodal map $TBV_i \mapsto -TBV_i$ (axis preserved, direction reversed). In *axis-consistent* settings (e.g., homogeneous sensors sharing a base frame), TBVs correspond one-to-one— $x \leftrightarrow x$, $y \leftrightarrow y$, $z \leftrightarrow z$ —and differ only by a rotation (e.g., two Xsens IMUs on the same rigid body [55]). In *axis-ambiguous* settings (e.g., robot world/hand-eye or cross-vendor rigs), axes may be *relabelled* and/or *sign-flipped* (e.g., source x -TBV matches target $-y$ -TBV).

I. Signed Permutations

Let $P \in \mathbb{R}^{3 \times 3}$ be a permutation matrix (relabels $\{x, y, z\}$) and $S = \text{diag}(s_x, s_y, s_z)$ with $s_i \in \{\pm 1\}$ (independent sign flips). Define

$$L \triangleq PS \in O(3), \quad L \in \{-1, 0, 1\}^{3 \times 3}, \quad LL^T = I \quad (2)$$

which, by *left* multiplication, reorders/flips the *rows* of any 3×3 triad (TBVs under our row-vector convention). There are $3! \times 2^3 = 48$ hypotheses; if a right-handed triad is required, restrict to the 24 with $\det(L) = +1$, noting $\det(L) = \text{sgn}(P) s_x s_y s_z$.

J. From TBV Pairings to Configurations

With target TBVs $TBV_x(\mathcal{A})$, $TBV_y(\mathcal{A})$, $TBV_z(\mathcal{A})$ and source TBVs $TBV_x(\mathcal{B})$, $TBV_y(\mathcal{B})$, $TBV_z(\mathcal{B})$, any $L = PS$ induces a bijection π and signs (s_x, s_y, s_z) , yielding

$$\begin{aligned} & (TBV_x(\mathcal{A}), s_x TBV_{\pi(x)}(\mathcal{B})), \\ & (TBV_y(\mathcal{A}), s_y TBV_{\pi(y)}(\mathcal{B})), \\ & (TBV_z(\mathcal{A}), s_z TBV_{\pi(z)}(\mathcal{B})). \end{aligned} \quad (3)$$

For each pair we run our spherical pattern matching algorithms ($SPMC$, FRS , $SPMC_FRS$) to obtain per-axis rotations (R_x, R_y, R_z) and scores (c_x, c_y, c_z) , then select

$$L^* \in \arg \max_{L \in \mathcal{L}_{48}} (c_x(L) + c_y(L) + c_z(L)).$$

K. Right-Handed Constraint (Practical Rule)

Most frames are proper; enforce

$$\begin{aligned} \det(L) &= \det(P) s_x s_y s_z = +1 \\ \iff s_x s_y s_z &= \det(P), \end{aligned} \quad (4)$$

which halves the search to 24 configurations. *Examples:* (i) Identity π : $\det(P) = +1 \Rightarrow s_x s_y s_z = +1$ (even number of flips). (ii) Swapping x and y : $\det(P) = -1 \Rightarrow s_x s_y s_z = -1$ (odd number of flips).

L. Constructing L in Practice

We build L from a mapping string (e.g., “ $Ax \rightarrow -By$, $Ay \rightarrow +Bx$, $Az \rightarrow +Bz$ ”). Let rows $\{1, 2, 3\}$ correspond to $\{Ax, Ay, Az\}$ and columns $\{1, 2, 3\}$ to $\{Bx, By, Bz\}$. Then

$$L_{i,j} = \begin{cases} +1, & \text{if } A\text{-}i \text{ matches } +B\text{-}j, \\ -1, & \text{if } A\text{-}i \text{ matches } -B\text{-}j, \\ 0, & \text{otherwise,} \end{cases} \quad (5)$$

e.g., $L = \begin{bmatrix} 0 & -1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 1 \end{bmatrix}$ effects $[-By; +Bx; +Bz]$ when left-multiplying a 3×3 triad.

III. PROBLEM STATEMENT

Let $\mathcal{A} = \{\mathbf{a}_i\}_{i=1}^N$ and $\mathcal{B} = \{\mathbf{b}_j\}_{j=1}^M$ be two sets of orientations with $\mathbf{a}_i, \mathbf{b}_j \in SO(3)$. We posit that a subset of $\mathcal{B}$ is related to $\mathcal{A}$ by an unknown global rotation $\mathbf{R} \in SO(3)$ and a (possibly) unknown signed permutation $\mathbf{L} \in \mathcal{L} \subset O(3)$ that models axis relabels and sign flips (Section 2); the axis-consistent case corresponds to $\mathbf{L} = \mathbf{I}$. Formally, there exist an index set $\mathcal{J} \subseteq \{1, \dots, M\}$ and an injective association $\sigma: \mathcal{J} \rightarrow \{1, \dots, N\}$ such that

$$\mathbf{b}_j = \mathbf{L}^\top \mathbf{a}_{\sigma(j)} \mathbf{R} \epsilon_j, \quad j \in \mathcal{J}, \quad (6)$$

where $\epsilon_j \in SO(3)$ models measurement noise/drift with $\epsilon_j \approx \mathbf{I}$ for inliers; for $j \notin \mathcal{J}$ (outliers), $\mathbf{b}_j$ is unconstrained.

Our goal is to estimate $(\mathbf{L}, \mathbf{R})$ *without known correspondences* and under potential size mismatch $M \neq N$. A robust formulation is

$$(\mathbf{L}^*, \mathbf{R}^*) \in \arg \min_{\substack{\mathbf{L} \in \mathcal{L} \\ \mathbf{R} \in SO(3)}} \min_{\substack{\mathcal{J} \subseteq \{1, \dots, M\} \\ \sigma: \mathcal{J} \rightarrow \{1, \dots, N\}}} \sum_{j \in \mathcal{J}} \rho(d(\mathbf{b}_j, \mathbf{L}^\top \mathbf{a}_{\sigma(j)} \mathbf{R})),$$

where $d(\cdot, \cdot)$ is the geodesic distance on $SO(3)$ and ρ is a robust loss (e.g., Huber). In our correspondence-free approach, we do not solve the inner minimization over $(\mathcal{J}, \sigma)$ explicitly; instead, we maximize a permutation- and cardinality-invariant correlation objective induced by the TBVs on S^2 (Sections II-F-3), enumerating $\mathbf{L}$ (proper cases when appropriate) and fusing per-axis estimates into a single rotation.

Figure 3 provides an example of our method.

