
Unified Fourier-based Kernel and Nonlinearity Design for Equivariant Networks on Homogeneous Spaces

Yinshuang Xu^{*1} Jiahui Lei^{*1} Edgar Dobriban^{1,2} Kostas Daniilidis¹

Abstract

We introduce a unified framework for group equivariant networks on homogeneous spaces derived from a Fourier perspective. We consider tensor-valued feature fields, before and after a convolutional layer. We present a unified derivation of kernels via the Fourier domain by leveraging the sparsity of Fourier coefficients of the lifted feature fields. The sparsity emerges when the stabilizer subgroup of the homogeneous space is a compact Lie group. We further introduce a nonlinear activation, via an elementwise nonlinearity on the regular representation after lifting and projecting back to the field through an equivariant convolution. We show that other methods treating features as the Fourier coefficients in the stabilizer subgroup are special cases of our activation. Experiments on $SO(3)$ and $SE(3)$ show state-of-the-art performance in spherical vector field regression, point cloud classification, and molecular completion.

1 Introduction

Following the success of convolutional neural networks (CNNs) (Fukushima, 1980; LeCun et al., 1989), researchers made great strides in designing equivariant networks for groups beyond the standard translation operation. Equivariance preserves symmetries and drastically reduces sample complexity, making data augmentation unnecessary and consequently reducing training and testing time.

^{*}Equal contribution ¹Department of Computer and Information Science, University of Pennsylvania, Philadelphia, PA 19104, United States ²Department of Statistics and Data Science, University of Pennsylvania, Philadelphia, PA 19104, United States. Correspondence to: Yinshuang Xu <xuyin@seas.upenn.edu>, Jiahui Lei <leijh@seas.upenn.edu>, Edgar Dobriban <dobriban@wharton.upenn.edu>, Kostas Daniilidis <kostas@cis.upenn.edu>.

The standard approach for equivariance by design is group-convolutional networks (G-CNNs), specified by convolution kernels (Cohen & Welling, 2016). Kondor & Trivedi (2018) prove that such group convolutions are both sufficient and necessary for linear layers to be equivariant to the action of compact groups. Existing group convolutional nets construct equivariant kernels in a case by case fashion; for example, Cohen et al. (2018a) present a general constraint for the kernel of G-CNNs with features on a homogeneous space that is solved algebraically, while Finzi et al. (2021) solve for kernels numerically based on a similar finite-dimensional constraint.

We propose a unified recipe for kernel design. We lift the feature fields from the initial homogeneous space to the corresponding Mackey functions on the acting group, as introduced in (Cohen et al., 2018a). Since a Mackey function satisfies a well-known constraint, it has redundant information. Specifically, when the irreducible representation of the stabilizer subgroup is trivial, Kondor & Trivedi (2018) proves that the Fourier Transform of the Mackey function has a certain block sparsity pattern. This naturally leads to the question whether the Fourier transform of Mackey functions for nontrivial irreducible representations over more general groups may have a similar property. We prove that when the stabilizer subgroup is a compact Lie group, the Fourier coefficients are sparse and nonzero for specific field types, as stated in proposition 3.2. This spectral sparsity, appearing in both the input and the output of the convolution, implies that the kernel itself can be taken as sparse. This enables us to characterize the space of kernels without using the classical equivariance constraints and lays the foundation of our framework for linear layer design.

Regarding the nonlinearity, most existing group convolution methods on homogeneous spaces (Cohen et al., 2018a) apply norms or gated nonlinearities. The exceptions are works applying the Clebsch-Gordan decomposition of tensor products (Kondor et al., 2018) or interpreting the features as Fourier coefficients (De Haan et al., 2020; Poulenard & Guibas, 2021), since the feature vectors are in the vector space under the irreducible representation of the stabilizer subgroup. We propose a general formulation for nonlinear layers that consists of lifting the fields to the Mackey func-

tions on the group, applying an elementwise nonlinearity, and finally projecting back to a field on the homogeneous space through a convolution. We prove the equivariance of this activation, and that the recently proposed equivariant nonlinearity in Poulernard & Guibas (2021) is a special case.

