SPHERICAL PARAMETRIZED MANIFOLD CONVOLUTION

Anonymous authors
Paper under double-blind review

ABSTRACT

In geometric deep learning, manifold convolution aims to extend the convolution operator from CNNs to non-Euclidean domains. In this field, equivariance plays an important role as the target property that allows convolution filters to be applied at different locations on the manifold; for CNNs, this corresponds to translation symmetry. In curved spaces such as smooth 2D surfaces, this symmetry group is no longer valid. Therefore, the goal of equivariant convolution theory is to bridge this gap. Current approaches typically either impose restrictive topological or symmetry assumptions to obtain exact equivariance, or abandon global symmetry in favor of local geometric constructions. The former limits applicability to real-world objects with diverse shapes, while the latter must address gauge equivariance and often requires expensive differential-geometric computations. In this work, we propose to combine both worlds: equivariant convolution on general manifolds and the canonical geometry of the sphere, which offers closed-form formulas for computing connections, exponential maps, and geodesics. We use a conformal spherical parametrization, which introduces minimal distortion while allowing to precisely understand the approximation errors. The kernels are polar splines applied across rotational orientation banks, which satisfy equivariance under discrete rotations, while the pullback to the manifold accounts for the local distortion induced by the spherical parametrization. We study the behavior of our method through targeted experiments on a custom synthetic dataset, and show that the convolution is approximately equivariant under rotations. Further, our method maintains competitive performance under deformation, compared to other methods, while offering greater flexibility.

1 INTRODUCTION

The central problem in generalizing Convolutional Neural Networks (CNNs) to manifolds is how to define a convolution that remains consistent across locations despite changes in local orientation and geometry. From this, two major challenges arise. First, manifold convolution suffers from rotational ambiguity of gauges, since there is generally no canonical way to align local reference frames across points on the manifold (Weiler et al., 2021; Cohen et al., 2019a). Imposing such an alignment introduces an arbitrary directional biasFeng et al. (2019). Second, the convolution must adapt to variations in local geometry while still allowing meaningful parameter sharing across the domain. Addressing the first challenge requires steerable kernels, whose transformation law compensates for the action of the symmetry group on manifold signals so that convolution commutes with group actions (Schonsheck et al., 2022; Poulenard & Ovsjanikov, 2018; Dagès et al., 2025). Addressing the second challenge requires tools from differential geometry to account for the distortions induced by the manifold structure. Solving these problems would extend the expressive power of CNNs beyond Euclidean data to manifold-structured domains. This is not merely a methodological extension, but a necessary step toward applying learning algorithms in a geometrically consistent way across many scientific and practical domains.

There are two main methodological directions for defining convolutions on non-Euclidean domains. The first consists of manifold convolutions, which define filters directly on the manifold or on its tangent spaces. For example, Fey et al. (2018) uses continuous kernels adapted to mesh geometry, but does not provide equivariance. Earlier approaches, define local convolutions using geodesic constructions, but these can be computationally expensive (Masci et al., 2018; Dagès et al., 2024). In

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

many cases, rotational ambiguity is addressed either by restricting the kernel to be isotropic, which limits the ability to learn directional patterns, or by using aggregation procedures such as max-pooling over orientations. Some methods like Poulenard & Ovsjanikov (2018); Haan et al. (2021); Wiersma et al. (2020) provide more effective equivariant architectures, but their constructions remain geometrically and numerically complex and can be sensitive to discretization and mesh quality. The second direction consists of group convolutions, which define convolution on canonical spaces with known symmetries. Examples include convolution using the rotation group $SO(3)$Cohen et al. (2018); Esteves et al. (2018), spectral formulations in Fourier spaceKondor et al. (2018), 3D steerable filters (Weiler et al., 2018), and continuous kernels on Lie groups (Bekkers, 2021). These methods achieve strong or exact equivariance by construction, but they are restricted to canonical geometries and do not naturally transfer to other genus-0 surfaces.

Gauge equivariant formulations represent a theoretically complete treatment of equivariance on manifolds (Weiler et al., 2026; Cohen et al., 2019b; Kondor & Trivedi, 2018), grounding convolution in the language of steerable kernels and frame bundles. The gauge equivariant formulation is an extended theory that covers manifolds and group CNNs (Gerken et al., 2023; Cohen et al., 2019a; Aronsson, 2022). However, this theoretical construction comes at a practical cost as these methods assume a known local structure determined by the G-structure of the manifold, which makes scalability a significant challenge. Efficient implementations exist only for specific geometries, such as the sphere, the icosahedron, and the Möbius strip (Cohen et al., 2019a; Defferrard et al., 2020; Mitchel et al., 2022). Extending these methods to general irregular meshes remains an open and largely unresolved problem. Furthermore, gauge equivariant methods require the data to exhibit symmetries tied to the underlying manifold, limiting their applicability when such a structure is absent or difficult to determine.

To address the challenge of generalizing Convolutional Neural Networks (CNNs) to manifolds, we propose an equivariant convolution on the sphere that achieves approximate equivariance for genus-zero surfaces by incorporating a conformal spherical parametrization. In other words, we use the sphere as a proxy manifold to perform convolutions on arbitrary closed and orientable surfaces, taking advantage of the closed-form equations for differential objects on the sphere. We perform the convolution on the tangent space using a continuous kernel inspired by B-spline CNNs (Fey et al., 2018; Bekkers, 2021). We propose a fixed support set of tangent vectors arrange on symmetrical rings that sample the sphere through a consistent geometric pipeline using inverse differential maps (Wiersma et al., 2020). Our approach combines the symmetry of the sphere with the flexibility of spline-based kernels, providing an efficient and generalizable convolutional framework for genus-0 surfaces. The proposed construction yields approximate equivariance under rotations induced on the sphere and improves robustness to non-linear geometric deformations of the underlying surface.

1. We propose a novel proxy-sphere framework for manifold convolution on genus-zero surfaces, combining spherical parametrization with continuous tangent-space kernels.
2. We provide a systematic empirical evaluation, including benchmarks against existing manifold and equivariant CNNs.

2 Related work

The present work sits exactly between surface CNNs and Gauge Equivariant CNNs, and closes a gap that is only partially addressed by both approaches.

2.1 Surface CNN

The Surface CNNs literature argues that standard Euclidean CNNs cannot work directly on surfaces, because surfaces are non-uniform, irregular domains; therefore, the usual CNN idea of translating the same kernel over a regular grid cannot be applied directly (Bronstein et al., 2017). Surface CNNs are often defined as convolution on local surface neighborhoods, as geodesic patches or tangent plane parametrization (Masci et al., 2018; Poulenard & Ovsjanikov, 2018; Yang et al., 2020; Hu et al., 2022; Sharp et al., 2020). Other surface learning methods use direct mesh neighborhoods and diffusion propagation rather than tangent convolution (Maron et al., 2017; Hanocka et al., 2019; Mitchel et al., 2021). In the case of local patch methods, they suffer from rotational ambiguity and are sensitive to effects of discretizations or local coordinate noise. SplineCNN offers a flexi-

ble alternative among intrinsic convolution methods, which uses continuous B-spline kernels over pseudo-coordinates (Fey et al., 2018). We adapt this idea of a continuous kernel to the tangent space over tangent coordinates. SplineCNNs, however, do not provide equivariance by construction, so if the local coordinates or gauges rotate, the kernel response changes unless the model learns this behavior from data. Bekkers (2021) addresses this limitation through equivariant B-spline constructions, although not in the full setting of surface convolutions on curved manifolds considered here. Similarly, prior work using circular harmonics can impose a predictable transformation law under rotations (Wiersma et al., 2020), while gauge equivariance is achieved by explicit construction of anisotropic filters on surfaces in (Poulenard & Ovsjanikov, 2018), yet one must keep track of orientation and transport features consistently between tangent spaces. Our method follows this, but specializes in genus-0 manifolds by using a spherical parametrization $\eta : M \to S^2$ as a proxy domain. This way, pmConv combines the flexibility of spline kernels with the geometric idea that local tangent-space filters should transform in a controlled way under gauge changes.

2.2 Spherical convolution

Spherical CNNs and other group CNNs exploit the sphere's homogeneous SO(3) symmetry to obtain exact or strong equivariance, but they also show the main limitation that motivates our work: the method is naturally defined on S^2 or $SO(3)$, and the conversion from a general surface to a spherical signal is usually treated as a preprocessing step rather than as part of the convolutional geometry. The theoretical basis for this line of work comes from Kondor & Trivedi (2018), who observe that equivariant linear maps for compact group actions are essentially convolutions, formalized for compact groups. This idea was then developed in spherical deep learning by Cohen et al. (2018) and Esteves et al. (2018), who define learnable kernels using the group structure of $SO(3)$, so that spherical convolution becomes the natural analogue of planar convolution for signals on the sphere. In a related spectral direction, Kondor et al. (2018) propose Clebsch-Gordan Nets, building spherical neural networks directly in Fourier space using irreducible representations of $SO(3)$ and Clebsch-Gordan products to combine features. Together, these works are important because they show that equivariance can be made exact when the domain and the symmetry group are cleanly matched. The same property, however, also imposes a restriction: the sphere has a global $SO(3)$ action, while a general closed orientable surface does not generally admit the same global symmetry.

Convolution on spheres is explored in Coors et al. (2018), adapting the sampling locations of planar convolutional filters to the sphere for omnidirectional images, explicitly compensating for spherical image distortions and enabling transfer from perspective CNNs. However, their motivation is spherical image processing rather than convolution on arbitrary genus-0 surfaces. Möbius convolution is closer to our setting, since it studies conformal equivariance on the sphere and evaluates on genus-zero shapes through conformal mapping; nevertheless, its filters are not spline-based, and its main focus remains Möbius-equivariance on spherical signals, rather than a simple spline-based framework that can be applied directly to arbitrary closed and orientable surfaces (Mitchel et al., 2022). B-splines on Lie groups instead define coordinates through logarithmic maps, rather than as offsets on a graph, providing a more flexible equivariant convolution given that the convolution itself is defined on group (Bekkers, 2021). Consequently, B-spline kernel can be made equivariant, provided one supplies the right transformation law for its support coordinates. Our approach aligns with this, using spline interpolation as the practical kernel parametrization on tangent coordinates, where equivariance depends on the local gauge rather than on a global Lie group. In summary, our method is positioned between these lines of work: we keep the practical continuous-kernel spirit of SplineCNN, but use the spherical parametrization $\eta : M \to S^2$ as a proxy to define tangent-space convolution on genus-0 manifolds. The sphere here serves purely as a geometric reference domain, not as the target of the signal processing itself, while local gauge changes induced by rotations, parametrization, and deformation are still accounted for.

3 Background

Parametrized manifold convolution (pmConv) is an equivariant formulation of convolution for signals on manifolds that uses the sphere as a geometric proxy. Rather than performing differential-geometric computations directly on the manifold, pmConv carries them out on the sphere, whose geometry is simple and well understood (Figure 1). As in a classical CNN, the convolution at a

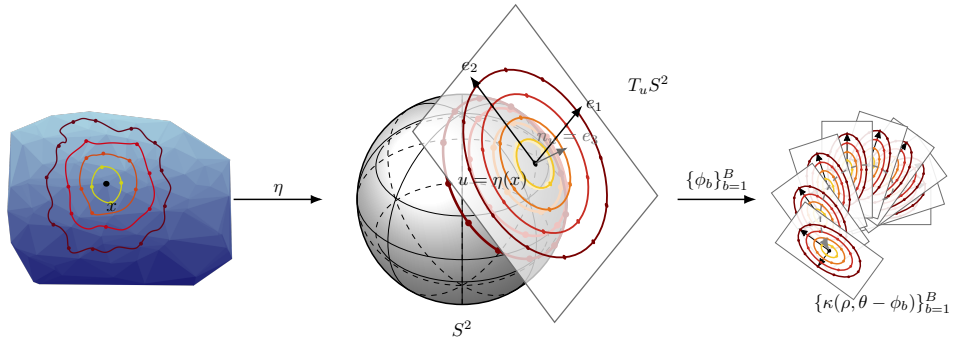

Figure 1: Overview of the pmConv construction. A point $x \in M$ is mapped to $u = \eta(x) \in S^2$. The polar support of tangent vectors $w \in T_uS^2$ samples points through the exponential map $\text{Exp}_u^S(w)$ and the inverse of η^{-1}. A rotational bank of copies of the kernel evaluation are stored as $\{\kappa(\rho, \theta - \phi_b)\}$ for $b \in [1, B]$.

point $x \in M$ is computed as the correlation between a kernel and the signal $f : M \to \mathbb{R}$ over a local support around p. For mathematical convenience, the kernel is defined on the tangent plane of the sphere, and the local support on this tangent plane is mapped onto M, where the signal is evaluated. The resulting change of domain is handled by a spherical conformal map $\eta : M \to S^2$, which allows us to move between the manifold and the sphere at each stage of the computation as needed.