IV. METHOD: $SO(3)$ PATTERN ALIGNMENT

Goal. Given a *source* orientation set $\mathcal{B} \subset SO(3)$ and a *target* set $\mathcal{A} \subset SO(3)$, estimate a single rotation that aligns $\mathcal{B}$ to $\mathcal{A}$ so subsequent computations occur in the target frame. In the axis-ambiguous case, we also recover an unknown signed permutation L that accounts for axis relabels and sign flips (Section II).

A. From $SO(3)$ to three spherical alignments

a) *TBVs and S^2 distributions.*: From each rotation we extract row-centric Transformed Basis Vectors (TBVs; Section II-C). Each $R \in SO(3)$ yields $\text{TBV}_x(R), \text{TBV}_y(R), \text{TBV}_z(R) \in S^2$, producing three spherical point sets per orientation set: $S2_x, S2_y, S2_z$ for both $\mathcal{A}$ and $\mathcal{B}$ (Section II-F).

b) *Per-axis spherical matching.*: Given any TBV pair $(U, V) \subset S^2$, we apply a correspondence-free spherical matcher—SPMC, FRS, or the hybrid SPMC_FRS [44]—to obtain (i) a rotation $\hat{R} \in SO(3)$ that aligns V to U , and (ii) a scalar correlation score c . This step is robust to set-size mismatch and outliers and runs in $\mathcal{O}(n)$ time per call [44]. Fig. 1 visually illustrates all the steps of $SO3_SPMC$. Appendix B provides some numerical simulations of $SO3_SPMC$ showing its accuracy and runtime on example synthetic data.

B. Axis-consistent vs. PASI pipelines

a) *Axis-consistent (one-to-one axes).*: When axes match by label and sign ($x \leftrightarrow x, y \leftrightarrow y, z \leftrightarrow z$), we align each axis independently:

$$\begin{aligned} (R_x, c_x) &= \text{MATCH}(A_x, B_x), \\ (R_y, c_y) &= \text{MATCH}(A_y, B_y), \\ (R_z, c_z) &= \text{MATCH}(A_z, B_z). \end{aligned}$$

The three estimates (R_x, R_y, R_z) are fused into a single rotation $\bar{R} \in SO(3)$ via a projected arithmetic mean $\bar{R} = \Pi_{SO(3)}(\frac{1}{3}(R_x + R_y + R_z))$ or a Karcher mean. Optionally, a one-shot Procrustes refinement is applied on the reconstructed rotations (Algorithm 1).

b) *Permutation-and-Sign Invariant (PASI).*: When axes may be relabelled and/or sign-flipped (Section II-F), we enumerate signed permutations $L = PS$ (Eq. (2)). We restrict to the 24 proper cases with $\det(L) = +1$ (Eq. (4)). The PASI procedure is:

1) *Precompute per-axis matches (18 calls).* For all $i, j \in \{x, y, z\}$ and $s \in \{\pm 1\}$,

$$(R_{ij}^{(s)}, c_{ij}^{(s)}) \leftarrow \text{MATCH}(A_i, s B_j).$$

2) *Score each signed permutation.* Each $L = PS$ induces the pairings in Eq. (3) via the permutation π and signs (s_x, s_y, s_z) . Score

$$S(L) = c_{x, \pi(x)}^{(s_x)} + c_{y, \pi(y)}^{(s_y)} + c_{z, \pi(z)}^{(s_z)}.$$

3) *Select and fuse.* Choose $L^* = \arg \max_L S(L)$, collect the corresponding (R_x, R_y, R_z) , and fuse to a single $\bar{R} \in SO(3)$ (projected mean or Karcher mean), followed by optional Procrustes refinement.

Because the 18 per-axis matches are computed once and reused across all 24 hypotheses, PASI adds only a lightweight lookup/summation overhead. The three PASI variants—**PASI_SO3_SPMC**, **PASI_SO3_FRS**, **PASI_SO3_SPMC_FRS**—differ only in the choice of MATCH (SPMC, FRS, SPMC_FRS [44]).




Fig. 3. **RWHE on $SO(3)$ via TBVs and PASI.** (a) Transformed Basis Vectors (TBVs; S^2 points) of the end-effector ($\mathcal{A}$) and camera ($\mathcal{B}$) orientation trajectories from the ETH Hand-Eye Calibration dataset (robot_arm_real split) [10]. The sets exhibit an *axis-ambiguous* relation: the best signed-permutation is $L^* = \begin{bmatrix} 0 & -1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 1 \end{bmatrix}$, i.e., $Ax \rightarrow -By$, $Ay \rightarrow +Bx$, $Az \rightarrow +Bz$. (b) Our PASI variant enumerates proper signed permutations $L = PS$ ($\det(L) = +1$), selects L^* by maximizing the summed per-axis spherical correlation, then estimates a continuous rotation $\bar{R}$ with a spherical matcher (SPMC/FRS/hybrid), yielding the two-sided alignment $\mathcal{A} \approx L^* \bar{B} \bar{R}^\top$.

C. Recovered mapping and convention

Axis-ambiguous alignment (final form). Under the row-vector convention, the recovered mapping between sets is

$$\mathcal{A} \approx L^* \bar{B} \bar{R}^\top, \quad (7)$$

where L^* is the selected proper signed permutation acting on TBV rows, and $\bar{R} \in SO(3)$ is the continuous rotation aligning the sign-corrected source to the target. When axes are consistent, $L^* = I$ and Eq. (7) reduces to the axis-consistent case.

D. Algorithm families and baseline structure

a) *Six instantiations:* All six methods share the same core: (i) extract TBVs; (ii) treat each TBV axis as an S^2 point set; (iii) align via spherical registration; (iv) fuse per-axis rotations. *Axis-consistent:* **SO3_SPMC** (fastest; may be mean-shift sensitive under heavy outliers), **SO3_FRS** (slower; more robust), **SO3_SPMC_FRS** (SPMC proposal + few FRS refinements) [44]. *PASI:* **PASI_SO3_SPMC**, **PASI_SO3_FRS**, **PASI_SO3_SPMC_FRS** wrap the axis-consistent core with the signed-permutation enumeration above.

b) *Connection to pseudocode:* Algorithm 1 (*Axis-Consistent Baselines on $SO(3)$*) provides the common structure used by **SO3_SPMC**, **SO3_FRS**, and **SO3_SPMC_FRS**; only **MATCH** changes. The PASI counterparts follow the *same* inner steps but prepend the 24-way signed-permutation selection.