In summary, our contributions are:

1. We provide a unifying Fourier view of group convolution on homogeneous spaces dealing with different field types, and we prove that the input, output, and the kernel of the convolution are block-sparse in the Fourier domain when the stabilizer subgroup is a compact Lie group. Given the irreducible representations of the group, this harmonic view leads to an efficient method for designing linear group convolution layers.
2. We present a novel approach to equivariant nonlinearities by applying elementwise nonlinearities to the feature fields lifted to Mackey functions on the group. We propose and implement a general approach for projecting functions on the group to fields on the homogeneous space in the final step of the nonlinearity.

The main goal and novelty of the paper is to use the Fourier coefficients for a unified derivation of kernels and nonlinearities. We reach state-of-the-art results in standard equivariance benchmarks in 3D shape classification, molecular completion, and spherical vector field regression.

2 Related Work

Equivariant Networks The most common method to design equivariant networks is via group convolution, on the group or on the homogeneous space where the features lie.

Existing work constructs group convolutional networks on images (Cohen & Welling, 2016; Worrall et al., 2017), point clouds (Thomas et al., 2018; Chen et al., 2021), voxel grid (Weiler et al., 2018; Worrall & Brostow, 2018), graphs (Maron et al., 2018; Keriven & Peyré, 2019), spherical images (Cohen et al., 2018b; Esteves et al., 2018; 2020; Cobb et al., 2020) and sets (Maron et al., 2020; Esteves et al., 2019). The above works that perform the group convolution directly on the homogeneous space implicitly lift the function to the group before the convolution, and project back to the homogeneous space after the convolution.

Another way to achieve equivariance is through averaging over the group orbits (Puny et al., 2021; Atzmon et al., 2021). Finzi et al. (2020) proposed a general method for any Lie group with a surjective exponential map. For regular groups, Bekkers (2019); Sosnovik et al. (2019) propose a direct construction for group convolution. Recently, equivariance was introduced for attention networks (Fuchs et al., 2020; Romero & Cordonnier, 2020; Hutchinson et al., 2021; Satorras et al., 2021). Cohen et al. (2018a) and Kondor &

Trivedi (2018) showed for homogeneous spaces and compact groups, respectively, that an equivariant map can always be written as a convolution. Finzi et al. (2021) introduced a numerical algorithm for equivariant linear layers based on solving linear systems involving the generators of the Lie algebra. We provide a general spectral approach for the theoretical analysis, and efficient implementations of those works.

Equivariance and Fourier Transform Several works use the relationship between group convolutional networks and the Fourier Transform. Kondor & Trivedi (2018) provide a Fourier view of the group convolution for a scalar field on the quotient space, or on the whole group. Cohen et al. (2018b) and Esteves et al. (2018) apply group convolutions on the Fourier transform of $SO(3)$ and the spherical harmonics, respectively, while Kondor et al. (2018) directly convolve and apply a non-linear activation in the spectral domain through the Clebsch-Gordan decomposition, using the compactness of $SO(3)$. Finally, Esteves et al. (2020) deals with both vector and scalar signals on the sphere through the spectral domain, and is a special case of our general framework.

Equivariant Nonlinearity It is nontrivial to design an expressive and equivariant nonlinear layer, since an equivariant activation has to commute with the group action. Several works (Cohen & Welling, 2016; Cohen et al., 2018b; Worrall & Brostow, 2018) lift the signals from the homogeneous space to the group, and apply an elementwise activation to the group function. However, these methods deal with scalar fields and obtain the invariant features through global pooling without projecting back to the homogeneous space. In (Thomas et al., 2018; Weiler et al., 2018; Worrall et al., 2017; Esteves et al., 2020) a nonlinearity is applied over invariant features such as norms. Because the direction of the tensor field remains unchanged, such a nonlinearity is not expressive enough, as explained in (Poulernard & Guibas, 2021). Kondor et al. (2018) and Anderson et al. (2019) apply a polynomial activation that can result in training instability, as mentioned in (Anderson et al., 2019). Weiler & Cesa (2019) introduce nonlinearities with respect to various representations of $E(2)$, while Deng et al. (2021) generalize the classical ReLU activation to vectors (i.e., order-one tensors) by truncating an equivariant projection of the vector—a combination of a gated activation and a linear layer. De Haan et al. (2020) and Poulernard & Guibas (2021) treat the features as Fourier coefficients, applying the Inverse Fourier Transform to functions on the stabilizer subgroup, and projecting back to the features by the Fourier Transform. We prove that such a nonlinearity is a special case of our method.

