

Supplementary Material for

4D Local Modeling Toward Dynamic Global Perception for Ambiguity-free Rotation-Invariant Point Cloud Analysis

Jiaxun Guo¹ Wentao Fan^{2*} Manar Amayri¹ Nizar Bouguila¹
¹CIISE, Concordia University

²Guangdong Provincial/Zhuhai Key Laboratory IRADS and Department of Computer Science,
 Beijing Normal-Hong Kong Baptist University

jiaxun.guo@mail.concordia.ca, wentaofan@bnbu.edu.cn,
 {manar.amayri, nizar.bouguila}@concordia.ca

This supplementary material provides additional theoretical and extended experimental results to support our proposed method of the main paper. We begin with theoretical insights into the rotation invariance of our proposed global-aware 4D pose feature (Section A), followed by a discussion on spatial ambiguities in point-pair representations (Section B) and challenges in using static global rotations (Section C). In Section D, we present a detailed formulation of the Bingham distribution on the hypersphere of unit quaternions, including entropy computation, parameter construction, acceptance-rejection sampling method, and the extended discussion on the Bingham design. Section E outlines the epoch-wise update strategy employed to stabilize the learning of global rotation. Finally, in Section F, we conduct additional experiments to assess the robustness of our method. These include evaluations under varying neighbourhood sizes, different values of the hyperparameter δ , and analysis of the learned global rotation R_g .

A Theoretical Analysis on Rotation Invariance of Global-aware 4D Pose Feature

We provide the theoretical analysis to establish the rotation invariance of the proposed Ga4DPF in RI convolution. Formal proofs demonstrate that 4DPPF and Ga4DPF remain invariant under arbitrary rotation, ensuring their robustness and consistency under arbitrary pose variations.

Theorem 1. *Let R_g denote a consistent global rotation, and the global anchor of the reference point \mathbf{p}_r in 4D space be defined as $\mathbf{p}_r^\dagger = ST(\mathbf{p}_r R_g; \mathbf{S})$. Then Ga4DPF(\mathbf{p}_r) is rotation invariant, i.e.*

$$\begin{aligned} \text{Ga4DPF}(\mathbf{p}_r) &= \text{Ga4DPF}(\mathbf{p}_r R) \\ &= (4\text{DPPF}(\mathbf{p}_r' R', \mathbf{p}_j' R'), 4\text{DPPF}(\mathbf{p}_r^\dagger R', \mathbf{p}_j^\dagger R')), \end{aligned} \quad (1)$$

*Corresponding author.

where $\forall R \in SO(3)$ is arbitrary rotation, R' is the representation of $R \in SO(4)$ in 4D space, \mathbf{p}_r' and \mathbf{p}_j' are the lifted reference point and the corresponding neighbourhood.

Proof. The 4DPPF is defined as

$$\begin{aligned} 4\text{DPPF}(\mathbf{p}_r', \mathbf{p}_j') &= (\|d\|_2, \cos(\alpha_1), \cos(\alpha_2), \cos(\alpha_3)), \\ \alpha_1 &= \angle(\Delta_r^1, d), \quad \alpha_2 = \angle(\Delta_j^1, d), \quad \alpha_3 = \angle(\Delta_r^1, \Delta_j^1), \end{aligned} \quad (2)$$

where $d = \mathbf{p}_j' - \mathbf{p}_r'$, $\angle(\cdot, \cdot)$ denotes the angle between two 4D vectors, and $\Delta = ST(\partial; \mathbf{S})$ represents the orthonormal basis vectors in 4D space. For ablations requiring signed angles in an oriented 3-subspace of \mathbb{R}^4 , we fix a canonical vector $\mathbf{e}_4 = (0, 0, 0, 1)^\top$, and define

$$s(\Delta_r^1, \Delta_j^1 | d) = \text{sign}(\det([d, \Delta_r^1, \Delta_j^1, \mathbf{e}_4])), \quad (3)$$

which determines the orientation of the 3-plane spanned by $(d, \Delta_r^1, \Delta_j^1)$. The corresponding signed angle $\angle(\Delta_r^1, \Delta_j^1)$ is defined via:

$$\angle(\Delta_r^1, \Delta_j^1) = \frac{\sqrt{1 - \cos^2(\Delta_r^1, \Delta_j^1)} \cdot s(\Delta_r^1, \Delta_j^1 | d)}{\cos(\Delta_r^1, \Delta_j^1)}. \quad (4)$$

Equations 3 and 4 define signed angles in 4D, enabling 4DPPF to encode orientation and resolve symmetries.