V. COMPLEXITY ANALYSIS

Prior work [44] proved that SPMC, FRS, and the hybrid SPMC_FRS each run in $\mathcal{O}(n)$ time, where n is the number of points involved in the spherical match (proportional to the number of input rotations per axis). By contrast, spherical/ $SO(3)$ cross-correlation methods based on harmonic transforms or rotation-space discretization scale as $\mathcal{O}(N_r^3 \log N_r)$, where N_r denotes either the resolution of the $SO(3)$ grid or the retained bandwidth (number of spherical

Algorithm 1: Axis-Consistent Baselines on $SO(3)$ (Common Structure for SO3_SPMC, SO3_FRS, SO3_SPMC_FRS)

Inputs. Target rotations $\{R_k^A\}_{k=1}^n$, Source rotations $\{R_\ell^B\}_{\ell=1}^m$. A per-axis spherical matcher **MATCH**($\cdot, \cdot$) instantiated as SPMC, FRS, or SPMC_FRS.

Outputs. Estimated global alignment $R_{\text{est}} \in SO(3)$ (maps Source to Target under the row-vector convention).

Procedure.

- 1) *Extract TBVs per set.*
 $(A_x, A_y, A_z) \leftarrow \text{TBVsFromSO3}(\{R_k^A\})$, $(B_x, B_y, B_z) \leftarrow \text{TBVsFromSO3}(\{R_\ell^B\})$. Each A_i, B_i is an S^2 point set (row-centric TBVs; Section II-C).
- 2) *Align each axis via the chosen matcher.*
 $(_, \tilde{B}_x) \leftarrow \text{MATCH}(A_x, B_x)$, $(_, \tilde{B}_y) \leftarrow \text{MATCH}(A_y, B_y)$,
 $(_, \tilde{B}_z) \leftarrow \text{MATCH}(A_z, B_z)$.
- 3) *Form mean frames.*
 $\text{MeanAligned} = [\mu(\tilde{B}_x); \mu(\tilde{B}_y); \mu(\tilde{B}_z)]$, $\text{MeanSource} = [\mu(B_x); \mu(B_y); \mu(B_z)]$, where $\mu(\cdot)$ is the Euclidean mean (normalize if desired).
- 4) *Estimate global rotation (Procrustes projection).*
 $R_{\text{frame}} = \Pi_{SO(3)}(\text{MeanAligned}^\top \text{MeanSource})$, $R_{\text{est}} = R_{\text{frame}}^\top$.

Instantiation.

- **SO3_SPMC:** **MATCH** $\equiv$ SPMC (fastest; can be mean-shift sensitive under heavy outliers).
- **SO3_FRS:** **MATCH** $\equiv$ FRS (slower; more robust).
- **SO3_SPMC_FRS:** **MATCH** $\equiv$ SPMC_FRS (SPMC proposal + few FRS refinements).

Notes. (i) Structure is identical across the three baselines; only **MATCH** changes. (ii) To handle axis ambiguity, wrap this baseline with the PASI enumeration (Section IV) to obtain **PASI_SO3_**.

harmonic coefficients) [33], [34]; this cubic-log scaling limits practical resolution and thus accuracy at scale [35].

a) *Axis-consistent baselines:* Let $n_A = |\mathcal{A}|$, $n_B = |\mathcal{B}|$, and $n \triangleq n_A + n_B$. Each per-axis spherical match (SPMC, FRS, or SPMC_FRS) costs $\Theta(n)$. The axis-consistent baselines (**SO3_SPMC**, **SO3_FRS**, **SO3_SPMC_FRS**) perform three




Fig. 4. (a) TBV S^2 distributions for the target set A_5 (EuRoC Machine Hall) and the seven source variants B_1 – B_7 after applying rotation R_4 . (b) Representative alignment produced by **SO3_SPMC_FRS**.

such matches (one for each of x, y, z):

$$T_{\text{axis-consistent}} = 3\Theta(n) + \mathcal{O}(n) = \mathcal{O}(n),$$

where the additional $\mathcal{O}(n)$ term accounts for simple reductions (means) and a constant-time 3×3 projection to $SO(3)$.

b) PASI variants: Naïvely, enumerating all 24 proper signed permutations and matching three axes per hypothesis would cost $24 \times 3\Theta(n)$. Instead, as detailed in Section IV, we *precompute* the 18 per-axis matches for all $(i, j, s) \in \{x, y, z\} \times \{x, y, z\} \times \{\pm 1\}$, and then evaluate the 24 hypotheses by table lookup and summation:

$$T_{\text{PASI}} = 18\Theta(n) + \mathcal{O}(24) + \mathcal{O}(n) = \mathcal{O}(n).$$

The final projection (or Karcher mean) and optional Procrustes refinement add only linear accumulation over rotations plus a constant-time 3×3 SVD.

c) Hybrid and refinement costs: For **SO3_SPMC_FRS**, let K be the (small) number of FRS refinements; the cost is $\Theta(n) + K\Theta(n) = \Theta(n)$ with a modest constant. The Procrustes step forms $\sum_{\ell} \hat{R}_{\ell}^{B, \text{aligned}} (R_{\ell}^B)^{\top}$ in $\mathcal{O}(n_B)$ and applies a 3×3 SVD (constant).

All proposed $SO(3)$ pattern matchers—**SO3_SPMC**, **SO3_FRS**, **SO3_SPMC_FRS** and their PASI counterparts—retain $\mathcal{O}(n)$ time with small constants (3 matches for axis-consistent; 18 precomputed matches for PASI), versus $\mathcal{O}(N_r^3 \log N_r)$ for cross-correlation on $\mathbb{S}^2/SO(3)$ [33], [34], [35].

VI. EXPERIMENTAL EVALUATION

We assess both *axis-consistent* and *axis-ambiguous* $SO(3)$ alignment settings (cf. Section IV). For the axis-consistent case we run controlled simulations derived from EuRoC MAV Machine Hall sequences MH_01–MH_05 [45] and test the three axis-consistent aligners. For the axis-ambiguous

case we solve the *rotational* component of Robot World-Hand-Eye calibration on the ETH Hand-Eye benchmark, split `robot_arm_real` [10], comparing our PASI variants against time-aligned RANSAC baselines (RANSAC Classic (RC) [56], RANSAC with pose filtering [57], and exhaustive-search (ES) with RANSAC verification over axis configurations).

A. Simulation: Axis-Consistent Scenario

From EuRoC MH_01–MH_05 we extract orientations to form five target sets $\mathcal{A} = \{A_1, \dots, A_5\}$ [45]. For each A_k , we create seven corrupted source sets $\{B_1, \dots, B_7\}$.

- B_1 : no noise, no outliers (same cardinality as A_k).
- B_2 : add zero-mean Gaussian noise, std. 0.01 rad; no outliers.
- B_3 : keep noise std. 0.01; replace 10% of samples with random rotations (outliers).
- B_4 – B_7 : same noise as B_2 ; increase outliers to 25%, 50%, 75%, and 90%, respectively.