3.1 Preliminaries

Spherical parametrization Let (M, ξ) be a compact, orientable two-dimensional Riemannian manifold discretized as a triangular mesh, with M the manifold and ξ the Riemannian metric $\xi_x : T_xM \times T_xM \to \mathbb{R}$. We use the spherical parametrization introduced in (Yueh et al., 2019) denoted as:

$$\eta : M \to S^2, \tag{3.1}$$

which defines a bijective map η between M and S^2. In reality, η is a piecewise-linear function on each triangle of the mesh and is affine in barycentric coordinates. The differential $d\eta_x : T_xM \to T_{\eta(x)}S^2$ has full-rank, meaning it is invertible for all $x \in M$. Further, it is a harmonic function with a fixed boundary meaning that η minimizes the Dirichlet energy:

$$E(\eta) = \int_M \|\nabla \eta(x)\|^2 \, dx, \tag{3.2}$$

which in practice yields an approximately conformal parametrization. This property is crucial because it enables a formalized analysis of the errors introduced by spherical map discretization on the equivariance of the final convolution operation.

Convolution support tangent vectors. The convolution is evaluated on the neighborhood on the tangent space of the sphere for a point $u \in S^2$ mapped from $x \in M$:

$$u = \eta(x) \in S^2. \tag{3.3}$$

In principle, we do the convolution on the tangent space of a proxy manifold, in this case S^2, hence we sample points on the sphere by the exponential map from a stencil on T_uS^2, given a local sphere tangent frame $\{e_1(u), e_2(u), e_3(u)\}$ in $\mathbb{R}^3$ such that $e_3(u) = u$ and $\{e_1(u), e_2(u)\}$ spans T_uS^2. Any $w \in T_uS^2$ is written as:

$$w = \alpha^1 e_1(u) + \alpha^2 e_2(u), \qquad \alpha(w; x) = (\alpha^1, \alpha^2) \in \mathbb{R}^2. \tag{3.4}$$

We define a fixed stencil on a regular polar grid that divides the circle into angular sections K_l at fixed radii ρ_l thus for a ring l, the tangent stencil is:

$$\alpha_{\ell m} = \rho_l \begin{bmatrix} \cos(2\pi m / K_l) \\ \sin(2\pi m / K_l) \end{bmatrix} \tag{3.5}$$

Canonical sampling on the sphere. In pmConv, the sphere is used as the canonical computational domain. Therefore, a tangent displacement is first interpreted on $T_u S^2$, then mapped to the sphere through the spherical exponential, and finally pulled back to M. This is defined as:

$$\pi_x(w) := \eta^{-1}\big(\mathrm{Exp}_u^{S^2}(w)\big), \qquad w \in T_u S^2. \tag{3.6}$$

For the unit sphere, we can write this in closed-form:

$$\mathrm{Exp}_u^{S^2}(w) = \cos(\|w\|)\, u + \sin(\|w\|)\, \frac{w}{\|w\|}, \qquad w \neq 0, \tag{3.7}$$

the condition $\mathrm{Exp}_u^{S^2}(0) = u$ is forced to avoid the generally invalid identity. Now, if the stencil lived in the $T_x M$, the correct mapping of the tangent vectors on v would be to find a value S^2 via the push-forward $d\eta : T_x M \to T_u S^2$ as:

$$\eta(\mathrm{Exp}_x^M(v)) = \mathrm{Exp}_{\eta(x)}^{S^2}(d\eta_x(v)), \tag{3.8}$$

which differs from our mapping $\pi_x(w)$ and the equivalence only holds under much stronger assumptions, such as local isometry. We analyze this for η as a conformal map, so the push-forward is reduced to a combination of scale and rotation, i.e. a linear transformation, at least locally, see Appendix A.

Parallel transport of feature fields. A correct formulation of convolution requires transporting features on M between two points $x, q \in M$, as features at different points live in different fibers of a frame bundle, a notion introduced by Weiler et al. (2021) and analogous to the tangent bundle in differential geometry. Consequently, the feature $f(\mathrm{Exp}_x(v))$ must be parallel transported from $q := \mathrm{Exp}_x^M(v)$ to x before it is correlated with the kernel at x.

The precise mathematical approach is to transport those features on geodesics on M. However, a transport can be used as an approximation in this case to take advantage of the simpler geometry of the sphere, and the transport is written as:

$$\Pi_{q\to x}^M \approx (d\eta_x)^{-1} \Pi_{\eta(q)\to\eta(x)}^{S^2} d\eta_q, \tag{3.9}$$

where $\Pi_{q\to x}^M$ is the parallel transport of the feature over the geodesic on M from the point q to x and likewise for $\Pi_{\eta(q)}^{S^2}$. This defines a diagram:

$$\begin{array}{ccc} T_q M & \xrightarrow{\;d\eta_q\;} & T_{\eta(q)} S^2 \\ {\scriptstyle \Pi_{q\to x}^M}\big\downarrow & & \big\downarrow{\scriptstyle \Pi_{\eta(\mathbf{q})\to\eta(x)}^{S^2}} \\ T_x M & \xleftarrow[(d\eta_x)^{-1}]{} & T_{\eta(x)} S^2 \end{array} \tag{3.10}$$

This is a reasonable construction based on the fact that η is conformal, and on local domains the tangent vectors are transformed linearly by a scaling factor and a rotation $d\eta_x = s_x R_x$. The error of approximation is studied in the Appendix A.

The parallel transport can be reduced to linear transform by representation theory $(\mathbb{R}^C, \rho)$, hence $\rho : G \to GL(C)$, for details, see Chapters 8 and 11 of Weiler et al. (2021). In essence, parallel transport induces a frame change on the feature bundle. The notation for the group representation of the transport in the literature is given as $\rho(g_{x\leftarrow q}^A) f^A(q)$, we use this $\Pi_{q\to x}^{\rho,M} f^A(q)$ with the hyper-index to keep equations a bit more compact, but it is the same linear map acting on feature fields $\Pi_{q\to x}^{\rho,M} : \mathbb{R}^C \to \mathbb{R}^C$.

3.2 Parametrized manifold convolution

Now we are ready to define the convolution operator of the parametrized manifold convolution. Let $f : M \to \mathbb{R}^{C_{\text{in}}}$. Given a compact support set $\Omega_u^{S^2} \subset T_u S^2$, the parametrized manifold convolution is defined by:

$$(\mathrm{pmConv}\, f)(x) = \int_{\Omega_u^{S^2}} \kappa_x(w) \Pi_{x'\to x}^{\rho,M} f\big(\pi_x(w)\big) dw, \tag{3.11}$$

where κ_x is a learnable kernel defined on the tangent plane $T_{\eta(x)}S^2$. This equation adapts the gauge equivariant convolution Cohen et al. (2019a) to our setting. The term $\Pi^{\rho;M}_{x'\to x}$ denotes the parallel transport of features from the point $x' = \mathrm{Exp}^M_x((d\eta_x)^{-1}(w))$ on M obtained via exponential map of the tangent vector $(d\eta_x)^{-1}(w)$, back to the point x, where the convolution is calculated. We have defined the point x' in the equation 3.6, but adjusted the notation to obtain a succinct equation.

Relation to the tangent-space on M. If one starts from a kernel k_x on T_xM, then by the linear change of variables $v = (d\eta_x)^{-1}w$ one may write

$$\kappa_x(w) = k_x\big((d\eta_x)^{-1}w\big)\,\big|\det(d\eta_x)^{-1}\big|. \tag{3.12}$$

The kernel is evaluated at the tangent coordinate $w \in \Omega^{S^2}_u$. Each such w corresponds to a tangent vector $(d\eta_x)^{-1}w \in \Omega^M_x \subset T_xM$, so the kernel's support is defined indirectly on T_xM through this pulled-back set. Note that, because η is only approximately conformal, $d\eta_x$ is not exactly a pure scale-and-rotation map: its non-conformal component introduces local anisotropic stretching and compression. Consequently, a fixed compact support region on $T_{\eta(x)}S^2$ does not correspond to a fixed isotropic region on T_xM. For spline kernels with compact support, this geometric limitation causes mismatch on the intended and effective support.

Spline Kernel We define a two 1-dimensional B-spline basis $B^{(r)}_\mu(\rho)$ and $B^{(\theta)}_\nu(\theta)$ with indices μ and ν, whose tensor product defines a 2D-basis. The tensor product spline is evaluated at w in polar coordinates, thus $\alpha(w;x) = (\rho,\theta)$. Each product has a matrix of learnable weights $\Theta_{\mu\nu} \in \mathbb{R}^{C_{\text{out}}\times C_{\text{in}}}$, but only a few of them are active at the given coordinate, depending on m the degree of the splines as $(m+1)^2$. Note that for simplicity, we write them as sum; please see the Appendix B for the implementation details. Hence, for an input channel i and output channel o is written

$$\kappa_{x,io}(\rho,\theta) = \sum_{\mu=1}^{k_r}\sum_{\nu=1}^{k_\theta} \Theta_{\mu\nu,io}\, B^{(r)}_\mu(\rho) B^{(\theta)}_\nu(\theta), \tag{3.13}$$

where $\Theta \in \mathbb{R}^{(k_r k_\theta)\times C_{in}\times C_{out}}$, k_r and k_θ are the number of kernel basis for the radial r and angular θ coordinates, respectively. This is based on the spline convolution proposed in Fey et al. (2018), which we detail in Appendix B, consistent with our implementation.

The equivariant solution uses a discrete set of rotated copies of the kernel. We use the idea from Freeman et al. (1991); Cohen & Welling (2016) but apply it to splines instead. Accordingly, each branch b has a kernel $\kappa^{(b)}_{io}(\rho,\theta) = \kappa_{io}(\rho,\theta-\varphi_b)$ with $\varphi_b = \frac{2\pi b}{B}$.

4 Discrete implementation on meshes

Consider the mesh represented as the graph $\mathcal{G} = (\mathcal{V},\mathcal{E},\mathcal{F})$ with vertices $x_i \in \mathcal{V}$, edges as pairs of indices $(i,j) \in \mathcal{E}$, and faces as triples of $(i,j,k) \in \mathcal{F}$; for simplicity, we sometimes write x instead of x_i. In the discrete setting, the spherical parametrization is simply a vertex-correspondence map that preserves the edges and faces of the original mesh; thus $u := \eta(x)$ is a vertex of the sphere mesh. The convolution stencil in T_uS^2 is discretized by an angular index $m \in [1, K_l]$ and a radial index $l \in [1, L]$, as in equation 3.5. The convolution points are given by the support tangent vectors w_{lm} and their projection onto the manifold, $x'_{lm} := \pi_x(w_{lm})$, as defined in equation 3.6. Note that the point x'_{lm} lies in into a triangle $\sigma_{lm} \in \mathcal{F}$ with vertex indices (p,q,r), and the signal value is computed by a barycentric interpolation:

$$f(x'_{lm}) = \sum_{i\in\{p,q,r\}} \lambda_{i,lm}\, f(x_{i,lm}). \tag{4.1}$$

The convolution is then discretized in terms of the support values on the mesh, f_{lm}, on each orientation branch, indexed by $b \in [1,B]$, and each output channel $o \in [1, C_{out}]$, as

$$(\mathrm{pmConv}\, f)_{xbo} := \sum_{l=1}^{L}\sum_{m=1}^{K_l} \Pi^{M,\rho}_{x'_{lm}\to x}\, f\big(x'_{lm}\big)\, \kappa_x\big(w_{lm}(x)\big)\, \omega_{lm}, \tag{4.2}$$

where ω_{lm} denotes the quadrature weight for the contribution of the convolution at the support vector w_{lm}. We take the quadrature to be proportional to $\sin(r_l)\,\Delta r\,\Delta\theta$ in polar coordinates, the natural area element on the sphere, see Appendix D. This weighting ensures that the discrete sum approximates a unit-normalized integral over the spherical tangent support.

The transport term $\Pi^M_{x'_{lm}\to x}$, which denotes the parallel transport of the feature from x'_{lm} to x along the chosen path on M, is given by the approximation in equation 3.9. In terms of the discretized coordinates, this is written as

$$\Pi^M_{x'_{lm}\to x} \approx (d\eta_x)^{-1}\,\Pi^{S^2}_{\eta(x'_{lm})\to\eta(x)}\,d\eta_{x'_{lm}}, \tag{4.3}$$

where $\Pi^{S^2}_{\eta(x'_{lm})\to\eta(x)}$ is the parallel transport on S^2 between the spherical images of the sampled point and the convolution point. On the sphere, the parallel transport has a closed-form expression in terms of two 2D tangent frames, given as a rotation matrix. The two Jacobians in equation 4.3 are computed from the tangent frames of the face σ_{lm}'s vertices, by taking the Moore-Penrose pseudoinverse of the edge vectors of that face (see Appendix C).