Because R' is an orthogonal matrix with $\det(R') = 1$, it preserves: (1) $\|R'd\| = \|d\|$. (2) $\langle R'x, R'y \rangle = \langle x, y \rangle$. (3) $\det(R') = 1$ ensures the sign in Equation 3 is preserved. Thus, each component of Equation 2 is invariant. Consequently,

$$4\text{DPPF}(\mathbf{p}_r', \mathbf{p}_j') = 4\text{DPPF}(\mathbf{p}_r' R', \mathbf{p}_j' R'). \quad (5)$$

Since each individual 4DPPF value is unchanged under R' , the aggregated descriptor is also unchanged:

$$\text{Ga4DPF}(\mathbf{p}_r) = \text{Ga4DPF}(\mathbf{p}_r R) \quad (6)$$

□

B Spatial Arrangement (Point-Pair) Ambiguity in PPF

We analyzed the ambiguity in the spatial arrangement of the classic point pair feature when a neighbouring point $\mathbf{p}_j(\theta)$ rotates around the primary axis of the LRF centred at \mathbf{p}_r as Figure 1.

Theorem 2. *Let \mathbf{p}_r be the reference point with local axis ∂_r^\perp . Consider a set of neighbors $\mathbf{p}_j(\theta)$ obtained by rotating a fixed point $\mathbf{p}_j(0)$ around ∂_r^\perp by angle θ , the classic point pair feature (PPF) satisfies:*

$$PPF(\mathbf{p}_r, \mathbf{p}_j(0)) = PPF(\mathbf{p}_r, \mathbf{p}_j(\theta)). \quad (7)$$

Proof. Recall the definition:

$$PPF(\mathbf{p}_r, \mathbf{p}_j) = (\|d\|_2, \cos(\alpha_1), \cos(\alpha_2), \cos(\alpha_3)), \quad (8)$$

$$\alpha_1 = \angle(\partial_r, d), \quad \alpha_2 = \angle(\partial_j, d), \quad \alpha_3 = \angle(\partial_r, \partial_j),$$

where $d = \mathbf{p}_j - \mathbf{p}_r$. For any vectors $a, b \in \mathbb{R}^3$, $\angle(a, b) = \arccos\left(\frac{\langle a, b \rangle}{\|a\|_2 \|b\|_2}\right)$. Let $R(\theta) \in SO(3)$ be the rotation matrix around axis ∂_r^\perp . Since \mathbf{p}_r lies on the axis of $R(\theta)$ and $R(\theta)^\top R(\theta) = I$, then:

$$R(\theta)\mathbf{p}_r = \mathbf{p}_r,$$

$$\langle R(\theta)\mathbf{p}_j, \mathbf{p}_r \rangle = \langle \mathbf{p}_j, R(\theta)^\top \mathbf{p}_r \rangle = \langle \mathbf{p}_j, \mathbf{p}_r \rangle, \quad (9)$$

$$\|R(\theta)\mathbf{p}_j\|_2 = \|\mathbf{p}_j\|_2.$$

Thus, we have

$$\begin{aligned} \|d(\theta)\|_2 &= \|R(\theta)\mathbf{p}_j - \mathbf{p}_r\|_2 \\ &= \sqrt{\|R(\theta)\mathbf{p}_j\|_2^2 + \|\mathbf{p}_r\|_2^2 - 2\langle R(\theta)\mathbf{p}_j, \mathbf{p}_r \rangle} \\ &= \sqrt{\|\mathbf{p}_j\|_2^2 + \|\mathbf{p}_r\|_2^2 - 2\langle \mathbf{p}_j, \mathbf{p}_r \rangle} \\ &= \|d\|_2, \\ \alpha_1(\theta) &= \angle(\partial_r, d(\theta)) = \arccos\left(\frac{\langle \partial_r, d(\theta) \rangle}{\|\partial_r\|_2 \|d(\theta)\|_2}\right) = \alpha_1 \\ \alpha_2(\theta) &= \angle(\partial_j(\theta), d(\theta)), \\ &= \arccos\left(\frac{\langle \partial_j(\theta), d(\theta) \rangle}{\|\partial_j(\theta)\|_2 \|d(\theta)\|_2}\right) = \alpha_2 \\ \alpha_3(\theta) &= \angle(\partial_r, \partial_j(\theta)), \\ &= \arccos\left(\frac{\langle \partial_r, \partial_j(\theta) \rangle}{\|\partial_r\|_2 \|\partial_j(\theta)\|_2}\right) = \alpha_3. \end{aligned} \quad (10)$$