To further stress-test, we sample 100 random global rotations $\{R_1, \dots, R_{100}\} \subset SO(3)$ and apply each to every B_j , yielding $7 \times 100 = 700$ sources per template A_k (i.e., 3500 trials across A_1 – A_5). We denote configurations as A_5 – B_5 – R_4 : (A_5 is the template; B_5 has 50% outliers and noise 0.01; rotated by R_4). Figure 4(a) shows TBVs of a target A_k and its seven sources under R_4 ; Fig. 4(b) presents a representative registration using **SO3_SPMC_FRS**. We report geodesic rotation error, success rates under angle thresholds, and runtime.

B. RWHE Rotational Subproblem: Axis-Ambiguous Scenario

We adopt the ETH Hand-Eye benchmark (`robot_arm_real`) [10]. The platform is a UR10 manipulator with an Intel RealSense SR300 RGB-D camera rigidly mounted near the end-effector; an AprilTag target




Fig. 5. Axis-Consistent Scenario: Quantitative results on simulated $SO(3)$ data derived from the EuRoC Machine Hall sequences. For each target set, the seven source configurations $B1 \rightarrow B7$ (increasing outlier ratio) are shown as box plots; the median of each box is annotated under the plot.

fixed in the environment provides camera poses. The dataset logs time-stamped 6-DoF trajectories for the hand in the robot base frame $T_{BH}(t)$ and the eye in the world/target frame $T_{WE}(t)$ as CSVs with fields $\{t, x, y, z, q_x, q_y, q_z, q_w\}$ (Hamilton convention). Here we use *orientation only*: we form two *unpaired* rotation sets, $\mathcal{A} = \{R_{BH}(t)\}$ (target: end-effector in base) and $\mathcal{B} = \{R_{WE}(t)\}$ (source: camera in world).

Because the frames may differ by axis relabels and sign flips, the axis-consistent methods are not applicable; instead we employ the *PASI* variants: **PASI_SO3_SPMC**, **PASI_SO3_FRS**, and **PASI_SO3_SPMC_FRS**. Each *PASI* method precomputes the 18 per-axis spherical matches, enumerates the 24 proper signed permutations $L = PS(\det(L) = +1)$, scores hypotheses by summed correlations, selects L^* , and fuses the per-axis rotations into a single $R^* \in SO(3)$ aligning $\mathcal{B}$ to $\mathcal{A}$. We compare against time-aligned RANSAC baselines: RC [56], pose-filtered RANSAC [57], and ES+RANSAC verification. The metrics include geodesic error, success rates, and runtime; translation is out-of-scope, and R^* can serve as a robust initializer for full RWHE pipelines.

VII. RESULTS AND DISCUSSION

A. Simulation: Axis-Consistent Scenario

Fig. 5 summarizes the error distributions across all five EuRoC-derived targets and seven corruption levels. Three observations are consistent across datasets:

(1) **SO3_SPMC is best overall.** **SO3_SPMC** achieves the lowest median geodesic error on the majority of configura-

tions, often with tight interquartile ranges. This aligns with prior findings that SPMC excels when the underlying spherical signals are “complete” (i.e., well-overlapped, above $\approx 65\%$ overlap) [44].

(2) **SO3_FRS is consistently suboptimal.** Across datasets, **SO3_FRS** exhibits higher medians and wider tails, reflecting its greater robustness cost when the axis-consistent assumption holds and coverage is good. While FRS can be advantageous under heavy contamination, here it underperforms SPMC in accuracy, with the overall rotation error close to 90° .

(3) **SO3_SPMC_FRS is competitive but has rare failures.** The hybrid is generally close to **SO3_SPMC** on medians (often $< 5^\circ$), but we observe occasional degradations; e.g., on A3-B4 the median reaches $\sim 12.87^\circ$. A plausible cause is that the FRS refinement, initialized by a good SPMC proposal, can be attracted to a nearby local optimum under certain multimodal or symmetric TBV distributions, leading to a suboptimal final rotation.

Averaging over datasets (Fig. 5b), both **SO3_SPMC** and **SO3_SPMC_FRS** perform well, with **SO3_SPMC** having the edge in accuracy and stability. **SO3_SPMC** has an overall rotation error between 0.15° for the MH-01-Easy set and 0.67° for the MH-05-Difficult set, while the overall rotation error for **SO3_SPMC_FRS** is between 0.40° - 1.00° for the various sets. Our takeaway from this results is when axis correspondences are known (axis-consistent regime), **SO3_SPMC** is the recommended default.

TABLE I
AXIS-AMBIGUOUS SCENARIO: BASELINE RESULTS ON THE
HAND-EYE/RWHE TASK. RMSE REPORTED IN DEGREES; RUNTIME IN
SECONDS.

Method	RMSE ($^{\circ}$)	Runtime (s)
RC_no_filter_best	0.6267	56.0
RC_filter_best	0.6677	40.7
RS_filter_best	0.6543	44.3
ES_best	0.6543	142.0
PASI_SO3_SPMC (ours)	0.6821	0.7
PASI_SO3_FRS (ours)	3.0008	5.6
PASI_SO3_SPMC_FRS (ours)	0.6122	6.2

B. Robot World-Hand-Eye on $SO(3)$: Axis-Ambiguous Scenario

Fig. 3(a) visualizes the TBVs (as S^2 point sets) for the end effector ($\mathcal{A}$) and the camera ($\mathcal{B}$). From the patterns it is evident that while z -TBVs align across sets, the x - y axes are permuted and sign-flipped. The selected signed permutation is

$$L^* = \begin{bmatrix} 0 & -1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 1 \end{bmatrix},$$

indicating $A_x \leftrightarrow -B_y$, $A_y \leftrightarrow +B_x$, and $A_z \leftrightarrow +B_z$. Hence, axis-consistent baselines are inappropriate here, whereas the PASI variants explicitly handle such relabel/sign ambiguities.

Fig. 3(b) shows qualitative overlays for **PASI_SO3_SPMC**, **PASI_SO3_FRS**, and **PASI_SO3_SPMC_FRS**. Table I gives the RMSE and Runtime for each method, while Fig. 6 reports angular-error histograms. For *evaluation only*, we time-align trajectories to form pairs and compute geodesic errors; the *algorithms themselves* consume the raw, unpaired orientation sets (different cardinalities, no synchronization).

Table I and Fig. 6 lead to three conclusions: (i) **PASI_SO3_SPMC** and **PASI_SO3_SPMC_FRS** are clearly superior to alternatives in accuracy; (ii) our PASI pipeline avoids any correspondence search/synchronization and is $\approx 6 - 60\times$ faster than the best time-aligned RANSAC-with-filtering baseline, depending on the variant; (iii) among our methods, **PASI_SO3_SPMC** is the fastest (~ 0.7 s), while **PASI_SO3_SPMC_FRS** is the most accurate (RMSE 0.6122° , median 0.51°). **PASI_SO3_FRS** is much less accurate (median and RMSE near 3°).