The parallel transport of the features yields an approximate $SO(2)$ rotation, obtained from the transport map $M_{lm} = \Pi^M_{x'_{lm}\to x}$ and a corresponding rotation angle τ_{lm}. For the orientation banks, this rotation is quantized to the nearest permutation shift, $s_{lm} = \lfloor \tau_{lm} B/2\pi \rfloor \bmod B$, forming an element of the cyclic representation C_B. We can then write the parallel transport of the orientation-bank feature as

$$\Pi^{M,\rho}_{x'_{lm}\to x}\, f_{lm,b} = f_{lm,\,b-s_{lm} \bmod B}. \tag{4.4}$$

In this formulation, we obtain an approximate equivariance, with three main sources of error. We analyze these in Appendix A and establish the result formally in Proposition A.2. It is worth noting that, although much of the differential-geometry computation is heavy, it admits closed-form solutions that speed up the calculations; moreover, many of these values are constant and can be precomputed during data preprocessing. Since η is piecewise linear on the mesh, both η and its inverse can be evaluated triangle-wise via fast barycentric interpolation. Likewise, the differential $d\eta_x$ is constant on each triangle in the piecewise-linear approximation, so the pushforward and pullback operators in equation 3.9 can be computed efficiently on a per-triangle basis.

5 Experiments

5.1 MNIST on manifolds

We first evaluated our model on classifying handwritten digits from the MNIST dataset, projected onto the manifold via the tangent map at each point, following a benchmark established in previous work on equivariant convolutional networks (Cohen et al., 2018; Kondor et al., 2018; Haan et al., 2021; Cohen et al., 2019a). We use a subset of 1000 handwritten-digit images, projected onto the template manifold T. In case $\eta = \mathrm{id}$, we use a unit sphere with Fibonacci sampling; for the other cases, the template is an ellipsoid. Following the construction of Cohen et al. (2018), each image is mapped by perspective projection onto a Driscoll-Healy S^2 grid (Driscoll & Healy, 1994), with a bandwidth of 30, resulting in a spherical grid. The points are linearly interpolated from the polar coordinates of the template mesh with respect to the center of mass.

Starting from the template $D : T \to M$, we introduce systematic deformations based on soft As-Rigid-As-Possible (ARAP) deformations (Sorkine et al., 2007). This allows us to investigate the errors of the spherical parametrization. We also include large irregular deformations, as in Haan et al. (2021), by adding noise to the vertices and applying a smooth Laplacian filter (Desbrun et al., 1999). We define five cases of deformation at different parameter settings.

Table 1 reports different levels of deformation and the evaluation under three gauge-rotation regimes: the identity transformation e, full rotations on $SO(3)$, and rotations restricted to a discrete subgroup of the stabilizer group of the z-axis, $\mathrm{Stab}(e_z) \simeq SO(2)$. We discretize the group to the size of the rotational bank used, $C_{12} \leq \mathrm{Stab}(e_z)$. Because GEM-CNN's implementation uses ResNet blocks, we implement both a 3-single-layer variant and a 3-block ResNet architecture.

Across the deformation cases, we observe that pmConv and the GEM-CNN baseline achieve comparable performance, but only with the matching ResNet backbone. In the case of the simple layered

Table 1: MNIST classification accuracy (%) under increasing geometric deformation and three gauge transformations: identity (e), arbitrary $SO(3)$ rotations, and discrete rotations $C_{12} \leq \mathrm{Stab}(e_z)$. Results are reported as mean $\pm$ standard deviation. In **bold-black** best value, and **bold-blue** second best value.

	None			Smooth			Random								
							σ=2.0			σ=3.0			σ=4.5		
Model	e	$SO(3)$	C_{12}	e	$SO(3)$	C_{12}	e	$SO(3)$	C_{12}	e	$SO(3)$	C_{12}	e	$SO(3)$	C_{12}
SplineCNN	73.5±2.8	12.8±4.1	34.9±0.7	26.8±1.9	24.0±2.3	23.7±0.8	25.5±2.6	22.8±4.3	24.1±3.6	25.3±3.3	22.5±3.5	24.1±4.0	24.7±3.8	24.3±2.8	24.9±3.4
GCN	24.8±1.8	24.8±1.8	24.8±1.8	24.7±2.1	24.7±2.1	24.7±2.1	24.8±2.8	24.8±2.8	24.8±2.8	24.3±3.3	24.3±3.3	24.3±3.3	25.0±3.1	25.0±3.1	25.0±3.1
GEM-CNN-ResNet	**92.5**±1.0	**92.5**±1.0	**92.5**±1.0	**92.3**±4.3	**92.3**±4.3	**92.3**±4.3	**91.0**±3.5	**91.0**±3.5	**91.0**±3.5	**86.5**±3.1	**86.5**±3.1	**86.5**±3.1	65.2±4.9	65.2±4.9	65.2±4.9
GEM-CNN	61.5±0.9	61.5±0.9	61.5±0.9	72.2±4.2	72.2±4.2	72.2±4.2	65.3±1.0	65.3±1.0	65.3±1.0	56.7±3.3	56.7±3.3	56.7±3.3	39.3±8.7	39.3±8.7	39.3±8.7
MLP	23.7±2.4	23.7±2.4	23.7±2.4	24.8±1.4	24.8±1.4	24.8±1.4	24.2±3.5	24.2±3.5	24.2±3.5	24.3±3.3	24.3±3.3	24.3±3.3	24.7±3.1	24.7±3.1	24.7±3.1
pmConv	**94.0**±2.0	**91.5**±3.0	**94.0**±2.0	85.0±2.6	83.5±3.8	85.0±2.6	83.3±2.5	82.0±2.8	83.3±2.6	83.2±2.8	83.2±3.1	83.2±2.8	**82.2**±3.5	**80.5**±1.8	**82.2**±3.5
pmConv-ResNet	**94.2**±1.8	88.0±5.2	**94.2**±1.8	**88.8**±4.8	**84.7**±3.5	**88.8**±4.8	**85.5**±2.2	**82.3**±4.3	**85.5**±2.2	**88.0**±1.4	**85.0**±1.4	**88.1**±1.3	**84.3**±2.9	**80.5**±2.3	**84.2**±2.8

Table 2: Orientation regression performance under two gauge regimes: identity e and full $SO(3)$ rotation. Lower MAE and $1-\cos$ indicate better orientation prediction.

	e		$SO(3)$		**Change**	
Model	**MAE**↓	$\mathbf{1-\cos}$↓	**MAE**↓	$\mathbf{1-\cos}$↓	$\mathbf{\Delta_{e,SO(3)}}$**MAE**↓	$\mathbf{\Delta_{e,SO(3)}}(\mathbf{1-\cos})$↓
SplineCNN	$0.460_{\pm 0.031}$	$0.487_{\pm 0.052}$	$0.737_{\pm 0.046}$	$0.919_{\pm 0.066}$	−2.763e−01	−4.320e−01
GCN	$0.751_{\pm 0.006}$	$0.945_{\pm 0.014}$	$0.751_{\pm 0.006}$	$0.945_{\pm 0.014}$	−1.589e−07	−2.384e−07
GEM-CNN-ResNet	$0.144_{\pm 0.023}$	$0.096_{\pm 0.031}$	$0.144_{\pm 0.023}$	$0.096_{\pm 0.031}$	1.490e−08	−9.934e−09
GEM-CNN	$0.569_{\pm 0.066}$	$0.656_{\pm 0.107}$	$0.569_{\pm 0.066}$	$0.656_{\pm 0.107}$	5.960e−08	−1.987e−08
MLP	$0.759_{\pm 0.017}$	$0.963_{\pm 0.027}$	$0.759_{\pm 0.017}$	$0.963_{\pm 0.027}$	0.000e+00	0.000e+00
pmConv	$0.225_{\pm 0.017}$	$0.179_{\pm 0.030}$	$0.224_{\pm 0.014}$	$0.178_{\pm 0.026}$	4.248e−04	5.899e−04
pmConv-ResNet	$0.176_{\pm 0.011}$	$0.121_{\pm 0.012}$	$0.177_{\pm 0.011}$	$0.122_{\pm 0.012}$	−6.409e−04	−1.111e−03

convolution, pmConv remains the strongest, which supports the idea that our model attains good results with fewer parameters. Under rigid and moderately noisy deformations, GEM-CNN (ResNet) generally obtains the highest accuracy; however, under the strongest deformation, $\sigma = 4.5$, GEM-CNN (ResNet) drops to $65.2 \pm 4.9\%$, whereas pmConv (ResNet) retains $84.3 \pm 2.9\%$. The same behavior is visible without the ResNet backbone, where pmConv retains $82.3 \pm 2.8\%$ compared to $39.3 \pm 8.7\%$ for GEM-CNN.

Moreover, pmConv proved to be stable under changes of gauge, with only small variation across the different regimes, as expected for $SO(3)$ rotations; the invariance is exact, however, for the discrete C_{12} case. GEM-CNN preserves exact equivariance, as rotational transformations are isometries. This is consistent with the results reported in Haan et al. (2021).

5.2 Anisotropic Gaussian Orientation

The second experiment is a controlled equivariance task. We generate an anisotropic signal on the tangent space at a randomly selected point, projected using the spherical logarithmic map from the spherical coordinates of the deformed mesh; details are provided in the Appendix E.

The models are trained to regress the angular direction of the anisotropic Gaussian with respect to the z-axis, given the notion of a gauge reference. Now, if we rotate the gauge and do not change the signal projection, the model should compensate for the gauge transformation, and the local tangent frames should not affect the convolution values; hence the last MLP layer perceives the same input signal and therefore computes the same angle value. This gives us a controlled approach to test for gauge equivariance and to check directly whether the rotated gauge affects the prediction. Henceforth, we can measure equivariance directly from the angular error between the predicted direction and the known transformed direction.

Table 2 shows the performance on the controlled orientation regression task, using the metrics $1-\cos$ and circular MAE, which account for the π-symmetry of the anisotropic Gaussian. GEM-CNN (ResNet) achieves the lowest error, with an MAE of 0.144 ± 0.023 and $1-\cos = 0.096 \pm 0.031$, and keeps these values exactly the same under rotation. The strongest variant, pmConv (ResNet), also accurately recovers the orientation, reaching an MAE of 0.176 ± 0.011 with $1-\cos = 0.121 \pm 0.012$,

and exhibits only a small change under rotation. Importantly, the pmConv variants remain stable under changes of gauge, with only small differences between the identity and $SO(3)$ conditions. In contrast, the other baselines either show consistently high error or deteriorate substantially under rotation. These results show that pmConv preserves directional information under gauge transformations, even though GEM-CNN still gives the strongest and exactly equivariant performance on this controlled task.

6 CONCLUSIONS

In this work, we introduce pmConv, a manifold convolution that uses a conformal spherical parametrization to take advantage of the simple geometry of the sphere and compute the local geometrical constructions involved in the convolution. As the conformal maps are locally acting on tangent vectors as stretching and rotation, we developed a formal analysis of the error of equivariance. We propose an orientation bank of polar splines that enable an approximated discretized rotation equivariance under isometric transformations of the gauge. The pre-computable geometry allows fast inference with a relatively simple architecture and can, without residual connections, already achieve competitive performance.

In the deformation experiments, pmConv retains accuracy at $\sigma = 4.5$ noise distortions, the baseline models degrade sharply. We argue that this is because the operations remain stable as the kernel is evaluated on the sphere. However, this needs to be confirmed by additional ablation experiments to the ones presented in Appendix F.

The orientation bank introduces a controlled equivariant exactness. In the implementation, the pmConv is configured with 12 orientations, and hence has perfect equivariance for rotation on the cyclic rotation group C_{12}. In the idealized discrete construction, such rotations correspond to cyclic permutations of the orientation channels; for arbitrary $SO(3)$ transformations, however, pmConv is only approximately equivariant. The distinction visible in the orientation regression confirms low relatively small equivariance errors. The remaining sources of error are specifically discretizations of conformal map as piece linear map, the approximation connection of tangent spaces, and the tangent quadrature errors.