Therefore,

$$PPF(\mathbf{p}_r, \mathbf{p}_j(0)) = PPF(\mathbf{p}_r, \mathbf{p}_j(\theta)), \quad (11)$$

which completes the proof. □




Figure 1. Illustration of the *Spatial Arrangement Ambiguity*. When a neighboring point $\mathbf{p}_j(\theta)$ rotates around the primary axis $\partial_1(\mathbf{p}_r)$ of the LRF by a rotation $R(\theta)$, the resulting pairwise geometric descriptor $PPF(\mathbf{p}_r, \mathbf{p}_j(\theta))$ remains unchanged. This demonstrates that classical PPF is invariant to such spatial transformations, leading to ambiguity in local geometric representation.

C Degeneracy for Static Global Rotation

We analyze the failure case when the global rotation R_g used to generate the anchor $\mathbf{p}_r^\dagger = \mathbf{p}_r R_g$ is static, and coincides with the equivalent rotation R_e that make local receptive fields identically as Figure 2. This highlights the necessity of the proposed dynamic global rotation mechanism.

Theorem 3. *If the global rotation satisfies $R_g = R_e$, then for any point p_j in the receptive field $\Gamma(\mathbf{p}_r)$, the SiPPF feature*

$$Ga4DPF(\mathbf{p}_r) = (4DPPF(\mathbf{p}'_r, \mathbf{p}'_j), 4DPPF(\mathbf{p}_r^\dagger, \mathbf{p}'_j)) \quad (12)$$

is invariant to replacing $\Gamma(\mathbf{p}_r)$ with its rotated counterpart $\Gamma(\mathbf{p}_r)R_e$.

Proof. Let $\mathbf{p}_j^E = \mathbf{p}_j R_e$ be the identical counterpart of \mathbf{p}_j under the equivalent rotation R_e . Since $R_e = R_g$, we have:

$$\mathbf{p}_r^\dagger = \mathbf{p}_r R_g = \mathbf{p}_r R_e. \quad (13)$$

Then the identical anchor becomes:

$$\mathbf{p}_r^{\dagger E} = (\mathbf{p}_r R_g) R_e = (\mathbf{p}_r R_e) R_e = \mathbf{p}_r^\dagger R_e. \quad (14)$$

This means the anchor of \mathbf{p}_r^E is exactly the identical counterpart of \mathbf{p}_r^\dagger :

$$\mathbf{p}_r^{\dagger E} = \mathbf{p}_r^\dagger R_e. \quad (15)$$

The equality *does not hold in general*, but is specific to the degenerate case $R_g = R_e$. From the rotation invariance of $4DPPF$ (Equation 5), we have:

$$\begin{aligned} 4DPPF(\mathbf{p}'_r, \mathbf{p}'_j) &= 4DPPF(\mathbf{p}_r^{\prime E}, \mathbf{p}_j^{\prime E}), \\ 4DPPF(\mathbf{p}_r^\dagger, \mathbf{p}'_j) &= 4DPPF(\mathbf{p}_r^{\dagger E}, \mathbf{p}_j^{\prime E}). \end{aligned} \quad (16)$$




Figure 2. Illustration of the *Degeneracy for Static Global Rotation*. When the global rotation R_g coincides with the equivalent rotation R_e , the Ga4DPF becomes insensitive to global pose clue from the anchor \mathbf{p}_r^\dagger , failing to distinguish the local receptive field $\Gamma(\mathbf{p}_r)$ (Left Wingtip) from its identical counterpart $\Gamma(\mathbf{p}_r)R_e$ (Right Wingtip).