The works most closely related to ours are (Cohen et al., 2018a) and (Kondor & Trivedi, 2018). The main differ-

ence is that we provide a Fourier perspective for the group convolution on the homogeneous space, enabling efficient kernel design. Moreover, we provide a new nonlinearity and—to our knowledge, for the first time—implement a general equivariant linear map from the regular representation to the induced representation. [Kondor & Trivedi \(2018\)](#) analyze G-CNNs through Fourier analysis, but only focus on scalar fields. In contrast, we consider all field types (scalars, vectors, tensors) and our conclusion is consistent with [\(Kondor & Trivedi, 2018\)](#) for the special case of scalar fields.

3 Method

We give in App. A.1 the basic definitions that we need for a group acting on a homogeneous space, as well as the definitions of cosets, bundles, fibers, and twist functions. In App. A.2 we describe irreducible representations, and leave here only the tools we need for modeling features as tensor fields (Sec. 3.1) and the Fourier Transform (Sec. 3.2). Our work proposes two components for equivariant neural networks: the linear group convolution layer (Sec. 3.3) and the nonlinear activation layer (Sec. 3.4).

3.1 Induced representation and Mackey Functions

For a feature (or field) $f : G/H \rightarrow V$ over the homogeneous space G/H taking values in the vector space V , its type (scalar, vector or higher order tensor) is determined by an irreducible representation of the stabilizer subgroup H . We use ρ to denote the unitary irreducible representation of H . Then $g \in G$ acts on the field as

$$(\mathcal{L}_g f)(x) = \rho(\mathfrak{h}(g^{-1}, x)^{-1})f(g^{-1}x) \quad (1)$$

where \mathfrak{h} is the twist function introduced in App. A.1. For example, for $G = SE(2)$, where $H = SO(2)$, the stabilizer subgroup has irreducible representations $\rho^m : SO(2) \mapsto \mathbb{C}$ for $m \in \mathbb{Z}$, given by $\rho^m(\theta) = e^{im\theta}$, for $\theta \in SO(2)$. For scalar fields, we have $\rho(\theta) = \rho^0(\theta) = 1$; for vector fields, we have $\rho(\theta) = \rho^1(\theta) = e^{i\theta}$; and for a physical quantity like momentum of inertia over \mathbb{R}^2 , we have $\rho(\theta) = \rho^2(\theta) = e^{2i\theta}$. In general, equation (1) describes the action of G on fields and \mathcal{L} is called the induced representation, denoted as $\mathcal{L} = \text{Ind}_H^G \rho$.

The field $f : G/H \rightarrow V$ can be lifted to a function $f^\uparrow^G : G \rightarrow V$ on the group G through an isomorphism Λ , via

$$f^\uparrow^G(g) = (\Lambda f)(g) = \rho(\mathfrak{h}(g)^{-1})f(gH)$$

and projected back via Λ^{-1}

$$f(x) = f(s(x)H) = (\Lambda^{-1}(f^\uparrow^G))(s(x)H) = f^\uparrow^G(s(x)),$$

where \mathfrak{h} is the twist function and s is the section map (App. A.1). For the lifting in [Kondor & Trivedi \(2018\)](#),

ρ is a trivial representation, since only the scalar field is considered. The action $\mathcal{L}' = \rho_{\text{reg}}^G$ of G on f^\uparrow^G is a regular representation $(\mathcal{L}'_g(f^\uparrow^G))(k) = f^\uparrow^G(g^{-1}k)$ ([Folland \(2016\)](#), Ch. 3).