An important limitation of pmConv is its topological assumption: the input needs to be a genus-0 surface. In addition to that, exact conformal maps do not always exist, or in some cases they can introduce very large stretching on tubular structures (Baden et al., 2018; Dorobantu et al., 2023). The spherical deformations are sometimes so critical that they introduce collapsed triangles with areas close to zero. Nevertheless, there exist several practically relevant problems in which signals are defined on not-so-strongly deforming closed surfaces. Potential applications include droplet, bubble, and vesicle dynamics, where scalar or vector fields evolve on deforming interfaces; cellular and membrane mechanics; cardiac mechanics, where subject- and pathology-dependent deformation is measured on ventricular surfaces; and hippocampal or other subcortical shape analysis in studies of neurodegeneration. Related ideas may also be applicable to signals defined directly on spherical domains, such as climate and atmospheric data. Evaluating pmConv on such real physical and biomedical manifold-valued problems is an important direction for future work.

AI USE STATEMENT

In this work, we used generative AI tools to assist in the writing of proofs, assist with translation. We have not used generative AI tools for theoretical models or conceptual frameworks, formulate mathematical claims, provide critical ingredients for proving mathematical claims, propose or refine hypotheses, support qualitative and thematic data analysis, interpret results; all of which are not applicable to this work. We used AI tools partially to:

1. Design or provide feedback on research methodology or experiments: more in the initial brain-storming phases.
2. Implement methods: as a code assistant, mostly for debugging, searching for performance gaps, debugging plots, and collecting and reformatting tables. The core design and implementation of the method was done manually.

3. Clean and reformat dataset: a few relevant fields in the data structure were suggested from manually designed metadata in the synthetic dataset generation.

Additionally, we used generative AI tools to draft parts of a research paper, transcribe research material like pictures of hand-written proofs, and discover related research papers on the relevant literature. We have reviewed all AI-assisted work. For example, we have checked for errors in the equations generated, cross-referenced literature with traditional tools like *Zotero*, and thoroughly read all cited papers from their original sources. The code was verified and audited by the author(s). We take responsibility for the final content of this work, including text, claims, or artifacts produced with the aid of generative AI. This disclosure follows the recommended text from the recommended AI Policy.

Reproducibility statement

The code will be made publicly available upon publication. For review, it is included in the supplementary material. The datasets can be generated using the commands documented in the accompanying `README.md`. Details are ommited from this version for anonymity.

Author Contributions

Omitted from this version for anonymity.

Acknowledgments

Omitted from this version for anonymity.

References

Jimmy Aronsson. Homogeneous vector bundles and G-equivariant convolutional neural networks. *Sampling Theory, Signal Processing, and Data Analysis*, 20(2):10, July 2022. ISSN 2730-5724. doi: 10.1007/s43670-022-00029-3. URL https://doi.org/10.1007/s43670-022-00029-3.

Alex Baden, Keenan Crane, and Misha Kazhdan. Möbius Registration. *Computer Graphics Forum*, 37(5):211–220, 2018. ISSN 1467-8659. doi: 10.1111/cgf.13503. URL https://onlinelibrary.wiley.com/doi/abs/10.1111/cgf.13503.

Erik J. Bekkers. B-Spline CNNs on Lie Groups, March 2021. URL http://arxiv.org/abs/1909.12057. arXiv:1909.12057 [cs].

Michael M Bronstein, Joan Bruna, Yann LeCun, Arthur Szlam, and Pierre Vandergheynst. Geometric deep learning: going beyond euclidean data. *IEEE Signal Processing Magazine*, 34(4):18–42, 2017.

T. S. Cohen, M. Weiler, B. Kicanaoglu, and M. Welling. Gauge Equivariant Convolutional Networks and the Icosahedral CNN. In *Proceedings of Machine Learning Research*, volume 97. PLMR, 2019a. URL http://proceedings.mlr.press/v97/cohen19d.html.

Taco Cohen and Max Welling. Group Equivariant Convolutional Networks. In Maria Florina Balcan and Kilian Q. Weinberger (eds.), *Proceedings of The 33rd International Conference on Machine Learning*, volume 48 of *Proceedings of Machine Learning Research*, pp. 2990–2999, New York, New York, USA, 20–22 Jun 2016. PMLR. URL https://proceedings.mlr.press/v48/cohenc16.html.

Taco Cohen, Mario Geiger, and Maurice Weiler. A General Theory of Equivariant CNNs on Homogeneous Spaces. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d' Alché-Buc, E. Fox, and R. Garnett (eds.), *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019b. URL https://proceedings.neurips.cc/paper_files/paper/2019/file/b9cfe8b6042cf759dc4c0cccb27a6737-Paper.pdf.

Taco S. Cohen, Mario Geiger, Jonas Koehler, and Max Welling. Spherical CNNs, February 2018. URL http://arxiv.org/abs/1801.10130. arXiv:1801.10130 [cs].

Benjamin Coors, Alexandru Paul Condurache, and Andreas Geiger. SphereNet: Learning Spherical Representations for Detection and Classification in Omnidirectional Images. In Vittorio Ferrari, Martial Hebert, Cristian Sminchisescu, and Yair Weiss (eds.), *Computer Vision – ECCV 2018*, pp. 525–541, Cham, 2018. Springer International Publishing. ISBN 978-3-030-01240-3. doi: 10.1007/978-3-030-01240-3_32.

Thomas Dagès, M. Lindenbaum, and A. Bruckstein. Metric Convolutions: A Unifying Theory to Adaptive Image Convolutions. In *2025 IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 13974–13984, 2024. doi: 10.1109/iccv51701.2025.01297.

Thomas Dagès, Michael Lindenbaum, and Alfred M Bruckstein. Metric Convolutions: A Unifying Theory to Adaptive Image Convolutions. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 13974–13984, 2025.

Michaël Defferrard, Martino Milani, Frédérick Gusset, and Nathanaël Perraudin. DeepSphere: a graph-based spherical CNN, December 2020. URL http://arxiv.org/abs/2012.15000. arXiv:2012.15000 [cs].

Mathieu Desbrun, Mark Meyer, Peter Schröder, and Alan H Barr. Implicit fairing of irregular meshes using diffusion and curvature flow. In *Proceedings of the 26th annual conference on Computer graphics and interactive techniques*, pp. 317–324, 1999.

Victor Dorobantu, Charlotte Borcherds, and Yisong Yue. Conformal Generative Modeling on Triangulated Surfaces. 2023. doi: 10.48550/ARXIV.2303.10251. URL https://arxiv.org/abs/2303.10251. Publisher: arXiv Version Number: 1.

J. R. Driscoll and D. M. Healy. Computing fourier transforms and convolutions on the 2-sphere. *Adv. Appl. Math.*, 15(2):202–250, June 1994. ISSN 0196-8858. doi: 10.1006/aama.1994.1008. URL https://doi.org/10.1006/aama.1994.1008.

Carlos Esteves, Christine Allen-Blanchette, Ameesh Makadia, and Kostas Daniilidis. Learning SO(3) Equivariant Representations with Spherical CNNs. In Vittorio Ferrari, Martial Hebert, Cristian Sminchisescu, and Yair Weiss (eds.), *Computer Vision – ECCV 2018*, pp. 54–70, Cham, 2018. Springer International Publishing. ISBN 978-3-030-01261-8. doi: 10.1007/978-3-030-01261-8_4.

Yutong Feng, Yifan Feng, Haoxuan You, Xibin Zhao, and Yue Gao. MeshNet: Mesh Neural Network for 3D Shape Representation. *Proceedings of the AAAI Conference on Artificial Intelligence*, 33(01):8279–8286, July 2019. ISSN 2374-3468. doi: 10.1609/aaai.v33i01.33018279. URL https://ojs.aaai.org/index.php/AAAI/article/view/4840.

Matthias Fey, Jan Eric Lenssen, Frank Weichert, and Heinrich Müller. Splinecnn: Fast geometric deep learning with continuous b-spline kernels. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 869–877, 2018.

William T Freeman, Edward H Adelson, et al. The design and use of steerable filters. *IEEE Transactions on Pattern analysis and machine intelligence*, 13(9):891–906, 1991.

Jan E. Gerken, Jimmy Aronsson, Oscar Carlsson, Hampus Linander, Fredrik Ohlsson, Christoffer Petersson, and Daniel Persson. Geometric deep learning and equivariant neural networks. *Artificial Intelligence Review*, 56(12):14605–14662, December 2023. ISSN 1573-7462. doi: 10.1007/s10462-023-10502-7. URL https://doi.org/10.1007/s10462-023-10502-7.

Pim de Haan, Maurice Weiler, Taco Cohen, and Max Welling. Gauge Equivariant Mesh CNNs: Anisotropic convolutions on geometric graphs, November 2021. URL http://arxiv.org/abs/2003.05425. arXiv:2003.05425 [cs].

Rana Hanocka, Amir Hertz, Noa Fish, Raja Giryes, Shachar Fleishman, and Daniel Cohen-Or. MeshCNN: A Network with an Edge. *ACM Transactions on Graphics*, 38(4):1–12, August 2019. ISSN 0730-0301, 1557-7368. doi: 10.1145/3306346.3322959. URL http://arxiv.org/abs/1809.05910. arXiv:1809.05910 [cs].

Shi-Min Hu, Zheng-Ning Liu, Meng-Hao Guo, Jun-Xiong Cai, Jiahui Huang, Tai-Jiang Mu, and Ralph R. Martin. Subdivision-Based Mesh Convolution Networks. *ACM Transactions on Graphics*, 41(3):1–16, June 2022. ISSN 0730-0301, 1557-7368. doi: 10.1145/3506694. URL http://arxiv.org/abs/2106.02285. arXiv:2106.02285 [cs.CV].

Risi Kondor and Shubhendu Trivedi. On the Generalization of Equivariance and Convolution in Neural Networks to the Action of Compact Groups, November 2018. URL http://arxiv.org/abs/1802.03690. arXiv:1802.03690 [stat].

Risi Kondor, Zhen Lin, and Shubhendu Trivedi. Clebsch–gordan nets: a fully fourier space spherical convolutional neural network. In *Advances in Neural Information Processing Systems*, volume 31, pp. 10138–10147, Montr{\'{e}}al, Canada, 2018. URL https://proceedings.neurips.cc/paper/2018/hash/a3fc981af450752046be179185ebc8b5-Abstract.html.

Haggai Maron, M. Galun, Noam Aigerman, M. Trope, Nadav Dym, Ersin Yumer, Vladimir G. Kim, and Y. Lipman. Convolutional neural networks on surfaces via seamless toric covers. *ACM Transactions on Graphics (TOG)*, 36:1 – 10, 2017. doi: 10.1145/3072959.3073616.

Jonathan Masci, Davide Boscaini, Michael M. Bronstein, and Pierre Vandergheynst. Geodesic convolutional neural networks on Riemannian manifolds, June 2018. URL http://arxiv.org/abs/1501.06297. arXiv:1501.06297 [cs].

Thomas W. Mitchel, Vladimir G. Kim, and M. Kazhdan. Field convolutions for surface cnns. In *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 9981–9991, 2021. doi: 10.1109/iccv48922.2021.00985.

Thomas W. Mitchel, Noam Aigerman, Vladimir G. Kim, and Michael Kazhdan. Möbius Convolutions for Spherical CNNs. In *ACM SIGGRAPH 2022 Conference Proceedings*, SIGGRAPH '22, pp. 1–9, New York, NY, USA, July 2022. Association for Computing Machinery. ISBN 978-1-4503-9337-9. doi: 10.1145/3528233.3530724. URL https://dl.acm.org/doi/10.1145/3528233.3530724.

Adrien Poulenard and Maks Ovsjanikov. Multi-directional geodesic neural networks via equivariant convolution. *ACM Trans. Graph.*, 37(6):236:1–236:14, December 2018. ISSN 0730-0301. doi: 10.1145/3272127.3275102. URL https://dl.acm.org/doi/10.1145/3272127.3275102.

Stefan C Schonsheck, Bin Dong, and Rongjie Lai. Parallel transport convolution: Deformable convolutional networks on manifold-structured data. *SIAM Journal on Imaging Sciences*, 15(1): 367–386, 2022. ISSN 1936-4954. Publisher: SIAM.

Nicholas Sharp, Souhaib Attaiki, Keenan Crane, and M. Ovsjanikov. Diffusionnet: Discretization agnostic learning on surfaces. *ACM Transactions on Graphics (TOG)*, 41:1 – 16, 2020. doi: 10.1145/3507905.

Olga Sorkine, Marc Alexa, et al. As-rigid-as-possible surface modeling. In *Symposium on Geometry processing*, volume 4, pp. 109–116, 2007.

Maurice Weiler, Mario Geiger, Max Welling, Wouter Boomsma, and Taco S Cohen. 3D Steerable CNNs: Learning Rotationally Equivariant Features in Volumetric Data. In *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc., 2018. URL https://proceedings.neurips.cc/paper/2018/hash/488e4104520c6aab692863cc1dba45af-Abstract.html.