Now plug into the Ga4DPF definition:

$$\begin{aligned}
 \text{Ga4DPF}(\mathbf{p}_r) &= (4\text{DPPF}(\mathbf{p}'_r, \mathbf{p}'_j), 4\text{DPPF}(\mathbf{p}_r^\dagger, \mathbf{p}'_j)) \\
 &= (4\text{DPPF}(\mathbf{p}_r^E, \mathbf{p}_j^E), 4\text{DPPF}(\mathbf{p}_r^\dagger, \mathbf{p}_j^E)) \\
 &= \text{Ga4DPF}(\mathbf{p}_r^E)
 \end{aligned} \tag{17}$$

This implies that if we replace a local receptive field $\Gamma(\mathbf{p}_j)$ with its identical counterpart $\Gamma(\mathbf{p}_j)R_e$, and simultaneously apply R_e to \mathbf{p}_r and its anchor \mathbf{p}_r^\dagger , the resulting Ga4DPF remains unchanged. Hence, Ga4DPF is invariant to symmetric or identical transformation when $R_g = R_e$, and fails to distinguish symmetric regions such as mirrored wings or chair legs. \square

D Bingham distribution on the hypersphere of unit quaternions

D.1 Bingham distribution

The probability density function of *Bingham* distribution [2, 8] for a unit quaternion $\mathbf{q} \in \mathbb{S}^3$ with parameter matrices \mathbf{V} and $\mathbf{\Lambda}$ is:

$$\mathcal{B}(\mathbf{q} | \mathbf{V}, \mathbf{\Lambda}) = \frac{1}{F(\mathbf{\Lambda})} \exp(\mathbf{q}^T \mathbf{V} \mathbf{\Lambda} \mathbf{V}^T \mathbf{q}), \tag{18}$$

where $\mathbf{\Lambda} = \text{diag}(\lambda_1, \lambda_2, \lambda_3, 0)$ is a diagonal matrix with negative parameters ($\lambda_1 \leq \lambda_2 \leq \lambda_3 < 0$), and the columns of the matrix $\mathbf{V} \in \mathbb{R}^{4 \times 4}$ are orthogonal unit vectors \mathbf{v} . $F(\mathbf{\Lambda})$ is the normalizing constant.

D.2 Entropy of Bingham distribution

Theorem 4. *The entropy of a Bingham distribution with PDF f is given by:*

$$h(f) = \log F - \mathbf{\Lambda} \cdot \frac{\nabla F}{F}. \tag{19}$$

Proof. Denote $C = \mathbf{V} \mathbf{\Lambda} \mathbf{V}^T$ for a more succinct form.

$$\begin{aligned}
 h(f) &= - \int_{\mathbf{q} \in \mathbb{S}^3} f(\mathbf{q}) \log f(\mathbf{q}) \\
 &= - \int_{\mathbf{q} \in \mathbb{S}^3} \frac{1}{F} \exp(\mathbf{q}^T C \mathbf{q}) (\mathbf{q}^T C \mathbf{q} - \log F) \\
 &= \log F - \frac{1}{F} \int_{\mathbf{q} \in \mathbb{S}^3} \mathbf{q}^T C \mathbf{q} \exp(\mathbf{q}^T C \mathbf{q}).
 \end{aligned}$$

Using $g(\mathbf{\Lambda})$ to represent the hyperspherical integral $\int_{\mathbf{q} \in \mathbb{S}^3} \mathbf{q}^T C \mathbf{q} \exp(\mathbf{q}^T C \mathbf{q})$,

$$\begin{aligned}
 g(\mathbf{\Lambda}) &= \int_{\mathbf{q} \in \mathbb{S}^3} \sum_{i=1}^3 \lambda_i (\mathbf{v}_i^T \mathbf{q})^2 \exp\left(\sum_{j=1}^3 \lambda_j (\mathbf{v}_j^T \mathbf{q})^2\right) \\
 &= \sum_{i=1}^3 \lambda_i \frac{\partial F}{\partial \lambda_i} = \mathbf{\Lambda} \cdot \nabla F.
 \end{aligned}$$

Thus, the entropy is $\log F - \mathbf{\Lambda} \cdot \frac{\nabla F}{F}$ as $\mathcal{L}_{\text{bingham}}$ in the main paper. \square