Overall, PASI delivers correspondence-free rotational RWHE alignment on $SO(3)$ that is both accurate and efficient in the presence of axis permutations and sign flips; the SPMC-based variant is preferred for speed, and the hybrid for the lowest error.

VIII. CONCLUSION

We presented a family of correspondence-free alignment algorithms on $SO(3)$ that (i) decompose rotations into *Transformed Basis Vectors* (TBVs), (ii) align the resulting S^2 point sets with fast spherical matchers (SPMC, FRS, and a hybrid) [44], and (iii) fuse per-axis estimates into a single rotation. To handle unknown axis conventions, we introduced *Permutation-and-Sign Invariant* (PASI) variants that enumerate proper




Fig. 6. Axis-Ambiguous Scenario: Histogram of rotational calibration error on the ETH Hand-Eye benchmark (*robot_arm_real*) [10]. SO3_SPMC (blue), SO3_FRS (orange), SO3_SPMC_FRS (green). All algorithms operate on the raw, unpaired rotation streams; for evaluation only, we time-align the trajectories and compute per-pair geodesic angular error on $SO(3)$ (in degrees), enabling a fair comparison to baselines.

signed permutations and select the maximizer, making the pipeline invariant to axis relabels and sign flips while preserving efficiency.

The methods accept unpaired sets of differing cardinalities, run in linear time in the number of samples, and remain accurate under high outlier rates. On EuRoC-derived simulations (axis-consistent), **SO3_SPMC** is the most accurate and stable default; the hybrid is competitive but can exhibit occasional local-optimum failures on multimodal TBV distributions. On the ETH Hand-Eye benchmark (axis-ambiguous), PASI variants achieve strong accuracy without any time alignment or correspondence search, with only a small constant-factor overhead from enumerating axis configurations.

Beyond benchmarking, our rotation-only initializer is directly useful whenever paired samples are unavailable, including IMU wearables/gloves for sensor-to-segment (S2S) calibration and cross-session re-calibration, IMU-magnetometer heading calibration (with ellipsoid-fit compensation), and VIO/SLAM initialization among others.

Limitations and future work: Our pipeline is not yet globally certifiable outside special cases; integrating certifiable $SO(3)$ solvers is a natural direction. Other avenues include streaming/online variants with uncertainty propagation, principled PASI search heuristics, learned activity-conditioned signatures for CT selection, and extensions that couple our rotational initializer with translation recovery for full $SE(3)$ calibration. We also expect that the overall method could be extended to data with higher dimensions, i.e. $SO(N)$.

Overall, TBV-based alignment with PASI offers a simple, scalable, and robust alternative to spherical/ $SO(3)$ correlation, making correspondence-free $SO(3)$ set alignment practical for real-world calibration, registration, and initialization workflows.

APPENDIX A SPHERICAL POINT PATTERN MATCHING ALGORITHMS

Following are summaries of the SPMC, FRS, and Hybrid SPMC + FRS algorithms.

Algorithm A1: Spherical Pattern Matching by Correlation (SPMC)

Input. Two spherical point sets $\mathcal{A} = \{\mathbf{a}_i\}_{i=1}^N$, $\mathcal{B} = \{\mathbf{b}_j\}_{j=1}^M \subset S^2$.
Output. $R_{\text{opt}} \in SO(3)$ aligning $\mathcal{B}$ to $\mathcal{A}$; correlation peak $C(s^*)$.
Steps.

- 1) *Mean-to-North alignment.* Let $\bar{\mathbf{a}} = \frac{1}{N} \sum_i \mathbf{a}_i$, $\bar{\mathbf{b}} = \frac{1}{M} \sum_j \mathbf{b}_j$. Compute $R_A, R_B \in SO(3)$ such that $R_A \bar{\mathbf{a}} = R_B \bar{\mathbf{b}} = [0, 0, 1]^\top$.
- 2) *Rotate sets to the North pole.* $\mathcal{A}^{\text{NP}} = \{R_A \mathbf{a}_i\}$, $\mathcal{B}^{\text{NP}} = \{R_B \mathbf{b}_j\}$. (Optional: flip hemisphere so all points have $z \geq 0$.)
- 3) *Build azimuth histograms.* Form 2D spherical occupancy histograms, then sum over polar bins to get 1D azimuth histograms $h_A, h_B \in \mathbb{R}^K$ (e.g., $K=360$).
- 4) *1D circular correlation.* For shifts $s \in \{0, \dots, K-1\}$,

$$C(s) = \sum_{\lambda=0}^{K-1} h_A(\lambda) h_B((\lambda + s) \bmod K).$$

Let $s^* = \arg \max_s C(s)$.

- 5) *Recover the rotation.* $R_{\text{shift}} = R_z(2\pi s^*/K)$, $R_{\text{opt}} = R_A^{-1} R_{\text{shift}} R_B$.

Notes. Exact recovery holds in the noise/outlier-free one-to-one case (means coincide; z -axis residual only).

Algorithm A2: Fast Rotation Search (FRS)

Input. $\mathcal{A} = \{\mathbf{a}_i\}_{i=1}^N$, $\mathcal{B} = \{\mathbf{b}_j\}_{j=1}^M \subset S^2$; max iters T ; tolerance ε .
Output. $R_{\text{est}} \in SO(3)$ aligning $\mathcal{B}$ to $\mathcal{A}$.
Steps.

- 1) *Fix target histograms.* Compute mean-to-North R_A ; set $\mathcal{A}^{\text{NP}} = \{R_A \mathbf{a}_i\}$. Build fixed azimuth histograms h_A^x, h_A^y, h_A^z (three views / axis-angle parameterizations).
- 2) *Initialize.* $R \leftarrow I_3$, $t \leftarrow 0$.
- 3) *Iterate (coarse-to-fine azimuth updates).* While $t < T$:
 - a) Form $\mathcal{B}^{(t)} = \{R \mathbf{b}_j\}$ and $\mathcal{B}^{\text{NP}} = \{R_A \mathcal{B}^{(t)}\}$.
 - b) Build moving azimuth histograms h_B^x, h_B^y, h_B^z .
 - c) For each axis $u \in \{x, y, z\}$, compute circular shift $s_u^* = \arg \max_s \sum_{\lambda} h_A^u(\lambda) h_B^u((\lambda + s) \bmod K)$.
 - d) Let $R_{\text{inc}} = R_z(2\pi s_z^*/K) R_y(2\pi s_y^*/K) R_x(2\pi s_x^*/K)$.
 - e) **Stop** if all $|s_u^*| \leq \varepsilon_K$ (equivalently, $\|\log(R_{\text{inc}})\| \leq \varepsilon$).
 - f) Update $R \leftarrow R_{\text{inc}} R$, $t \leftarrow t + 1$.
- 4) *Return.* $R_{\text{est}} = R$.