Maurice Weiler, Patrick Forré, Erik Verlinde, and Max Welling. Coordinate Independent Convolutional Networks – Isometry and Gauge Equivariant Convolutions on Riemannian Manifolds, June 2021. URL http://arxiv.org/abs/2106.06020. arXiv:2106.06020 [cs].

Maurice Weiler, Patrick Forré, Erik Verlinde, and Max Welling. *Equivariant and Coordinate Independent Convolutional Networks: A Gauge Field Theory of Neural Networks*. WORLD SCIENTIFIC, February 2026. ISBN 978-981-98-0662-1 978-981-98-0663-8. doi: 10.1142/14143. URL https://www.worldscientific.com/worldscibooks/10.1142/14143.

Ruben Wiersma, Elmar Eisemann, and Klaus Hildebrandt. CNNs on Surfaces using Rotation-Equivariant Features. *ACM Transactions on Graphics*, 39(4), August 2020. ISSN 0730-0301, 1557-7368. doi: 10.1145/3386569.3392437. URL http://arxiv.org/abs/2006.01570. arXiv:2006.01570 [cs].

Yuqi Yang, Shilin Liu, Hao Pan, Yang Liu, and Xin Tong. PFCNN: Convolutional Neural Networks on 3D Surfaces Using Parallel Frames. June 2020. doi: 10.48550/arXiv.1808.04952. URL http://arxiv.org/abs/1808.04952. arXiv:1808.04952 [cs.CV].

Mei-Heng Yueh, Tiexiang Li, Wen-Wei Lin, and Shing-Tung Yau. A Novel Algorithm for Volume-Preserving Parameterizations of 3-Manifolds. *SIAM Journal on Imaging Sciences*, 12(2):1071–1098, January 2019. doi: 10.1137/18M1201184. URL https://epubs.siam.org/doi/10.1137/18M1201184.

A Analysis on pmConv equivariance

A.1 Approximate local rotational equivariance with isotropic kernel

The relevant symmetry property of pmConv is not exact global group equivariance on M, but rather a *local approximate rotational equivariance* induced by the spherical parametrization. The mechanism is local because the operator is defined through the parametrization η, its differential $d\eta_p$, local gauges on tangent spaces, and a discrete directional sampling. In particular, the geometric part of the argument depends on how rotations in $T_{\eta(p)}S^2$ are transported to T_pM.

Gauge notation. Let $p \in M$, and let $e_p : \mathbb{R}^2 \to T_pM$ be a local oriented orthonormal frame (gauge) at p. Likewise, let $e^{S^2}_{\eta(p)} : \mathbb{R}^2 \to T_{\eta(p)}S^2$ be a local oriented orthonormal frame at $\eta(p) \in S^2$. The representation of a tangent vector in the gauge e_p is denoted by coordinates in $\mathbb{R}^2$. For neighboring points $q, p \in M$, let

$$g_{p \leftarrow q} \in SO(2)$$

denote the gauge transition / parallel transport element mapping the gauge at q to the gauge at p. Let

$$\rho_{\text{in}} : SO(2) \to \mathrm{GL}(\mathcal{F}_{\text{in}})$$

be the representation acting on the input feature fiber.

pmConv as a transported tangent-space integral. Let $\eta : M \to S^2$ be a bijective spherical parametrization. Define

$$A_p := (d\eta_p)^{-1} : T_{\eta(p)}S^2 \to T_pM.$$

At a point $p \in M$, the continuous parametrized manifold convolution is defined locally by

$$(\text{pmConv } f)(p) = \int_{T_{\eta(p)}S^2} \tilde{k}_p(w)\, \rho_{\text{in}}(g_{p \leftarrow q(w)})\, f(q(w))\; dw, \qquad q(w) = \text{Exp}_p(A_p w), \quad \text{(A.1)}$$

where $\tilde{k}_p$ is compactly supported in a neighborhood of the origin such that $\text{Exp}_p(A_p w)$ is well-defined on the support of $\tilde{k}_p$. In the discrete implementation, this becomes

$$(\text{pmConv } f)(p) \approx \sum_{k=1}^{K} \tilde{k}_p(\tilde{w}_k(p))\; \rho_{\text{in}}(g_{p \leftarrow q_k(p)})\, f(q_k(p))\; \omega_k, \quad \text{(A.2)}$$

where

$$\tilde{w}_k(p) = w_{lm}(p), \qquad u = \eta(p), \quad \text{(A.3)}$$

and

$$q_k(p) = \text{Exp}_p(A_p \tilde{w}_k(p)). \quad \text{(A.4)}$$

Here $\{\tilde{w}_k(p)\}_{k=1}^{K} \subset T_uS^2$ are the fixed polar-grid stencil vectors w_{lm} defined in Eq. equation 3.5, with the index k enumerating the ring/angle pairs (l, m). Thus, the pmConv pipeline can be interpreted as follows:

1. choose a fixed polar stencil of directions directly in $T_{\eta(p)}S^2$,
2. map them to T_pM through A_p,
3. sample the manifold through Exp_p,
4. transport the sampled feature from the gauge at $q_k(p)$ to the gauge at p,
5. weight the transported features by a kernel in tangent coordinates.

Small conformal distortion model. The symmetry argument uses the assumption that $d\eta_p$ is close to a similarity. Concretely, assume that in local oriented orthonormal frames,

$$d\eta_p = s_p R_p + E_p, \qquad R_p \in SO(2), \qquad s_p > 0, \qquad \|E_p\| \le \varepsilon_p s_p, \qquad 0 \le \varepsilon_p < 1. \tag{A.5}$$

Equivalently, $d\eta_p$ has small conformal distortion. This is the precise hypothesis needed below; harmonicity of η alone does not imply equation A.5. Under equation A.5,

$$A_p = (d\eta_p)^{-1} = s_p^{-1} R_p^{-1} + \Delta_p, \tag{A.6}$$

with the quantitative bound

$$\|\Delta_p\| \le \frac{\varepsilon_p}{s_p(1-\varepsilon_p)}. \tag{A.7}$$

Transported tangent-plane rotations. For $Q \in SO(2)$, define the transported rotation on T_pM by

$$\widehat{Q}_p = R_p^{-1} Q R_p \in SO(2). \tag{A.8}$$

If $d\eta_p$ were exactly conformal, i.e.

$$d\eta_p = s_p R_p,$$

then

$$A_p Q = s_p^{-1} R_p^{-1} Q = (R_p^{-1} Q R_p)\, s_p^{-1} R_p^{-1} = \widehat{Q}_p A_p. \tag{A.9}$$

Thus, in the ideal conformal case, a rotation in $T_{\eta(p)}S^2$ is transported exactly to a rotation in T_pM. In the approximate case,

$$A_p Q = \big(s_p^{-1} R_p^{-1} + \Delta_p\big) Q \tag{A.10}$$

$$= \widehat{Q}_p\, s_p^{-1} R_p^{-1} + \Delta_p Q \tag{A.11}$$

$$= \widehat{Q}_p A_p + \Big[\Delta_p Q - \widehat{Q}_p \Delta_p\Big]. \tag{A.12}$$

Therefore,

$$\|A_p Q - \widehat{Q}_p A_p\| \le 2\|\Delta_p\| \le \frac{2\varepsilon_p}{s_p(1-\varepsilon_p)}. \tag{A.13}$$

Isotropic kernel. In this subsection we assume that the kernel is isotropic on the sphere-side tangent plane, i.e.

$$\tilde{k}_p(w) = \phi_p(\|w\|), \tag{A.14}$$

for some radial profile ϕ_p. Hence

$$\tilde{k}_p(Qw) = \tilde{k}_p(w), \qquad \forall\, Q \in SO(2). \tag{A.15}$$

This removes kernel anisotropy from the equivariance defect, so that the only remaining defects come from:

1. the small conformal distortion of η,
2. the finite directional quadrature induced by the fixed stencil sampling,
3. the discrete approximation of the transport term $\rho_{\mathrm{in}}(g_{p\leftarrow q})$.

Gauge transport. The transport term is essential. The feature $f(q)$ is expressed in the local gauge at q, whereas the kernel at p is expressed in the local gauge at p. The factor

$$\rho_{\mathrm{in}}(g_{p\leftarrow q})$$

transports the feature from the gauge at q to the gauge at p before aggregation. We assume that the discretized transport is consistent with Levi–Civita parallel transport up to an error $\varepsilon_{\mathrm{conn}}(p)$ on the relevant neighborhood.

Continuous local equivariance estimate. Let

$$\mathcal{T}^{\mathrm{src}}_{p,Q} f = \int_{T_{\eta(p)}S^2} \tilde{k}_p(w)\, \rho_{\mathrm{in}}(g_{p\leftarrow q_Q(w)})\, f(q_Q(w))\; dw, \tag{A.16}$$

$$\mathcal{T}^{\mathrm{tgt}}_{p,Q} f = \int_{T_{\eta(p)}S^2} \tilde{k}_p(w)\, \rho_{\mathrm{in}}(g_{p\leftarrow \widehat{q}_Q(w)})\, f(\widehat{q}_Q(w))\; dw, \tag{A.17}$$

where

$$q_Q(w) = \mathrm{Exp}_p(A_p Q w), \qquad \widehat{q}_Q(w) = \mathrm{Exp}_p(\widehat{Q}_p A_p w).$$

Assume that f is Lipschitz on the relevant geodesic neighborhood with constant L_f, that Exp_p is Lipschitz on the relevant subset of T_pM with constant $L_{\mathrm{Exp},p}$, and that the transport approximation is Lipschitz-stable on the same neighborhood. Then

$$\left\| \mathcal{T}^{\mathrm{src}}_{p,Q} f - \mathcal{T}^{\mathrm{tgt}}_{p,Q} f \right\| \le C_p \frac{2\varepsilon_p}{s_p(1-\varepsilon_p)} \int_{T_{\eta(p)}S^2} \|w\|\, |\tilde{k}_p(w)|\, dw + C'_p\, \varepsilon_{\mathrm{conn}}(p), \tag{A.18}$$

for constants C_p, C'_p depending on local Lipschitz bounds and the operator norm of ρ_{in}.

Effect of the fixed stencil sampling. The stencil directions $\tilde{w}_k(p)$ form only a finite discretization of the tangent plane. However, the polar-grid quadrature weights ω_{lm} are constructed to be proportional to the natural area element $\sin(r_l)\,\Delta r\,\Delta\theta$ on the sphere, so the discrete sum approximates the corresponding tangent-plane integral as the stencil is refined ($L \to \infty$, $K_l \to \infty$ for all l). Accordingly, for smooth test functions g, one expects a quadrature estimate of the form

$$\left| \int g(w)\, dw - \sum_{k=1}^{K} g(\tilde{w}_k(p))\, \omega_k \right| \le \varepsilon_{\mathrm{quad}}(K)\, \|g\|_{C^m}, \qquad \varepsilon_{\mathrm{quad}}(K) \to 0 \quad \text{as } K \to \infty, \tag{A.19}$$

where K denotes the total number of stencil points (summed over rings l and angular positions m).

Proposition A.1 (Approximate local rotational equivariance in the isotropic case)**.** *Assume that $\eta : M \to S^2$ is a bijective spherical parametrization whose differential satisfies*

$$d\eta_p = s_p R_p + E_p, \qquad R_p \in SO(2), \qquad s_p > 0, \qquad \|E_p\| \le \varepsilon_p s_p, \qquad 0 \le \varepsilon_p < 1.$$

Assume further that:

1. *the kernel is isotropic, i.e. $\tilde{k}_p(w) = \phi_p(\|w\|)$,*
2. *the fixed polar stencil directions satisfy a quadrature error bound of the form equation A.19,*
3. *the transport term is discretized with consistency defect $\varepsilon_{\mathrm{conn}}(p)$,*
4. *f and Exp_p are Lipschitz on the relevant neighborhood.*

Then the pmConv operator is approximately equivariant under local tangent-plane rotations. More precisely, the local equivariance defect is controlled by the three terms

$$\varepsilon_p, \qquad \varepsilon_{\mathrm{quad}}(K), \qquad \varepsilon_{\mathrm{conn}}(p),$$

up to constants depending on local Lipschitz, transport, and quadrature regularity bounds. In particular, the defect tends to zero in the limit

$$\varepsilon_p \to 0, \qquad K \to \infty, \qquad \varepsilon_{\mathrm{conn}}(p) \to 0. \tag{A.20}$$

Proof. At a point $p \in M$, the continuous local pmConv operator is given by Equation equation A.1. Under the small-distortion assumption, Equations equation A.12–equation A.13 show that rotating tangent directions on $T_{\eta(p)}S^2$ and then transporting them to T_pM differs from transporting first and then rotating in T_pM by a term of size $O(\varepsilon_p)$. By the Lipschitz continuity of f and Exp_p, this yields the continuous estimate equation A.18. Because the kernel is isotropic, it is invariant under local rotations of its argument, so no additional kernel anisotropy defect appears. The remaining continuous error comes from the geometric transport through A_p and from the approximation of parallel transport by $\rho_{\mathrm{in}}(g_{p\leftarrow q})$. In the discrete implementation, the fixed polar stencil provides only a finite directional quadrature, and the associated defect is controlled by $\varepsilon_{\mathrm{quad}}(K)$. Combining these three effects, the pmConv response under a local tangent-plane rotation differs from the ideally transported rotated response by terms controlled by the conformal distortion of η, the finite directional quadrature error, and the transport discretization defect. □

Remark. The argument above is local rather than a global group-equivariance proof on M. Its mechanism is:

1. small conformal distortion of η transports tangent-plane rotations on S^2 to approximately matching tangent-plane rotations on M,
2. the isotropic kernel removes dependence on local frame orientation at the kernel level,
3. the fixed polar stencil approximates isotropic directional integration,
4. the transport term $\rho_{\text{in}}(g_{p\leftarrow q})$ enforces gauge consistency of the aggregated features,
5. together these imply approximate local rotational equivariance of pmConv.