D.3 Construction of \mathbf{V} and $\mathbf{\Lambda}$

We randomly initialize a 7-d vector $(\mathbf{z}_1, \mathbf{z}_2)$ where the first 4-d vector \mathbf{z}_1 are normalized as a quaternion first and then following *Birdal Strategy* [3] to construct the parameter metric \mathbf{V} as:

$$\mathbf{V}(\mathbf{z}_1) \triangleq \begin{bmatrix} z_{11} & -z_{12} & -z_{13} & z_{14} \\ z_{12} & z_{11} & z_{14} & z_{13} \\ z_{13} & -z_{14} & z_{11} & -z_{12} \\ z_{14} & z_{13} & -z_{12} & -z_{11} \end{bmatrix}. \tag{20}$$

The last 3-d vector \mathbf{z}_2 is first activated by the *softplus* function ϕ , and then transformed into a parameter matrix $\mathbf{\Lambda}$ with decreasing diagonal values obtained through the negative accumulative sum:

$$\begin{aligned}
 \lambda_1 &= -\phi(z_{21}) - \phi(z_{22}) - \phi(z_{23}) \\
 \lambda_2 &= -\phi(z_{21}) - \phi(z_{22}) \\
 \lambda_3 &= -\phi(z_{21}).
 \end{aligned} \tag{21}$$

D.4 Acceptance Rejection Sampling for Bingham distribution

Kent and Ganeiber [11] demonstrated that samples from the Bingham distribution can be generated by employing the ACG distribution as an envelope density within an acceptance-rejection sampling framework. Building on this approach, the following algorithm can be applied to simulate a Bingham-based latent space with the parameter matrix $\mathbf{A} = \mathbf{V} \mathbf{\Lambda} \mathbf{V}^T$:

1. Set acceptance list \mathbf{L} , $\mathbf{\Psi}^{-1} = \mathbf{I} + \frac{2}{b} \mathbf{A}$, and assume that there exists a constant $\mathbf{M}^* \geq \text{upper}\{\frac{f^*(\mathbf{q}_n)}{g^*(\mathbf{q}_n)}\}$;

Algorithm 1 Enhancing Rotation-Invariant 3D Learning with Global Pose Awareness and Attention Mechanisms

Input: Point cloud dataset $\mathcal{D} = \{P_i\}$, network f_θ , total epochs T

Output: Trained network parameters θ

- 1: Initialize network parameters θ
 - 2: Initialize Bingham parameters $\mathbf{V}, \mathbf{\Lambda}$
 - 3: Sample global rotation quaternion:
 $q_{mode} \sim \mathcal{B}(q \mid \mathbf{V}, \mathbf{\Lambda})$
 - 4: Convert to rotation matrix: $R_g = \text{Rot}(q_{mode})$
 - 5: **for** each epoch $t = 1, 2, \dots, T$ **do**
 - 6: Fix R_g for all mini-batches in epoch t
 - 7: **for** each mini-batch $B \subset \mathcal{D}$ **do**
 - 8: **for** each reference point $\mathbf{p}_r \in B$ **do**
 - 9: Compute anchor: $\mathbf{p}_r^\dagger = \mathbf{p}_r R_g$
 - 10: Extract Ga4DPF for each neighbor $p_j \in \Gamma(p_r)$:

$$\text{Ga4DPF}(\mathbf{p}_r) = (4\text{DPPF}(\mathbf{p}_r', \mathbf{p}_j'), 4\text{DPPF}(\mathbf{p}_r^\dagger, \mathbf{p}_j'))$$
 - 11: **end for**
 - 12: Forward pass: $\hat{y} = f_\theta(B; R_g)$
 - 13: Compute task loss: $\mathcal{L}_{\text{task}} = \mathcal{L}(\hat{y}, y)$
 - 14: Compute total loss:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{task}} + \delta \cdot \sqrt{(\mathcal{L}_{\text{bingham}} - 0.1 \cdot \mathcal{L}_{\text{task}})^2}$$
 - 15: Update model parameters: $\theta \leftarrow \theta - \eta \nabla_\theta \mathcal{L}_{\text{total}}$
 - 16: **end for**
 - 17: Update Bingham:
 - 18: Resample:
 $q_g = \text{GeometricMean}\{q_i\}_{i=1}^{N_s} \sim \mathcal{B}(q \mid \mathbf{V}, \mathbf{\Lambda})$
 - 19: Update $R_g = \text{Rot}(q_g)$
 - 20: **end for**
 - 21: **return** θ
-

2. Draw \mathbf{u} from $\text{Uniform}(0, 1)$ and candidate values \mathcal{Q} from the ACG distribution on spheres with parameter matrices $\mathbf{\Psi}$;

3. Update acceptance list \mathbf{L} , where $\mathbf{u} < \frac{f^*(\mathbf{q}_n; \mathbf{A}_n)}{M^* g^*(\mathbf{q}_n; \mathbf{\Psi}_n)}$, to accept \mathcal{Q} . Go to Step 1 if not $\mathbf{L.all}()$.