Notes. Using SPMC's mean-to-North initialization and multi-axis azimuth updates yields rapid convergence; small- T (e.g., 5-10) is often sufficient in practice.

Algorithm A3: Hybrid SPMC + FRS

Input. $\mathcal{A}, \mathcal{B} \subset S^2$; FRS hyperparameters (T, ε) .

Output. $R_{\text{est}} \in SO(3)$.

Steps.

- 1) *SPMC initialization.* Run SPMC (Alg. A1) to obtain R_0 .
- 2) *Warm-start FRS.* Initialize $R \leftarrow R_0$ and run FRS (Alg. S2) with small T and tight ε .
- 3) *Return.* $R_{\text{est}} = R$.

Notes. Combines SPMC's fast global azimuth recovery with FRS's local refinement; empirically more robust under moderate outliers than pure SPMC, while remaining faster than pure FRS.

APPENDIX B NUMERICAL SIMULATIONS OF SO_3 ALIGNMENT

While each of these algorithms was previously presented and validated in [44], we provide additional examples and validation of the SO_3_SPMC algorithm on several commonly occurring spherical patterns.

A. Numerical Simulations

A truly uniform sampling of $SO(3)$ induces (approximately) uniform TBV distributions on S^2 , yielding little discriminative structure. To obtain informative test cases, we generate structured $SO(3)$ samples by restricting Euler angles to bounded ranges.

a) *Data generation (set $\mathcal{A}$).*: We sample $n = 2000$ Euler-angle triples (roll, pitch, yaw) within the ranges

$$\text{roll} \in (-40^\circ, 10^\circ), \text{pitch} \in (-20^\circ, 20^\circ), \text{yaw} \in (-10^\circ, 50^\circ).$$

Each triple is converted to a rotation matrix using standard formulas [58], producing $\mathcal{A} = \{A_i\}_{i=1}^n \subset SO(3)$. Figure 7 illustrates a representative configuration and its S^2 (TBV) views.

b) *Ground-truth transform and source set.*: We draw an arbitrary ground-truth $R_{\text{gt}} \in SO(3)$ and synthesize

$$\mathcal{B} = \{B_i \equiv R_{\text{gt}} A_i\}_{i=1}^n.$$

We then run SO_3_SPMC to align the spherical patterns and recover R_{est} , the estimate of R_{gt} .

c) *Error metric.*: Given $R_1, R_2 \in SO(3)$, the geodesic (angular) error is

$$\theta(R_1, R_2) = \cos^{-1} \left(\frac{\text{trace}(R_1 R_2^\top) - 1}{2} \right) [\text{rad}],$$

converted to degrees for reporting. For the synthesized pair $(\mathcal{A}, \mathcal{B})$ we form the constant relative error

$$\Delta \equiv R_{\text{gt}} R_{\text{est}}^\top,$$

which equals $B_i (R_{\text{est}} A_i)^\top$ for each i (since $A_i A_i^\top = I$). We report the mean angular error (MAE) $\theta(\Delta)$; for completeness we also compute the per-pair angle and average (which coincides with $\theta(\Delta)$ in this noiseless construction).

d) *Scenarios and sizes.*: We evaluate three pattern families, each at two set sizes ($n = 2000$ and $n = 10$) to probe both data-rich and data-scarce regimes:

- **Scenario 1 (bounded integer ranges).** Random Euler triples with integer end-points: $(-40^\circ, 10^\circ)$, $(-20^\circ, 20^\circ)$, $(-10^\circ, 50^\circ)$.
- **Scenario 2 (bounded non-integer ranges).** Random Euler triples with non-integer end-points: $(-40.5^\circ, 10.5^\circ)$, $(-20.5^\circ, 20.5^\circ)$, $(-10.5^\circ, 50.5^\circ)$. This stresses potential interactions with histogram binning in SO_3_SPMC .
- **Scenario 3 (Gaussian around a mean).** Quaternions sampled from a Gaussian in angle space, centered at Euler means $[30^\circ, 10^\circ, 15.5^\circ]$ (roll, pitch, yaw), then mapped to $SO(3)$.




Fig. 7. Visualization of different stages of the SOM_SPMC algorithm with Scenario 1. The top, middle, and bottom rows correspond to different stages of the x, y, and z distributions respectively. The left-most column presents the initial S2-distributions, while the right-most column demonstrates the final alignment. From the plots in the last column, it is evident that all distributions overlap significantly, indicating a successful alignment.

TABLE II
MEAN ANGULAR ERROR (MAE, DEGREES) FOR SO3_SPMC UNDER VARIOUS SIMULATION SCENARIOS.

Scenario	2000 points	10 points
1 (bounded integer ranges)	0.060°	0.10°
2 (bounded non-integer ranges)	0.068°	0.15°
3 (Gaussian around a mean)	0.068°	0.08°

e) Protocol and reporting.: For each scenario and size, we run SO3_SPMC to estimate R_{est} , then report MAE in degrees. A summary across all settings appears in Table II. Representative S^2 visualizations of the TBV distributions and an end-to-end example are provided in Figure 7.

B. Ablation on Number of Points

We present how runtime and accuracy scale with the number of rotations. For each scenario, we form sets with cardinalities

$$n \in \{10^1, 10^2, 10^3, 10^4, 10^5, 10^6\}$$

and run SO3_SPMC to estimate R_{est} .

Figure 8 plots wall-clock time vs. n . After a small warm-up region ($n \leq 10^3$), the log-log slope is approximately 1 for 10^4 - 10^6 points, empirically confirming near-linear complexity consistent with the theory.

Figure 9 shows the rotational error vs. n . As n increases, alignment improves: larger samples stabilize the TBV histograms and reduce sensitivity to outliers (the error distribution's tail shrinks). At very small n (e.g., 10 points), the empirical means and histograms are more easily skewed, producing higher variance.

Overall, our numerical simulation results had accuracies of 0.068° or less for 2000 points, while the accuracies for 10 points were much larger, between 0.08°-0.15°. This may be due to the bin sized used, with 1° bins. For small numbers




Fig. 8. Runtime vs. number of rotations for SO3_SPMC. The log-log slope is ≈ 1 for 10^4 - 10^6 points.




Fig. 9. Rotational error vs. number of rotations for SO3_SPMC.

of data points, each point will be placed into a bin at a random location, and thus there will be some error between the original and rotated points. Using additional bins would provide increased accuracy, and this would be computationally tractable for small numbers of points. Alternatively, sub-pixel estimation schemes such as [59] could be used to improve the results.