A.2 Rotated spline kernels and improved approximate equivariance

To promote directional sensitivity while retaining rotational structure, we now replace the isotropic kernel by a bank of rotated spline kernels indexed by a discrete rotation group. Let

$$G_n := C_n = \{R_{\theta_r} : \theta_r = 2\pi r/n,\ r = 0, \dots, n-1\} \subset SO(2)$$

denote the cyclic subgroup of planar rotations. Let

$$K : \mathbb{R}^2 \to \text{Hom}(\mathcal{F}_{\text{in}}, \mathcal{F}_{\text{out}})$$

be a base spline kernel defined in gauge coordinates. We define its rotated copies by

$$K_r(\xi) = K(R_{-\theta_r}\xi), \qquad r = 0, \dots, n-1. \tag{A.21}$$

The input feature field is lifted to orientation channels,

$$f : M \to \mathcal{F}_{\text{in}} \otimes \mathbb{R}^n,$$

equivalently written as

$$f : M \times G_n \to \mathbb{R}^{C_{\text{in}}},$$

where the second index corresponds to the orientation channel. At a point $p \in M$, let $u = \eta(p)$ and let

$$A_p = (d\eta_p)^{-1} : T_u S^2 \to T_p M.$$

Given the fixed stencil directions

$$\tilde{w}_k(p) = w_{lm}(p) \in T_u S^2, \qquad k = 1, \dots, K,$$

defined by the polar grid of Eq. equation 3.5 (with k enumerating the ring/angle pairs (l, m)), the corresponding sampling locations on the manifold are

$$q_k(p) = \text{Exp}_p(A_p \tilde{w}_k(p)).$$

Let $\alpha(\tilde{w}_k(p); p) \in \mathbb{R}^2$ denote the coordinates of $\tilde{w}_k(p)$ in the chosen gauge at $u = \eta(p)$. We define the rotated-kernel pmConv operator by

$$(\text{pmConv}\, f)(p, r) = \sum_{k=1}^{K} \sum_{s=0}^{n-1} K_{r-s}\big(\alpha(\tilde{w}_k(p); p)\big) \left[\rho_{\text{in}}(g_{p\leftarrow q_k(p)}) f(q_k(p))\right]_s \omega_k, \tag{A.22}$$

where the index $r - s$ is understood modulo n.

Discrete rotation action on orientation channels. For $t \in \{0, \dots, n-1\}$, define the action of G_n on orientation-indexed features by cyclic shift:

$$[\rho_{\text{reg}}(t) f](p, r) = f(p, r - t \bmod n). \tag{A.23}$$

The rotated kernel bank satisfies the discrete transformation rule

$$K_r(R_{\theta_t}\xi) = K_{r-t}(\xi), \qquad t, r \in \{0, \dots, n-1\}. \tag{A.24}$$

Why the full operator is still only approximately equivariant. The discrete kernel bank has an exact channel-permutation symmetry in the chosen local gauge. However, the full pmConv operator remains only approximately equivariant on the manifold for four reasons:

1. the geometric transport from $T_{\eta(p)}S^2$ to T_pM is only approximately similarity-preserving when $\varepsilon_p > 0$,
2. the continuous rotation group $SO(2)$ is approximated by the discrete subgroup G_n,
3. the fixed stencil quadrature remains approximate,
4. transport discretization remains approximate.

Thus, rotated copies improve the rotational modeling of the kernel, but they do not remove the geometric approximation errors coming from the parametrization and discretization. A convenient measure of the residual angular discretization is

$$\varepsilon_{\mathrm{ang}}(n) = \sup_{\theta\in[0,2\pi)} \min_{r\in\{0,\dots,n-1\}} |\theta - \theta_r|, \tag{A.25}$$

with angles understood modulo 2π. Clearly,

$$\varepsilon_{\mathrm{ang}}(n) \to 0 \qquad \text{as } n \to \infty.$$

Proposition A.2 (Approximate equivariance with rotated spline kernels)**.** *Assume that the hypotheses of Proposition A.1 hold, except that the isotropic kernel is replaced by the rotated kernel bank equation A.21, and suppose moreover that the angular discretization defect $\varepsilon_{\mathrm{ang}}(n)$ tends to zero as $n \to \infty$. Then the rotated-kernel pmConv operator equation A.22 is approximately equivariant under local rotations. More precisely, its local equivariance defect is controlled by*

$$\varepsilon_p, \qquad \varepsilon_{\mathrm{quad}}(K), \qquad \varepsilon_{\mathrm{conn}}(p), \qquad \varepsilon_{\mathrm{ang}}(n),$$

up to constants depending on local Lipschitz, transport, and quadrature regularity bounds and on the smoothness of the spline kernel.

Proof. The channel structure of the rotated kernel bank is exactly compatible with the discrete group G_n through Equation equation A.24. Thus, within a fixed gauge, discrete rotations of the kernel argument induce corresponding cyclic shifts of the orientation channels. However, the full operator still depends on the geometric transport $A_p = (d\eta_p)^{-1}$ and on the sampled manifold points $q_k(p) = \mathrm{Exp}_p(A_p\tilde{w}_k(p))$. As in Proposition A.1, the deviation between source-side and target-side rotations is controlled by the small conformal distortion parameter ε_p. In addition, replacing continuous rotations by the discrete set of rotated copies introduces an angular discretization error of size $\varepsilon_{\mathrm{ang}}(n)$. The finite stencil quadrature and the discretized transport contribute the additional terms $\varepsilon_{\mathrm{quad}}(K)$ and $\varepsilon_{\mathrm{conn}}(p)$. Combining these contributions shows that the full pmConv operator is approximately equivariant under local rotations, with defect controlled by the four stated terms. As

$$\varepsilon_p \to 0, \qquad K \to \infty, \qquad \varepsilon_{\mathrm{conn}}(p) \to 0, \qquad n \to \infty,$$

the equivariance defect tends to zero. □

Remark A.1 (Geometric interpretation)**.** The rotated kernel bank provides an exact discrete rotational symmetry at the channel level in the chosen local gauge, while the spherical parametrization provides the geometric transport between sphere-side and manifold-side tangent planes. When $d\eta_p$ has small conformal distortion, rotations in $T_{\eta(p)}S^2$ are transported to approximately matching rotations in T_pM. Hence the rotated copies improve directional modeling and sharpen the rotational symmetry of the kernel, while the overall equivariance of the full pmConv remains approximate due to geometric distortion, finite quadrature, transport discretization, and angular discretization.

Remark A.2 (Gauge dependence)**.** Both propositions are stated in the chosen local gauge used to define the coordinates $\alpha(\,\cdot\,; p)$. The transport term $\rho_{\mathrm{in}}(g_{p\leftarrow q})$ compensates for feature transport between neighboring points, but the presentation still fixes a local gauge for the kernel coordinates. A fully gauge-equivariant reformulation would additionally track explicit gauge changes at the kernel level. The present result should therefore be interpreted as approximate local rotational equivariance of the geometric operator in a fixed gauge, together with gauge-consistent feature transport.

B SPLINE KERNELS ON THE TANGENT SPACE

Let $u \in S^2$ be a point on the sphere and let $T_u S^2$ denote its tangent space. After choosing an orthonormal tangent frame at u, we identify $T_u S^2 \simeq \mathbb{R}^2$ and write a tangent vector as $q = (q^{(1)}, q^{(2)})$.

We define a compact support region in tangent coordinates as

$$\Omega = [a_1, b_1] \times [a_2, b_2] \subset T_u S^2 \simeq \mathbb{R}^2. \tag{B.1}$$

To express the spline basis on a normalized domain, let

$$\nu : \Omega \to [0, 1]^2 \tag{B.2}$$

denote the affine rescaling map. The map $B_j^{(n)} : [0, 1] \to \mathbb{R}$ defines a parametrized B-spline basis function of degree n. Suppose that k_1 and k_2 basis functions are defined along the two tangent-coordinate axes. We introduce the multi-index set as a way to select the $(m + 1)^2$ active splines at a given coordinate Fey et al. (2018)

$$\mathcal{I} = \{0, \dots, k_1 - 1\} \times \{0, \dots, k_2 - 1\}. \tag{B.3}$$

An element $m \in \mathcal{I}$ is a multi-index

$$m = (m_1, m_2), \tag{B.4}$$

where m_1 and m_2 select the B-spline basis functions along the first and second tangent-space coordinates, respectively.

The corresponding tensor-product basis function is

$$\Phi_m(q) = B_{m_1}^{(n)} \left(\nu^{(1)}(q) \right) B_{m_2}^{(n)} \left(\nu^{(2)}(q) \right), \qquad q \in \Omega. \tag{B.5}$$

Thus, $\mathcal{I}$ is the index set of the tensor-product B-spline basis functions.

Kernel parameterization. Let C_{in} and C_{out} denote the numbers of input and output channels, respectively. For each pair of channels $(c_{\text{in}}, c_{\text{out}})$, we define a continuous tangent-space kernel

$$g_{c_{\text{in}}, c_{\text{out}}} : \Omega \to \mathbb{R} \tag{B.6}$$

as a linear combination of the tensor-product B-spline basis functions,

$$g_{c_{\text{in}}, c_{\text{out}}}(q) = \sum_{m \in \mathcal{I}} W_{m, c_{\text{in}}, c_{\text{out}}} \Phi_m(q). \tag{B.7}$$

Equivalently, the learnable coefficients can be collected into the parameter tensor

$$W \in \mathbb{R}^{k_1 \times k_2 \times C_{\text{in}} \times C_{\text{out}}}. \tag{B.8}$$

The coefficient $W_{m, c_{\text{in}}, c_{\text{out}}}$ is therefore the learnable control value associated with the tensor-product basis element indexed by m and the channel pair $(c_{\text{in}}, c_{\text{out}})$.

Convolution in the tangent space. Let $f : S^2 \to \mathbb{R}^{C_{\text{in}}}$ be the input signal on the sphere. For a tangent vector $q \in T_u S^2$, the exponential map

$$\text{Exp}_u : T_u S^2 \to S^2 \tag{B.9}$$

maps q to the corresponding point on the sphere reached by following the geodesic starting at u with initial tangent vector q.

The signal around u can therefore be represented in tangent coordinates by the pullback

$$f_u(q) = f\left(\mathrm{Exp}_u(q)\right), \qquad q \in \Omega. \tag{B.10}$$

The tangent-space convolution at u is defined by integrating the local signal against the kernel over Ω,

$$y^{(c_{\mathrm{out}})}(u) = \sum_{c_{\mathrm{in}}=1}^{C_{\mathrm{in}}} \int_{\Omega} f^{(c_{\mathrm{in}})}\left(\mathrm{Exp}_u(q)\right) g_{c_{\mathrm{in}},c_{\mathrm{out}}}(q)\, dq. \tag{B.11}$$

Here, dq denotes the area element in the Euclidean tangent coordinates induced by the chosen orthonormal frame of T_uS^2.