Here, $f^*(\mathbf{q}_n; \mathbf{A}_n) = \exp(-\mathbf{q}_n^T \mathbf{A}_n \mathbf{q}_n)$, and $g^*(\mathbf{q}_n; \mathbf{\Psi}_n) = (\mathbf{q}_n^T \mathbf{\Psi}_n^{-1} \mathbf{q}_n)^{-2}$. $\mathbf{\Psi}$ and \mathbf{A} are the set of parameter matrices of the ACG and Bingham distribution. Empirically, setting $b = 1$ as default works well in many situations. The introduction of the acceptance list \mathbf{L} has significantly accelerated the speed of parallel sampling and optimization across the entire network.

D.5 Why Bingham distribution?

There exist multitudes of approaches to represent probability distributions of rotation [10], including geometric meth-

ods [1] and measure-theoretic approaches [6]. The quaternion is continuous and minimally redundant parameterizations without singularities. Therefore, it is an adequate choice when the objective of tasks is direct regression rather than tracking [5, 9]. Indeed, the antipodal symmetry of the Bingham distribution [2] makes it particularly well-suited for capturing and explaining the unique topology of quaternions. There have been extensive applications of Bingham distributions to model probability distributions on quaternions [3, 4, 7]. As such, we adopt quaternions as our representation, with the Bingham distribution forming the basis for constructing our probabilistic rotation model. This approach, compared to using the matrix Fisher distribution to model rotation matrices, is mathematically more convenient for establishing the subsequent latent space.

E Epoch-wise Bingham Update Strategy

To ensure global consistency and rotation-invariance across mini-batches during training, we introduce an *epoch-wise Bingham update strategy*. Specifically, we employ a shared global rotation R_g to generate the anchor $\mathbf{p}_r^\dagger = \mathbf{p}_r R_g$, enabling relative pose computation with consistent reference directions.

Rather than sampling R_g from a uniform distribution $\mathcal{U}(\text{SO}(3))$ in each iteration, we learn a task-adaptive rotation distribution using the Bingham distribution.

At the beginning of each epoch, a global rotation anchor R_g is sampled via $\{q_i\}_{i=1}^{N_s} \sim \mathcal{B}(q \mid \mathbf{V}, \mathbf{\Lambda})$ and fixed across all mini-batches in that epoch. This consistent reference facilitates the learning of rotation-invariant representations, while still allowing dynamic updates across epochs.

The Bingham parameters $(\mathbf{V}, \mathbf{\Lambda})$ are updated throughout training via the following total loss:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{task}} + \delta \cdot \sqrt{(\mathcal{L}_{\text{bingham}} - 0.1 \cdot \mathcal{L}_{\text{task}})^2}, \quad (22)$$

which encourages the Bingham mode (i.e., the most likely rotation) to align with the task optimization direction.

The complete training procedure is summarized in Algorithm 1. This formulation balances local invariance and global discriminability of patch-level pose, and helps address ambiguities induced by symmetric shapes or local geometric degeneracies.

F Additional Experiments

F.1 Robustness to the Number of Neighbours

We investigate the impact of the neighbourhood size k in k NN-based local feature grouping, which determines the receptive field for each point during feature extraction. To evaluate the model’s robustness to different k values, we perform experiments on ModelNet40 [13] and ScanObjectNN [12] under the z/SO(3) setting. As summarized in