The lifting operation commutes with the group action, so $\rho_{\text{reg}}^G \circ \Lambda = \Lambda \circ \text{Ind}_H^G \rho$, which justifies calling the lifting an isomorphism. The function $f^\uparrow^G : G \rightarrow V$ lifted from the field $f : G/H \rightarrow V$ is called a Mackey function, satisfying $f^\uparrow^G(gh) = \rho(\mathfrak{h}^{-1})f^\uparrow^G(g)$.

3.2 Fourier Transform

It is known that the Fourier Transform exists for most groups of interest for equivariance: finite groups, compact groups, separable unimodular locally compact groups of Type I (see [Gross \(1978\)](#), [Folland \(2016\)](#), Ch. 7, and App. A.3 for definitions), and certain semidirect product groups with an Abelian normal subgroup, like $SE(n)$ ([Gauthier et al., 1991](#)). In these cases, there is a uniform (Haar) measure on the group G that is both a left and a right Haar measure; we will denote its differential by dg . Let $U(\cdot, p)$ be the unitary irreducible representation for each element $p \in \hat{G}$, where \hat{G} is the dual of the group G , the set of equivalence classes of unitary irreducible representations of G ([Folland \(2016\)](#), Ch. 7, [Chirikjian et al. \(2001\)](#), Ch. 8). In the cases of interest, there is also a Haar measure ν on \hat{G} . For any function $f : G \rightarrow \mathbb{C}$ that is square integrable with respect to the measure, the Fourier Transform of f is defined as

$$\hat{f}(p) = \mathcal{F}(f)(p) = \int_G f(g)\overline{U(g, p)}dg$$

while the inverse Fourier Transform reads

$$f(g) = \int_{\hat{G}} \text{tr}(\hat{f}(p)^T U(g, p))d\nu(p)$$

and the convolution and Parseval theorems both hold for $f \in \mathcal{L}^2(G)$. For vector valued functions, we apply both transforms component-wise.

3.3 Unified Kernel Derivation

As described in [Cohen et al. \(2018a\)](#), an equivariant linear layer, consisting of a convolution with a kernel, can be realized in three steps: 1) lifting the input field (feature map) $f_{in} : G/H \rightarrow V$ to the corresponding Mackey function $f_{in}^\uparrow^G : G \rightarrow V$ on the group; 2) applying group convolution to the lifted function $f_{in}^\uparrow^G$ with the kernel κ to find the output Mackey function $f_{out}^\uparrow^G$; 3) projecting $f_{out}^\uparrow^G$ back to the homogeneous space to obtain the output field (feature map) f_{out} . These three steps can be summarized as the following linear map yielding

$$\begin{aligned} f_{out}(x) &= (\Lambda_2^{-1}(\kappa * (\Lambda_1 f_{in}))) (x) \\ &= \int_G \kappa(g^{-1}s(x))(\Lambda_1 f_{in})(g)dg, \end{aligned}$$




Figure 1. **Convolution layer:** (a). Lifting the input field $f_{in} : G/H \rightarrow V$ to the Mackey function $f_{in}^{\uparrow G} : G \rightarrow V$. (b). Convolution with designed kernel to find a new Mackey function $f_{out}^{\uparrow G} : G \rightarrow V$ of order l_{out} . (c). Projection to the homogeneous space to find f_{out} . We highlight the sparsity in the spectral domain of the input, kernel and output.

where Λ_1 and Λ_2 are the lifting isomorphisms for the input and output fields, respectively, and $\kappa : G \rightarrow \text{Hom}(V_1, V_2)$. When G is a semidirect product, convolution simplifies to

$$(\Lambda_2^{-1}(\kappa * (\Lambda_1 f)))(x) = \int_{G/H} \kappa'(s(x)^{-1}y) f(y) dy,$$

where $\kappa'(x) = \kappa(s(x))$. Previous work (Cohen et al., 2018a) derives the constraints for the kernel κ when the input and output are Mackey functions by working in the spatial domain. In contrast, we analyze and design the linear layer from the spectral perspective. In section 3.3.1, we will prove that the Fourier coefficients of a Mackey function obey certain sparsity patterns, which also imply the sparsity of the kernel κ . Moreover, in certain cases our proposition leads to a more efficient implementation of convolutions in the spectral domain (compared to the spatial domain), as shown in Figure 1.