REFERENCES

- [1] G. Wahba, "A least squares estimate of satellite attitude," *SIAM review*, vol. 7, no. 3, pp. 409–409, 1965.
- [2] P. J. Besl and N. D. McKay, "Method for registration of 3-d shapes," in *Sensor fusion IV: control paradigms and data structures*, vol. 1611. Spie, 1992, pp. 586–606.
- [3] D. Chetverikov, D. Svirko, D. Stepanov, and P. Krsek, "The trimmed iterative closest point algorithm," in *2002 International Conference on Pattern Recognition*, vol. 3. IEEE, 2002, pp. 545–548.
- [4] H. Li and R. Hartley, "The 3d-3d registration problem revisited," in *2007 IEEE 11th international conference on computer vision*. IEEE, 2007, pp. 1–8.
- [5] Y.-S. Shiu and S. Ahmad, "Calibration of wrist-mounted robotic sensors by solving homogeneous transform equations of the form $AX=XB$," *IEEE Transactions on Robotics and Automation*, vol. 5, no. 1, pp. 16–29, 1989.
- [6] A. Li, L. Wang, and D. Wang, "Simultaneous robot-world and hand-eye calibration using dual-quaternions and Kronecker product," *The International Journal of Robotics Research*, vol. 29, no. 5, pp. 635–653, 2010.
- [7] M. Shah, S. V. Naik, J. G. L., and R. M. Voyles, "Hand-eye calibration using Kronecker products," *Mechanism and Machine Theory*, vol. 69, pp. 54–66, 2013.
- [8] A. Tabb and K. M. A. Yousef, "Solving the robot-world/hand-eye(s) calibration problem with iterative methods," *Machine Vision and Applications*, vol. 28, no. 5–6, pp. 569–590, 2017.
- [9] F. Furrer, M. Fehr, T. Novkovic, H. Sommer, I. Gilitschenski, and R. Siegwart, "Evaluation of combined time-offset estimation and hand-eye calibration on robotic datasets," in *Field and Service Robotics*, ser. Springer Proceedings in Advanced Robotics. Springer, 2018, vol. 5, pp. 145–159.
- [10] —, "Evaluation of combined time-offset estimation and hand-eye calibration on robotic datasets," in *Field and Service Robotics: Results of the 11th International Conference*. Springer, 2017, pp. 145–159.
- [11] T. Seel, T. Schauer, and J. Raisch, "Joint axis and position estimation from inertial measurement data by exploiting kinematic constraints," in *Proc. IEEE Int. Conf. on Control Applications (CCA)*, 2012, pp. 45–49.
- [12] T. Seel, J. Raisch, and T. Schauer, "IMU-based joint angle measurement for gait analysis," *Sensors*, vol. 14, no. 4, pp. 6891–6909, 2014.
- [13] M. Ekdahl, K. D. Dahl, J. R. C. Hewett, *et al.*, "Inertial measurement unit sensor-to-segment calibration comparison for sport-specific motion analysis," *Sensors*, vol. 23, no. 18, p. 7987, 2023.
- [14] F. Olsson, M. Kok, T. Seel, and K. Halvorsen, "Robust plug-and-play joint axis estimation using inertial sensors," *Sensors*, vol. 20, no. 12, p. 3534, 2020.
- [15] G. D. Raimondo, F. Magistri, A. Pedrocchi, *et al.*, "A scoping review of sensor-to-segment calibration for wearable inertial sensors," *Sensors*, vol. 22, no. 24, p. 9837, 2022.
- [16] Y. Liu, Y. Zhang, M. Tan, K. Xu, Y. Li, and X. Wang, "Novel calibration algorithm for a three-axis strapdown magnetometer," *Sensors*, vol. 14, no. 5, pp. 8485–8504, 2014.
- [17] C. Pang, L. Chen, Y. Liu, and S. Wang, "Calibration and compensation of three-axis magnetometer errors by using a nonlinear ellipsoid fitting algorithm," *Journal of Magnetism and Magnetic Materials*, vol. 346, pp. 177–182, 2013.
- [18] H. Wu, S. Liu, Z. Wei, Z. Zhang, and Y. Sun, "An improved magnetometer calibration and compensation method based on Levenberg–Marquardt algorithm," *Measurement and Control*, vol. 53, no. 9–10, pp. 1312–1321, 2020.
- [19] X. Cui, A. Zhang, C. Qi, and S. Zhang, "An accurate calibration method for magnetometers based on ellipsoid fitting," in *Proc. IEEE/ION Position, Location and Navigation Symposium (PLANS)*, 2018, pp. 1010–1015.
- [20] A. I. Mourikis and S. I. Roumeliotis, "A multi-state constraint kalman filter for vision-aided inertial navigation," in *Proc. IEEE Int. Conf. on Robotics and Automation (ICRA)*, 2007, pp. 3565–3572.
- [21] T. Qin, P. Li, and S. Shen, "Vins-mono: A robust and versatile monocular visual-inertial state estimator," *IEEE Transactions on Robotics*, vol. 34, no. 4, pp. 1004–1020, 2018.
- [22] S. Leutenegger, S. Lynen, M. Bosse, R. Siegwart, and P. Furgale, "Keyframe-based visual-inertial odometry using nonlinear optimization," *The International Journal of Robotics Research*, vol. 34, no. 3, pp. 314–334, 2015.
- [23] D. Scaramuzza and Z. Zhang, "Visual-inertial odometry," *Encyclopedia of Robotics*, pp. 1–8, 2019.
- [24] T. Carpentier and A. Einbond, "Spherical correlation as a similarity measure for 3d radiation patterns of musical instruments," in *16ème Congrès Français d'Acoustique*. HAL Open Science, April 2022. [Online]. Available: <https://openaccess.city.ac.uk/id/eprint/28202/>
- [25] J. Keat, "Analysis of least-squares attitude determination routine doaoap," Technical Report CSC/TM-77/6034, Comp. Sc. Corp. Tech. Rep., 1977.
- [26] F. L. Markley, "Attitude determination using vector observations and the singular value decomposition," *Journal of the Astronautical Sciences*, vol. 36, no. 3, pp. 245–258, 1988.
- [27] F. L. Markley and D. Mortari, "Quaternion attitude estimation using vector observations," *The Journal of the Astronautical Sciences*, vol. 48, pp. 359–380, 2000.
- [28] J. Wu, Z. Zhou, B. Gao, R. Li, Y. Cheng, and H. Fourati, "Fast linear quaternion attitude estimator using vector observations," *IEEE Transactions on Automation Science and Engineering*, vol. 15, no. 1, pp. 307–319, 2017.
- [29] H. Yang and L. Carlone, "A quaternion-based certifiably optimal solution to the wahba problem with outliers," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2019, pp. 1665–1674.
- [30] T. S. Cohen, M. Geiger, J. Köhler, and M. Welling, "Spherical cnns," *arXiv preprint arXiv:1801.10130*, 2018.
- [31] M. Kazhdan, T. Funkhouser, and S. Rusinkiewicz, "Rotation invariant spherical harmonic representation of 3 d shape descriptors," in *Symposium on geometry processing*, vol. 6, 2003, pp. 156–164.
- [32] M. Kazhdan and T. Funkhouser, "Harmonic 3d shape matching," in *ACM SIGGRAPH 2002 conference abstracts and applications*, 2002, pp. 191–191.
- [33] L. Sorgi and K. Daniilidis, "Normalized cross-correlation for spherical images," in *Computer Vision-ECCV 2004: 8th European Conference on Computer Vision, Prague, Czech Republic, May 11-14, 2004. Proceedings, Part II 8*. Springer, 2004, pp. 542–553.
- [34] A. Makadia and K. Daniilidis, "Rotation recovery from spherical images without correspondences," *IEEE transactions on pattern analysis and machine intelligence*, vol. 28, no. 7, pp. 1170–1175, 2006.
- [35] C. Esteves, J.-J. Slotine, and A. Makadia, "Scaling spherical cnns," *arXiv preprint arXiv:2306.05420*, 2023.
- [36] E. Wise, P. Kaveti, Q. Chen, W. Wang, H. Singh, J. Kelly, D. M. Rosen, and M. Giamou, "A certifiably correct algorithm for generalized robot-world and hand-eye calibration," *arXiv preprint arXiv:2507.23045*, 2025.
- [37] A. Li, L. Wang, and D. Wu, "Simultaneous robot-world and hand-eye calibration using dual-quaternions and Kronecker product," *Int. J. Phys. Sci.*, vol. 5, no. 10, pp. 1530–1536, 2010.
- [38] Y. Wang, W. Jiang, K. Huang, S. Schwertfeger, and L. Kneip, "Accurate calibration of multi-perspective cameras from a generalization of the hand-eye constraint," in *2022 International Conference on Robotics and Automation (ICRA)*. IEEE, 2022, pp. 1244–1250.
- [39] M. Shah, "Solving the robot-world/hand-eye calibration problem using the Kronecker product," *Journal of Mechanisms and Robotics*, vol. 5, no. 3, p. 031007, 2013.
- [40] A. Tabb and K. M. Ahmad Yousef, "Solving the robot-world hand-eye (s) calibration problem with iterative methods," *Machine Vision and Applications*, vol. 28, no. 5, pp. 569–590, 2017.
- [41] M. Horn, T. Wodtke, M. Buchholz, and K. Dietmayer, "Extrinsic infrastructure calibration using the hand-eye robot-world formulation," in *2023 IEEE Intelligent Vehicles Symposium (IV)*. IEEE, 2023, pp. 1–8.
- [42] D. Evangelista, E. Olivastri, D. Allegro, E. Menegatti, and A. Pretto, "A graph-based optimization framework for hand-eye calibration for multi-camera setups," *arXiv preprint arXiv:2303.04747*, 2023.
- [43] M. Ulrich and M. Hillemann, "Uncertainty-aware hand-eye calibration," *IEEE Transactions on Robotics*, vol. 40, pp. 573–591, 2023.
- [44] A. Sarker and A. T. Asbeck, "Correspondence-free fast and robust spherical point pattern registration," in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2025, pp. 28 156–28 166.
- [45] M. Burri, J. Nikolic, P. Gohl, T. Schneider, J. Rehder, S. Omari, M. W. Achtelik, and R. Siegwart, "The EuRoC micro aerial vehicle datasets," *The International Journal of Robotics Research*, vol. 35, no. 10, pp. 1157–1163, 2016.
- [46] D. W. Henderson and E. Moura, "Experiencing geometry: on plane and sphere," (*No Title*), 1996.
- [47] H. Flanders, *Differential forms with applications to the physical sciences*. Courier Corporation, 1963, vol. 11.
- [48] S. Liao, E. Gavves, and C. G. M. Snoek, "Spherical regression: Learning viewpoints, surface normals and 3D rotations on N-spheres," in *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019, pp. 9751–9759.