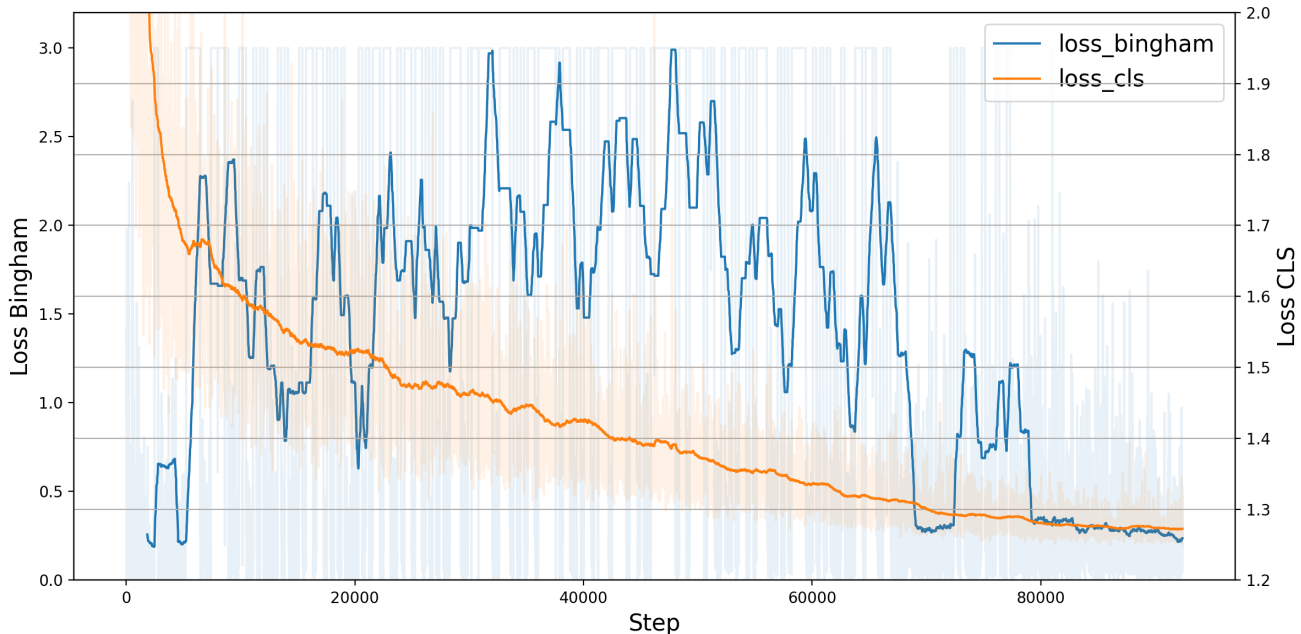


Figure 3. The Bingham loss $\mathcal{L}_{\text{bingham}}$ and classification loss \mathcal{L}_{cls} in the training of ModelNet40 with normals and z/SO(3) setting. $\mathcal{L}_{\text{bingham}}$ is optimized with \mathcal{L}_{cls} through our joint loss design, allowing to effectively learn a global consistent rotation R_g .

Table 1, the model consistently achieves high classification accuracy across a broad range of k , with the best performance at $k = 20$. Accuracy remains competitive for both smaller ($k = 10$) and larger ($k = 40$) neighbourhood sizes, indicating that the method is robust to variations in local receptive configuration.

Number of neighbors	$k=10$	$k=20$	$k=40$
ModelNet40 (pc)	91.7	91.9	91.9
ModelNet40 (pc+n)	92.5	92.8	92.7
ScanObjectNN	85.5	87.4	86.8

Table 1. Classification accuracy (%) under different neighbor size k on ModelNet40 and ScanObjectNN with z/SO(3) setting.

F.2 Performance Under Different Hyperparameter δ

To assess the robustness of our method with respect to the hyperparameter δ , we conduct experiments varying its value in the loss function (Equation 22), where δ controls the trade-off between the task loss $\mathcal{L}_{\text{task}}$ and the pose regularization term $\mathcal{L}_{\text{bingham}}$. As reported in Table 2, the model achieves consistently high classification accuracy across a wide range of δ values (0.6 to 1.0), with optimal performance observed at $\delta = 0.8$. The limited variance in performance across this range suggests that our approach is not overly sensitive to the choice of δ , demonstrating the stability of the overall formulation.

δ	0.6	0.7	0.8	0.9	1.0
ModelNet40 (pc)	91.5	91.7	91.9	91.8	91.8
ModelNet40 (pc+n)	92.4	92.7	92.8	92.7	92.6
ScanObjectNN	84.5	85.7	87.4	86.9	86.5

Table 2. Classification accuracy (%) under different hyperparameter δ on ModelNet40 and ScanObjectNN with z/SO(3) setting..