3.3.1 Sparsity of Mackey functions in the Fourier domain

In this section we prove that the Fourier transform of a Mackey function $f^{\uparrow G}$ lifted from the field f has a certain sparsity pattern. To prove this proposition, we first need to introduce a lemma about unitary irreducible representations of G .

Lemma 3.1. *Let G be a unimodular group, and its stabilizer subgroup H be a compact Lie group. For any $p \in \hat{G}$, the dual group, the restricted representation for the unitary irreducible representation $U(\cdot, p)$, $U(\cdot, p)|_H$, can be decomposed as the direct sum $\bigoplus_{i \in \mathcal{Q}(p)} \rho^i(h(g))$, for an index set $\mathcal{Q}(p)$ parametrized by $p \in \hat{G}$, where ρ^i are the irreducible representations of H .*

See App. B.1 for the proof. Using Lemma 3.1 we can now describe the sparsity of Mackey functions in the Fourier domain.

Proposition 3.2. *Assume G is a unimodular group and its stabilizer subgroup is a compact Lie group. A Mackey function $f^{\uparrow G} : G \rightarrow V$ lifted via $f^{\uparrow G}(g) = \rho(h(g)^{-1})f(gH)$ from a field $f : G/H \rightarrow V$ has the following sparsity pattern in the Fourier domain: $[\widehat{f^{\uparrow G}}](p)_{*,j}$ is nonzero only if*

the block at column j in the decomposition of $U(\cdot, p)|_H$ is equivalent to the dual representation of ρ .

See App. B.2 for the proof. Next, we show that this sparsity carries on to the convolution, ensuring that the spectrum of the kernel is also sparse.

Corollary 3.3. *Let $f_1^{\uparrow G}$ and $f_2^{\uparrow G}$ be Mackey functions lifted from the fields f_1 and f_2 , and ρ_1 and ρ_2 be the irreducible representations that determine the field type of f_1 and f_2 . For any group convolution $f_2^{\uparrow G}(g) = \int_G \kappa(\nu^{-1}g) f_1^{\uparrow G}(\nu) d\nu$, without loss of generality, the kernel κ has the following sparsity pattern on its Fourier coefficients: $[\widehat{\kappa}](p)_{i,j}$ is zero when the block at row i in the decomposition of $U(\cdot, p)|_H$ is not equivalent to the dual representation of ρ_1 , or the block at column j in the decomposition of $U(\cdot, p)|_H$ is not equivalent to the dual representation of ρ_2 .*

See App. B.4 for the proof. In particular, when ρ is the trivial representation, i.e., the field is scalar-valued, and G is compact, proposition 3.2 recovers the conclusion of (Kondor & Trivedi, 2018). In summary, we proved that the spectrum of the Mackey function and the corresponding kernel are sparse. This sparsity enables us to directly and analytically design kernels and implement the group convolution.

Meanwhile, we state that our design gives a complete characterization of the space of kernels for equivariant convolutions. Given an appropriate unitary irreducible representation, the block sparsity stated in Prop. 3.2 can be simplified to column sparsity. The Fourier coefficients of elements of vector functions are related. The converse also holds, hence the Fourier coefficients of a function f have such sparsity and related values if and only if f is a Mackey function. Further, with the convolution theorem, we find the sparsity in the spectrum of the kernel κ , and the relation of Fourier coefficients for every element in the matrix function. Finally, we prove that there is a bijection between the kernel space $\{\kappa : G \rightarrow \text{Hom}(V_1, V_2)\}$, where $\kappa(gh) = \rho_{out}(h^{-1})\kappa(g)$ and $\kappa(hg) = \kappa(g)\rho(h^{-1})$ for any $g \in G, h \in H$, and the kernel space with such a spectrum. See App. B.5 for details. To help the reader, now explain the sparsity pattern for $SO(3)$ and $SE(2)$.