Substituting the B-spline expansion of the kernel gives

$$y^{(c_{\mathrm{out}})}(u) = \sum_{c_{\mathrm{in}}=1}^{C_{\mathrm{in}}} \int_{\Omega} f^{(c_{\mathrm{in}})}\left(\mathrm{Exp}_u(q)\right) \left[\sum_{m\in\mathcal{I}} W_{m,c_{\mathrm{in}},c_{\mathrm{out}}} \Phi_m(q)\right] dq. \tag{B.12}$$

Using the definition of the tensor-product basis, this becomes

$$y^{(c_{\mathrm{out}})}(u) = \sum_{c_{\mathrm{in}}=1}^{C_{\mathrm{in}}} \int_{\Omega} f^{(c_{\mathrm{in}})}\left(\mathrm{Exp}_u(q)\right) \sum_{m\in\mathcal{I}} W_{m,c_{\mathrm{in}},c_{\mathrm{out}}} \prod_{r=1}^{2} B_{m_r}^{(n)}\left(\nu^{(r)}(q)\right) dq. \tag{B.13}$$

For numerical evaluation, let $\{q_\ell\}_{\ell=1}^{L} \subset \Omega$ be a set of tangent-space quadrature points with corresponding weights $\{\omega_\ell\}_{\ell=1}^{L}$. The convolution can then be approximated as

$$y^{(c_{\mathrm{out}})}(u) \approx \sum_{\ell=1}^{L} \omega_\ell \sum_{c_{\mathrm{in}}=1}^{C_{\mathrm{in}}} f^{(c_{\mathrm{in}})}\left(\mathrm{Exp}_u(q_\ell)\right) g_{c_{\mathrm{in}},c_{\mathrm{out}}}(q_\ell). \tag{B.14}$$

Expanding the spline kernel yields

$$y^{(c_{\mathrm{out}})}(u) \approx \sum_{\ell=1}^{L} \omega_\ell \sum_{c_{\mathrm{in}}=1}^{C_{\mathrm{in}}} f^{(c_{\mathrm{in}})}\left(\mathrm{Exp}_u(q_\ell)\right) \sum_{m\in\mathcal{I}} W_{m,c_{\mathrm{in}},c_{\mathrm{out}}} \prod_{r=1}^{2} B_{m_r}^{(n)}\left(\nu^{(r)}(q_\ell)\right). \tag{B.15}$$

In this formulation, the different indices and variables have distinct geometric roles. The point $u \in S^2$ is the center of the convolution, $q \in T_uS^2$ is a tangent-space displacement, $r \in \{1,2\}$ indexes the tangent-coordinate axes, and $m = (m_1, m_2) \in \mathcal{I}$ indexes the tensor-product B-spline basis. The exponential map $\mathrm{Exp}_u(q)$ converts a tangent-space displacement into the corresponding sampling location on the sphere, while $W_{m,c_{\mathrm{in}},c_{\mathrm{out}}}$ contains the learnable coefficients defining the convolutional kernel.

Coordinates normalization Following the formulation of Cox-de Boor B-spline, artificially, the domain of the kernels has a unitary domain. For the radial coordinate $\hat{\rho} = \rho / r_{\max}$. For the angular basis, $t = \frac{k_\theta}{2\pi}\big((\theta + \pi) \bmod 2\pi\big)$ with $n = \lfloor t \rfloor$ so the spline matches edge and keeps it seam continuous, more details. The radial spline is an open spline, while the angular spline is a periodic one.

Example open and close splines Consider the example of open and close spline for number of kernel basis $k_r = 4$, $k_\theta = 8$, and the spline degree is one, which the one used in our implementations. We denote $s = \hat{\rho}$ to make equations more compact, thus the open radial B-spline basis is the four hat functions

$$\big(B_1^{(r)}, B_2^{(r)}, B_3^{(r)}, B_4^{(r)}\big)(\rho) = \begin{cases} (1-3s,\ 3s,\ 0,\ 0), & 0 \le s < \frac{1}{3}, \\ (0,\ 2-3s,\ 3s-1,\ 0), & \frac{1}{3} \le s < \frac{2}{3}, \\ (0,\ 0,\ 3-3s,\ 3s-2), & \frac{2}{3} \le s \le 1. \end{cases} \tag{B.16}$$

For the periodic angular basis uses the close spline formulation and defines: $\delta = t - n$. The degree-one cyclic B-spline implementation has precisely two active angular functions,

$$B^{(\theta)}_{(n-1) \bmod k_\theta}(\theta) = 1 - \delta, \qquad B^{(\theta)}_{n \bmod k_\theta}(\theta) = \delta, \tag{B.17}$$

with every other $B_j^{(\theta)}(\theta)$ equal to zero. In particular, these basis values always sum to one and the modulo indices make the $-\pi/\pi$ seam continuous.

C Estimation of Jacobians for meshes

Consider the $\eta : M \to S^2$ the approximation of the differential d_η for a given a face $\sigma = (a, b, c) \in \mathcal{F}$ is computed from the mesh and sphere chord edges

$$e_1^M = x_b - x_a \tag{C.1}$$
$$e_2^M = x_c - x_a \tag{C.2}$$
$$e_1^S = u_b - u_a \tag{C.3}$$
$$e_2^S = u_c - u_a, \tag{C.4}$$

where $u_i = \eta(x_i)$. We take then the inner product with an orthonormal triangle frame (T_σ^M, B_σ^M) from the mesh face and an analogous frame (T_σ^S, B_σ^S) from the corresponding sphere face. Projecting the two edges into those frames gives

$$D_\sigma = \begin{bmatrix} \langle e_1^M, T_\sigma^M \rangle & \langle e_2^M, T_\sigma^M \rangle \\ \langle e_1^M, B_\sigma^M \rangle & \langle e_2^M, B_\sigma^M \rangle \end{bmatrix}, \tag{C.5}$$

$$G_\sigma = \begin{bmatrix} \langle e_1^S, T_\sigma^S \rangle & \langle e_2^S, T_\sigma^S \rangle \\ \langle e_1^S, B_\sigma^S \rangle & \langle e_2^S, B_\sigma^S \rangle \end{bmatrix}. \tag{C.6}$$

The jacobian and the inverse jacobians are compute as

$$J_\sigma = G_\sigma D_\sigma^+ \tag{C.7}$$
$$A_\sigma = J_\sigma^+ \tag{C.8}$$

Where $\cdot^+$ is the Moore-Penrose pseudo inverse. The pseudo inverse make the calculation robust to small triangles or degenerated triangles. Each vertex query is assigned to deterministic faces index, so when computing the transport we obtain a precomputed matrix every forward-pass.

D Polar support quadrature

In the convolution the quadrature is constructed from the tangent vector on $T_u S^2$ with $u \in \eta(x)$ for a point $x \in M$ and the map $\eta : M \to S^2$. As mentioned in Section 3.1, there is a set of tangent points of the exponential map of the set of tangent vector of the support, the subdivisions on the form by the grid on the manifold S^2 form the quadrature. We compute consider the radial ring r_l for $l \in [1, L]$ with K_l angular partitions, the area elements as the angular index $m \in [1, K_l]$ is given by

$$\hat{\omega}_{ml} = sin(r_l)\Delta r \Delta\theta \tag{D.1}$$

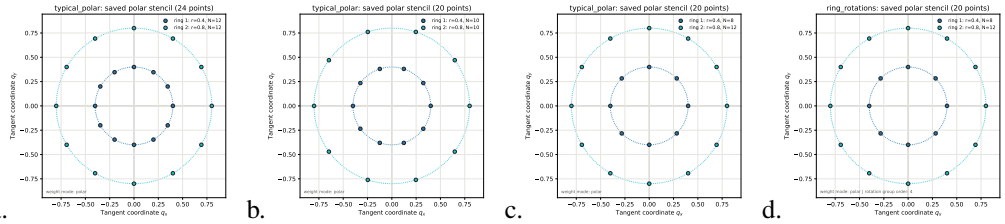

a. b. c. d.

Figure 2: Kernel Construction variant. a. Baseline kernel with 24 points. b. regular polar stencil with 20 points. c. regular polar with 20 points d. ring polar where first ring has C_8 symmetry, and second ring has C_{12} symmetry with a total of C_4, where $4 = \gcd(8, 12)$.

where $\Delta r = r_{l+1} - r_l$ and the $\Delta\theta = 2\pi/K_l$ assuming a equidistant partition of the angular axis. The term $sin(r_l)$ is takes for the area-scaling factor of the geodesic of the polar coordinate on the unit sphere around a point $u \in S^2$, computed as the determinant of the metric tensor.

The value used in the equation equation 4.2 is the normalization so the area in the domain is has a unitary value, thus we write

$$\omega_{lm} = \frac{\hat{\omega}_{ml}}{\sum_{l'm'} \hat{\omega}_{m'l'}} \tag{D.2}$$

E Anisotropic signal generation

We choose a center c on the deformed mesh M. For the calculation the signal, project the points to a canonical sphere and then use a spherical logarithmic map to compute the correspondence onto the tangent plane

$$\hat{c} = \frac{c}{\|c\|}, \qquad \hat{x} = \frac{x}{\|x\|}, \qquad z(x) = \mathrm{Log}_{\hat{c}}^{S^2}(\hat{x}) \in T_{\hat{c}}S^2. \tag{E.1}$$

The center is normalized at the equivalent z(c)=0 direction. In addition, two vertices sharing the same ray from the origin have the same local direction and the same sphere logarithmic map $\mathrm{Log}_{\hat{c}}^{S^2}$. A Hughes-Möller construction gives an orthonormal tangent basis on this sphere $(e_1, e_2) \in T_cS^2$. Then we randomly sample an angle $\delta \in [0, \pi)$ and define the anisotropic Gaussian directions as v with perpendicular vector $v_\perp$ as:

$$v_\parallel = R(\delta)e_1 = \cos(\delta)e_1 + \sin(\delta)e_2 \qquad v_\perp = v \times \hat{c} \tag{E.2}$$

The scalar value assignment to the vertex x is with $\sigma_\parallel \gg \sigma_\perp$ gives the signal its anisotropy:

$$f(x) = A \exp\left(-\frac{1}{2}\left(\frac{\langle z(x), v\rangle^2}{\sigma_\parallel^2} + \frac{\langle z(x), v_\perp\rangle^2}{\sigma_\perp^2}\right)\right). \tag{E.3}$$

F Ablations experiments

The ablation experiments examine the architectural design of pmConv with orientation banks. They are organized around five questions, each varying one key hyperparameter while the others are held at the baseline configuration ($K = 12$, radii $0.4, 0.8$, $N = 12$ points per ring). We have B^+ as the baseline after one run of hyperparameter optimization only for kernel values.

1. **Q1: kernel support scale.** We scale the radial extent of the kernel in three steps while keeping $\rho_2 = 2\rho_1$: $(\rho_1, \rho_2) \in \{(0.4, 0.8), (0.8, 1.6), (1.6, 3.2)\}$. On the unit sphere, the geodesic angle reached by a tangent vector v equals its length. From Eq. equation 3.7, $\cos\theta = \langle p, \mathrm{Exp}_p^{S^2}(v)\rangle = \cos|v|$, so $\theta = |v|$ for $|v| \le \pi$ and $\theta = 2\pi - |v|$ for $\pi < |v| \le 2\pi$. The largest radius, 3.2, therefore corresponds to about 183°: the outer ring passes the antipode of p, where the exponential map is no longer injective and the kernel loses its locality.
2. **Q2: angular resolution.** We keep two rings with radii 0.4 and 0.8 and vary the number of sample points per ring, $N \in \{4, 8, 10, 12, 32\}$. A finer angular resolution should allow a finer resolution of rotations, and hence better equivariance, which is what we test.