F.3 Learned Bingham Distribution for R_g

To understand the training dynamics of the Bingham distribution used to estimate the global rotation R_g , we analyze the evolution of the Bingham loss $\mathcal{L}_{\text{bingham}}$ alongside the classification loss \mathcal{L}_{cls} during training on ModelNet40 (with normals) under the z/SO(3) setting, as shown in Figure 3. The overall objective (Equation 22) is designed to couple $\mathcal{L}_{\text{bingham}}$ with the task loss $\mathcal{L}_{\text{task}}$, encouraging both to decrease in tandem. This coupling ensures that the optimization of the Bingham distribution is aligned with the task objective. As training progresses, the decline in $\mathcal{L}_{\text{task}}$ is mirrored by a corresponding reduction in $\mathcal{L}_{\text{bingham}}$, enabling the model to infer a stable and consistent global rotation R_g without explicit supervision.

References

- [1] Timothy D. Barfoot and Paul Timothy Furgale. Associating uncertainty with three-dimensional poses for use in estimation problems. *IEEE Trans. Robotics*, 30(3):679–693, 2014.

- [2] Christopher Bingham. An antipodally symmetric distribution on the sphere. *The Annals of Statistics*, pages 1201–1225, 1974. [3](#), [4](#)
- [3] Tolga Birdal, Umut Simsekli, Mustafa Onur Eken, and Slobodan Ilic. Bayesian pose graph optimization via bingham distributions and tempered geodesic mcmc. *Advances in neural information processing systems*, 31, 2018. [3](#), [4](#)
- [4] Mai Bui, Tolga Birdal, Haowen Deng, Shadi Albarqouni, Leonidas Guibas, Slobodan Ilic, and Nassir Navab. 6d camera relocalization in ambiguous scenes via continuous multimodal inference. In *European Conference on Computer Vision*, pages 139–157. Springer, 2020. [4](#)
- [5] Benjamin Busam, Tolga Birdal, and Nassir Navab. Camera pose filtering with local regression geodesics on the riemannian manifold of dual quaternions. In *Proceedings of the IEEE international conference on computer vision workshops*, pages 2436–2445, 2017. [4](#)
- [6] Luca Falorsi, Pim De Haan, Tim R Davidson, and Patrick Forré. Reparameterizing distributions on lie groups. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 3244–3253. PMLR, 2019. [4](#)
- [7] Jared Glover and Leslie Pack Kaelbling. Tracking the spin on a ping pong ball with the quaternion bingham filter. In *2014 IEEE international conference on robotics and automation (ICRA)*, pages 4133–4140. IEEE, 2014. [4](#)
- [8] Jared Marshall Glover. *The quaternion Bingham distribution, 3D object detection, and dynamic manipulation*. PhD thesis, Massachusetts Institute of Technology, 2014. [3](#)
- [9] F. Sebastian Grassia. Practical parameterization of rotations using the exponential map. *J. Graphics, GPU, & Game Tools*, 3(3):29–48, 1998. [4](#)
- [10] Adrian Haarbach, Tolga Birdal, and Slobodan Ilic. Survey of higher order rigid body motion interpolation methods for keyframe animation and continuous-time trajectory estimation. In *2018 International Conference on 3D Vision (3DV)*, pages 381–389. IEEE, 2018. [4](#)
- [11] John T Kent, Asaad M Ganeiber, and Kanti V Mardia. A new method to simulate the bingham and related distributions in directional data analysis with applications. *arXiv preprint arXiv:1310.8110*, 2013. [3](#)
- [12] Mikaela Angelina Uy, Quang-Hieu Pham, Binh-Son Hua, Thanh Nguyen, and Sai-Kit Yeung. Revisiting point cloud classification: A new benchmark dataset and classification model on real-world data. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 1588–1597, 2019. [4](#)
- [13] Zhirong Wu, Shuran Song, Aditya Khosla, Fisher Yu, Linguang Zhang, Xiaoou Tang, and Jianxiong Xiao. 3d shapenets: A deep representation for volumetric shapes. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1912–1920, 2015. [4](#)