Sparsity in the $SO(3)$ -spectrum Consider the group $G = SO(3)$, with the stabilizer subgroup $SO(2)$ and the homogeneous space \mathbb{S}^2 . A rotation matrix $R \in SO(3)$ can be parametrized by the Euler Z-Y-Z angles (α, β, γ) . The unitary irreducible representations of $SO(3)$ are indexed by integers l and have the form

$$D_{mn}^l(R) = D_{mn}^l(\alpha, \beta, \gamma) = e^{-im\alpha} d_{mn}^l(\beta) e^{-in\gamma},$$

where m, n are row and column indices, and d_{mn}^l are the elements of Wigner's small d -matrices. The lifting process and its inverse take the form

$$\begin{aligned} f \uparrow^G(R) &= f \uparrow^G(\alpha, \beta, \gamma) = e^{-ik\gamma} f(\alpha, \beta), \\ f(\alpha, \beta) &= f \uparrow^G(\alpha, \beta, 0), \end{aligned}$$

where k is the corresponding order of the irreducible representation of the field. The Fourier Transform $\widehat{f \uparrow^G}_{mn}$ (see App. C.2 for details) becomes:

$$\int_{(\alpha, \beta) \in \mathbb{S}^2} f(\alpha, \beta) e^{im\alpha} d_{mn}^l(\beta) d\alpha \sin(\beta) d\beta * 2\pi \delta(k - n),$$

where $\delta(x)$ is the Kronecker delta function, which is zero except at $x = 0$, where it equals unity. The convolution on $SO(3)$ takes the form $(l_1 * l_2)(g) = \int_{k \in G} l_1(\nu^{-1}g) l_2(\nu) d\nu$ and the convolution theorem states $\mathcal{F}(\kappa * f)_{mn}^l = \sum_j \hat{f}_{mj}^l \hat{\kappa}_{jn}^l$. When the output field corresponds to an m_2 -th order irrep of $SO(2)$ and the input field to an m_1 -th order irrep, as shown in Figure 2(a), we find the following sparsity structure of κ

$$\hat{\kappa}_{mn}^l = \hat{\kappa}_{m_1 n}^l \delta(m - m_1) \delta(n - m_2).$$

By applying the inverse Fourier Transform, we find the kernel $\kappa(R)$ to be:

$$\sum_{l=0}^{\infty} c^l D_{m_1 m_2}^l(R) = e^{-im_1 \alpha} \sum_{l=0}^{\infty} c^l d_{m_1 m_2}^l(\beta) e^{-im_2 \gamma}.$$

Projecting the output Mackey function to the output field leads to (App. C.2)

$$\begin{aligned} f_{out}(\alpha, \beta) &= C \int_{\mathbb{S}^2} e^{-im_2 h(R^{-1}(\alpha', \beta', 0)R(\alpha, \beta, 0))} \\ &\quad \kappa'(R_y(-\beta')R_z(-\alpha')x_{\alpha, \beta}) f_{in}(\alpha', \beta') d\alpha' \sin(\beta') d\beta', \end{aligned}$$

where κ' is a function on the sphere and $\kappa'(\alpha, \beta) = \kappa(\alpha, \beta, 0)$, $x_{\alpha, \beta}$ is point on the sphere and C is a constant. When we apply the convolution on the homogeneous space \mathbb{S}^2 , the twist $e^{-im_2 h(R^{-1}(\alpha', \beta', 0)R(\alpha, \beta, 0))}$ appears, which is consistent with the correlation on the homogeneous space derived in (Cohen et al., 2018a) for non-semidirect product groups. This shows the difficulty of implementing convolution on the homogeneous space directly in the spatial domain. Therefore, it is more efficient to implement the convolution in the spectral domain. When the input and the output are scalar fields, the filter is isotropic, which is consistent with the results in (Esteves et al., 2017).