Table 3: Orientation-bank ablation for aniso_orientation (ellipsoid with small ARAP deformations): mean ± standard deviation across seeds. Bold marks the best value within each question block, including the baseline.

						e		$SO(3)$		C_{12}	
ID	Construction	K	Radii	N per ring	Seeds	MAE	$1-\cos$	MAE	$1-\cos$	MAE	$1-\cos$
Q1: support scale											
B	regular polar	12	0.4,0.8	12,12	1/3	0.188	0.154	0.303	0.298	0.288	0.280
B	regular polar	8	0.4,0.8	16,16	3	$\mathbf{0.137_{\pm 0.020}}$	$\mathbf{0.099_{\pm 0.028}}$	$\mathbf{0.193_{\pm 0.042}}$	$\mathbf{0.154_{\pm 0.047}}$	$\mathbf{0.184_{\pm 0.039}}$	$\mathbf{0.145_{\pm 0.047}}$
Q1a	regular polar	12	0.8,1.6	12,12	3	$0.188_{\pm 0.017}$	$0.163_{\pm 0.021}$	$0.366_{\pm 0.017}$	$0.377_{\pm 0.017}$	$0.352_{\pm 0.027}$	$0.369_{\pm 0.032}$
Q1b	regular polar	12	1.6,3.2	12,12	3	$0.146_{\pm 0.029}$	$0.110_{\pm 0.039}$	$0.352_{\pm 0.036}$	$0.364_{\pm 0.054}$	$0.316_{\pm 0.029}$	$0.321_{\pm 0.042}$
Q2: angular resolution											
B	regular polar	12	0.4,0.8	12,12	1/3	0.188	0.154	0.303	0.298	0.288	0.280
B	regular polar	8	0.4,0.8	16,16	3	$0.137_{\pm 0.020}$	$0.099_{\pm 0.028}$	$\mathbf{0.193_{\pm 0.042}}$	$\mathbf{0.154_{\pm 0.047}}$	$\mathbf{0.184_{\pm 0.039}}$	$\mathbf{0.145_{\pm 0.047}}$
Q2a	regular polar	12	0.4,0.8	4,4	2/3	$\mathbf{0.073_{\pm 0.005}}$	$\mathbf{0.034_{\pm 0.004}}$	$0.553_{\pm 0.005}$	$0.634_{\pm 0.010}$	$0.473_{\pm 0.009}$	$0.547_{\pm 0.014}$
Q2b	regular polar	12	0.4,0.8	8,8	2/3	$0.172_{\pm 0.023}$	$0.139_{\pm 0.032}$	$0.405_{\pm 0.038}$	$0.438_{\pm 0.066}$	$0.357_{\pm 0.029}$	$0.371_{\pm 0.049}$
Q2c	regular polar	12	0.4,0.8	10,10	2/3	$0.176_{\pm 0.011}$	$0.135_{\pm 0.008}$	$0.288_{\pm 0.084}$	$0.276_{\pm 0.098}$	$0.289_{\pm 0.067}$	$0.280_{\pm 0.078}$
Q2d	regular polar	12	0.4,0.8	32,32	3	$0.142_{\pm 0.029}$	$0.098_{\pm 0.034}$	$0.246_{\pm 0.097}$	$0.216_{\pm 0.113}$	$0.222_{\pm 0.081}$	$0.187_{\pm 0.095}$
Q3: orientation-bank order											
B	regular polar	12	0.4,0.8	12,12	1/3	0.188	0.154	0.303	0.298	0.288	0.280
B	regular polar	8	0.4,0.8	16,16	3	$0.137_{\pm 0.020}$	$0.099_{\pm 0.028}$	$0.193_{\pm 0.042}$	$0.154_{\pm 0.047}$	$0.184_{\pm 0.039}$	$0.145_{\pm 0.047}$
Q3a	regular polar	1	0.4,0.8	12,12	3	$0.133_{\pm 0.015}$	$0.089_{\pm 0.019}$	$\mathbf{0.159_{\pm 0.031}}$	$\mathbf{0.117_{\pm 0.039}}$	$\mathbf{0.151_{\pm 0.024}}$	$\mathbf{0.111_{\pm 0.031}}$
Q3b	regular polar	4	0.4,0.8	12,12	3	$\mathbf{0.130_{\pm 0.051}}$	$\mathbf{0.089_{\pm 0.060}}$	$0.188_{\pm 0.072}$	$0.151_{\pm 0.086}$	$0.183_{\pm 0.068}$	$0.146_{\pm 0.081}$
Q3c	regular polar	10	0.4,0.8	12,12	3	$0.159_{\pm 0.041}$	$0.115_{\pm 0.052}$	$0.245_{\pm 0.090}$	$0.217_{\pm 0.109}$	$0.255_{\pm 0.095}$	$0.230_{\pm 0.124}$
Q3d	regular polar	32	0.4,0.8	12,12	3	$0.157_{\pm 0.029}$	$0.112_{\pm 0.031}$	$0.307_{\pm 0.121}$	$0.296_{\pm 0.167}$	$0.279_{\pm 0.103}$	$0.263_{\pm 0.135}$
Q4: radial resolution											
B	regular polar	12	0.4,0.8	12,12	1/3	0.188	0.154	0.303	0.298	0.288	0.280
B	regular polar	8	0.4,0.8	16,16	3	$\mathbf{0.137_{\pm 0.020}}$	$\mathbf{0.099_{\pm 0.028}}$	$\mathbf{0.193_{\pm 0.042}}$	$\mathbf{0.154_{\pm 0.047}}$	$\mathbf{0.184_{\pm 0.039}}$	$\mathbf{0.145_{\pm 0.047}}$
Q4a	regular polar	12	0.08,0.8	12,12	3	$0.222_{\pm 0.027}$	$0.191_{\pm 0.032}$	$0.341_{\pm 0.019}$	$0.345_{\pm 0.026}$	$0.331_{\pm 0.014}$	$0.333_{\pm 0.021}$
Q4b	regular polar	12	0.08,0.44,0.8	12,12,12	3	$0.147_{\pm 0.021}$	$0.107_{\pm 0.028}$	$0.244_{\pm 0.106}$	$0.218_{\pm 0.131}$	$0.238_{\pm 0.096}$	$0.213_{\pm 0.123}$
Q4c	regular polar	12	0.08,0.32,0.56,0.8	12,12,12,12	3	$0.150_{\pm 0.021}$	$0.104_{\pm 0.031}$	$0.334_{\pm 0.096}$	$0.330_{\pm 0.142}$	$0.302_{\pm 0.085}$	$0.291_{\pm 0.122}$
Q4d	regular polar	12	0.08,0.26,0.44,0.62,0.8	12,12,12,12,12	3	$0.161_{\pm 0.009}$	$0.114_{\pm 0.018}$	$0.285_{\pm 0.105}$	$0.269_{\pm 0.134}$	$0.258_{\pm 0.078}$	$0.234_{\pm 0.103}$
Q4e	regular polar	12	0.08,0.16,0.24,0.32,0.4,0.48,0.56,0.64,0.72,0.8	12,12,12,12,12,12,12,12,12,12	3	$0.159_{\pm 0.029}$	$0.114_{\pm 0.034}$	$0.307_{\pm 0.101}$	$0.295_{\pm 0.143}$	$0.289_{\pm 0.090}$	$0.276_{\pm 0.119}$
Q5: kernel construction											
B	regular polar	12	0.4,0.8	12,12	1/3	0.188	0.154	0.303	0.298	0.288	0.280
B	regular polar	8	0.4,0.8	16,16	3	$\mathbf{0.137_{\pm 0.020}}$	$\mathbf{0.099_{\pm 0.028}}$	$\mathbf{0.193_{\pm 0.042}}$	$\mathbf{0.154_{\pm 0.047}}$	$\mathbf{0.184_{\pm 0.039}}$	$\mathbf{0.145_{\pm 0.047}}$
Q5a	regular polar	4	0.4,0.8	10,10	2/3	$0.192_{\pm 0.048}$	$0.162_{\pm 0.052}$	$0.258_{\pm 0.055}$	$0.237_{\pm 0.056}$	$0.249_{\pm 0.054}$	$0.230_{\pm 0.064}$
Q5b	regular polar	4	0.4,0.8	8,12	2/3	$0.167_{\pm 0.016}$	$0.127_{\pm 0.021}$	$0.255_{\pm 0.036}$	$0.233_{\pm 0.043}$	$0.240_{\pm 0.032}$	$0.215_{\pm 0.033}$
Q5c	ring polar	4	0.4,0.8	8,12	2/3	$0.156_{\pm 0.007}$	$0.118_{\pm 0.002}$	$0.354_{\pm 0.081}$	$0.365_{\pm 0.112}$	$0.325_{\pm 0.072}$	$0.322_{\pm 0.103}$
Q5d	Cartesian	4	0.4,0.8	10,10	2/3	$0.172_{\pm 0.011}$	$0.140_{\pm 0.016}$	$0.323_{\pm 0.047}$	$0.323_{\pm 0.064}$	$0.298_{\pm 0.012}$	$0.290_{\pm 0.014}$

3. **Q3: orientation-bank order.** We vary the number of orientations in the bank, $K \in \{1, 4, 10, 12, 32\}$. The bank order should also affect equivariance: when K matches, or divides, the angular resolution of the rings, the rotated kernels fall on the stencil points and the bank should be more consistent.
4. **Q4: radial resolution.** We keep the outer radius at 0.8 and increase the number of rings from 2 to 10, evenly spaced between 0.08 and 0.8, with $N = 12$ points per ring.
5. **Q5: stencil construction.** We compare how the tangent vectors of the kernel are arranged (Figure 2). The ring-polar construction staggers points on neighbouring rings, so that the tangent vectors cover the disc more uniformly. We compare it against regular polar grids with the same (10+10) or a different (8+12) angular resolution per ring, without the staggering rule, all with $K = 4$. We also compare with the stencil with cartesian coordinates. We aim to get 20 points in all the cases.

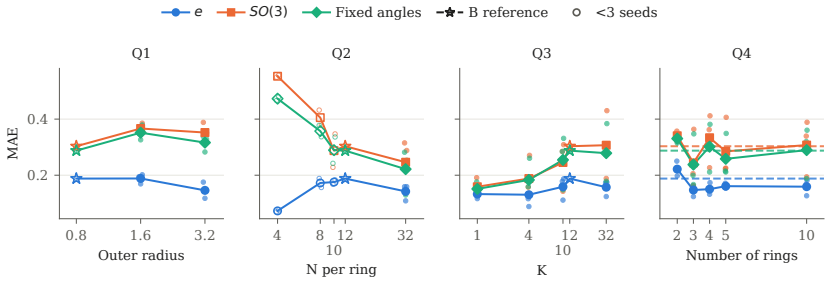

(a) MAE

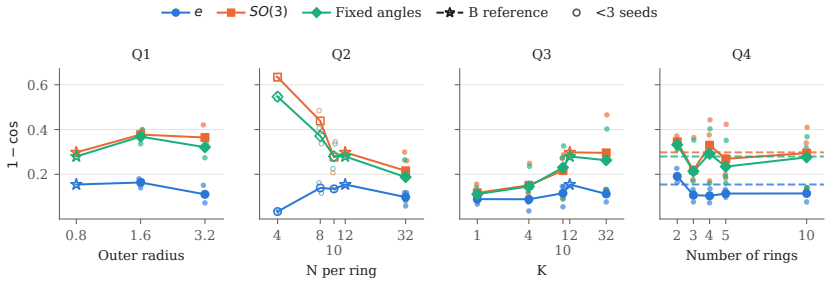

(b) $1 - \cos$

Figure 3: Sensitivity of pmConv to its design parameters (Q1–Q4) on regression anisotropic orientation, surface is a ellipsoid with small ARAP deformations. Each panel varies one parameter from the baseline B ($K = 12$, radii $0.4, 0.8$, $N = 12$), marked with a star. Lines show the mean over seeds and bars ± 1 std, under three test conditions: identity (e), random $SO(3)$ rotations and C_{12} rotations. In Q4, B is not on the sweep and is shown as dashed lines.

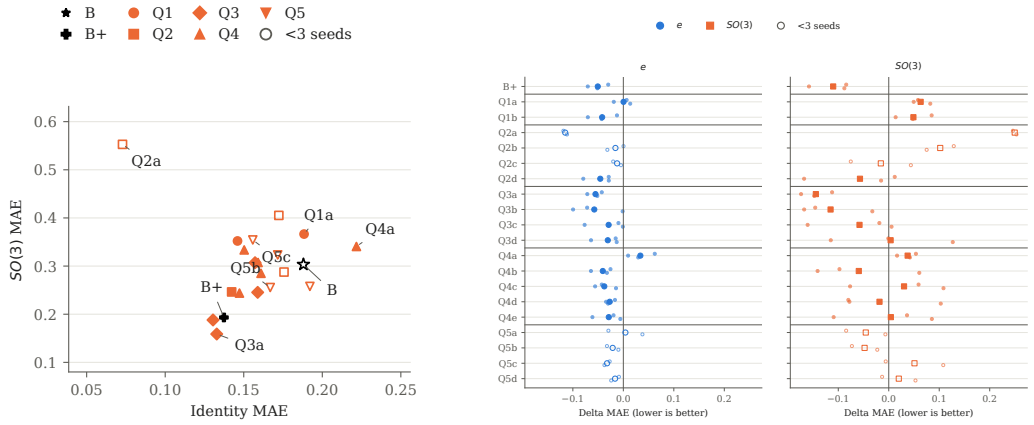

(a) Accuracy vs. rotation robustness

(b) Change in MAE relative to B

Figure 4: (a) MAE without rotation (e) against MAE under random $SO(3)$ rotations, one point per configuration; marker shape indicates the ablation question. (b) Difference in MAE from the baseline B for each ablation, grouped by question; lower is better. Hollow markers: fewer than three seeds.