- [49] K. S. Arun, T. S. Huang, and S. D. Blostein, "Least-squares fitting of two 3-d point sets," *IEEE Transactions on pattern analysis and machine intelligence*, vol. PAMI-9, no. 5, pp. 698–700, 1987.
- [50] J. R. Forbes and A. H. de Ruiter, "Linear-matrix-inequality-based solution to wahba's problem," *Journal of guidance, control, and dynamics*, vol. 38, no. 1, pp. 147–151, 2015.
- [51] B. K. Horn, "Closed-form solution of absolute orientation using unit quaternions," *Josa a*, vol. 4, no. 4, pp. 629–642, 1987.
- [52] B. K. Horn, H. M. Hilden, and S. Negahdaripour, "Closed-form solution of absolute orientation using orthonormal matrices," *Josa a*, vol. 5, no. 7, pp. 1127–1135, 1988.
- [53] K. Khoshelham, "Closed-form solutions for estimating a rigid motion from plane correspondences extracted from point clouds," *ISPRS Journal of Photogrammetry and Remote Sensing*, vol. 114, pp. 78–91, 2016.
- [54] J. Saunderson, P. A. Parrilo, and A. S. Willsky, "Semidefinite descriptions of the convex hull of rotation matrices," *SIAM Journal on Optimization*, vol. 25, no. 3, pp. 1314–1343, 2015.
- [55] A. Sarker, D.-R. Emenonye, A. Kelliher, T. Rikakis, R. M. Buehrer, and A. T. Asbeck, "Capturing upper body kinematics and localization with low-cost sensors for rehabilitation applications," *Sensors*, vol. 22, no. 6, p. 2300, 2022.
- [56] M. A. Fischler and R. C. Bolles, "Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography," *Commun. ACM*, vol. 24, no. 6, p. 381–395, June 1981. [Online]. Available: <https://doi.org/10.1145/358669.358692>
- [57] J. Schmidt, F. Vogt, and H. Niemann, "Robust hand-eye calibration of an endoscopic surgery robot using dual quaternions," in *Joint Pattern Recognition Symposium*. Springer, 2003, pp. 548–556.
- [58] G. G. Slabaugh, "Computing euler angles from a rotation matrix," CityUniv. London, London, U.K., Tech. Rep. 2000, 1999.
- [59] M. Guizar-Sicairos, S. T. Thurman, and J. R. Fienup, "Efficient subpixel image registration algorithms," *Optics letters*, vol. 33, no. 2, pp. 156–158, 2008.