Sparsity in the $SE(2)$ -spectrum Consider the group $G = SE(2) = \mathbb{R}^2 \times SO(2)$, where the homogeneous space is \mathbb{R}^2 and stabilizer subgroup is $SO(2)$. Any $g \in G$ can be parameterized as $(x, \theta) = (a, \phi, \theta)$, where we identify $x \in \mathbb{R}^2$ with its action t_x , etc; and where $a = |x|$ is the modulus of x , and for nonzero $a > 0$, $x/a = e^{i\theta}$, for $\theta \in [0, 2\pi)$. The unitary irreducible representations of $SE(2)$ are indexed by integers v and $p \in \mathbb{R}^+$, and with J_v denoting the v -th order Bessel function of the first kind, they have the form

$$U_{mn}(g, p) = i^{n-m} e^{-i(n\theta + (m-n)\phi)} J_{n-m}(pa).$$

The lifting process and its inverse is:

$$\begin{aligned} f \uparrow^G(a, \phi, \theta) &= e^{-im\theta} f(a, \phi) \\ f(a, \phi) &= f \uparrow^G(a, \phi, 0), \end{aligned}$$

where m is the corresponding order of the irreducible representation of the field. Then the Fourier Transform of a Mackey function $f \uparrow^G \in \mathcal{L}^2(G)$ is

$$\begin{aligned} \widehat{f \uparrow^G}_{mn} &= \int_{g \in G} f \uparrow^G(g) \overline{U_{mn}(g, p)} dg \\ &= \int_{g \in G} e^{-ik\theta} f(a, \phi) i^{m-n} e^{i(n\theta + (m-n)\phi)} J_{n-m}(pa) dg \\ &= \int_{(a, \phi) \in G/H} f(a, \phi) i^{m-n} J_{n-m}(pa) e^{i(m-n)\phi} \\ &\quad \int_H e^{-ik\theta} e^{in\theta} d\theta. \end{aligned}$$

Further, $\widehat{f \uparrow^G}_{mn}$ equals

$$\int_{(a, \phi) \in G/H} f(a, \phi) i^{m-n} J_{m-n}(pa) e^{i(m-n)\phi} * 2\pi \delta(n - k).$$

The convolution on $SE(2)$ has the form $(l_1 * l_2)(g) = \int_{k \in G} l_1(k^{-1}g) l_2(k) dk$. The convolution theorem for $SE(2)$ states that $\mathcal{F}(l_1 * l_2)(p) = \hat{l}_2(p) \hat{l}_1(p)$, which can be viewed as the multiplication of two infinite dimensional matrices.

As shown in figure 2(b), when the input and output of the convolution are Mackey functions, the Fourier coefficients are nonzero on the m_1 -th column and m_2 -th column. Therefore, we obtain the kernel

$$\mathcal{F}(\kappa)(p)_{mn} = \hat{\kappa}(p)_{mn} \delta(m - m_1) \delta(n - m_2).$$

By applying the inverse Fourier Transform, we find the kernel to be:

$$\kappa(g) = \int_0^\infty c_p i^{m_2 - m_1} e^{-i(m_2\theta + (m_1 - m_2)\phi)} J_{m_2 - m_1}(pa) p dp.$$




Figure 2. Convolution on $SO(3)$ and $SE(3)$ in the spectral domain. We illustrate the sparsity of the Fourier coefficients of the input, output and kernel.

Mackey functions can be projected to \mathbb{R}^2 to find

$$\begin{aligned} \kappa'(x) &= \int_0^\infty c_p i^{m_2-m_1} e^{-i(m_1-m_2)\phi} J_{m_2-m_1}(pa) p dp \\ &= e^{i(m_2-m_1)\phi} R(a), \end{aligned}$$

where c_p is a constant, and R is a radial function. This exactly recovers the form of the kernel in harmonic networks (Worrall et al., 2017).

In App. C.3, we provide another example, $SE(3)$; we have a similar analysis and conclusion. The block sparsity stated in Prop. 3.2 and Cor. 3.3 also exist in the $SE(3)$ -spectrum for Mackey functions and the corresponding kernels. The derived kernels are consistent with the form of tensor field networks in Thomas et al. (2018) and 3D steerable CNNs in Weiler et al. (2018).

3.4 Equivariant nonlinearity

The whole feature map consists of fields of different types, i.e., it can be written as $f(x) = \bigoplus_i f^{l_i}(x)$ where l_i is the type of the field. Therefore, the group action on a feature map is the direct sum of the associated induced representations, $\bigoplus_i \text{Ind}_H^G \rho^{l_i}$. For an equivariant nonlinearity $\sigma : V^{in} \rightarrow V^{out}$, we need

$$\bigoplus_j \text{Ind}_H^G \rho^{l_{out_j}} \circ \sigma = \sigma \circ \bigoplus_i \text{Ind}_H^G \rho^{l_{in_i}}. \quad (2)$$

The most common strategy to obtain equivariant nonlinear maps in modern deep learning is to apply a fixed nonlinear function ξ elementwise to each coordinate of the feature f in some fixed basis. In general, this does not satisfy equivariance. However, this nonlinearity is equivariant for the regular representation, i.e., $\rho_{reg}^G(g) \circ \xi = \xi \circ \rho_{reg}^G(g)$, as mentioned in Cohen & Welling (2016):

$$\begin{aligned} [\rho_{reg}^G(g) \circ \xi(l)](k) &= [\xi(l)](g^{-1}k) = \xi(l(g^{-1}k)) \\ &= \xi([\rho_{reg}^G(g)(l)](k)) = [\xi \circ \rho_{reg}^G(g)(l)](k), \end{aligned}$$

where $l : G \rightarrow V$. Therefore, we can leverage the lifting isomorphism Λ , which satisfies that $\rho_{reg}^G \circ \Lambda = \Lambda \circ \text{Ind}_H^G \rho$, to lift the features to the corresponding Mackey functions. When

there are multiple types of features $\{f^{l_i} : G/H \rightarrow V\}$, we simply apply the sum of these Mackey functions. We denote the composition of the lifting isomorphism and sum operation by $\bar{\Lambda}$, so that

$$\begin{aligned} l(g) &= [\bar{\Lambda}(\bigoplus_i f^{l_i})](g) = \sum_i [\Lambda_i(f^{l_i})](g) \\ &= \sum_i f^{l_i \uparrow G}(g) = \sum_i \rho^{l_i}(\mathfrak{h}(g)^{-1}) f^{l_i}(gH), \quad (3) \end{aligned}$$

where Λ_i is the lifting isomorphism for the field f^{l_i} . Clearly, $\rho_{reg}^G \circ \bar{\Lambda} = \bar{\Lambda} \circ \bigoplus_i \text{Ind}_H^G \rho^{l_i}$. In general, l is not a Mackey function, only a sum of different Mackey functions. Its spectral domain is shown in Figure 3.

After applying an elementwise nonlinearity to the group function, we need to project the group function to the homogeneous space. The signal may be any function on the group and its Fourier coefficients may be nonsparse, as shown in Figure 3. To maintain equivariance, we propose to convolve with a designed kernel as shown in Figure 3 to find a Mackey function, and then project to the homogeneous space via

$$f(x) = \int \kappa(g^{-1}s(x))l(g)dg.$$

In Figure 3, we know that the kernel κ is also a Mackey function, satisfying $\kappa(gh) = \rho(h^{-1})\kappa(g)$. Therefore $\kappa(g)$ can be expressed as $\kappa(g) = \rho(\mathfrak{h}(g)^{-1})\kappa(s(gH)) = \rho(\mathfrak{h}(g)^{-1})\kappa'(gH)$, for $\kappa' = \kappa \circ s$. Then the convolution becomes:

$$f(x) = \int \rho(\mathfrak{h}(g^{-1}s(x))^{-1})\kappa'(g^{-1}x)l(g)dg,$$

where ρ is the irreducible representation corresponding to the field f . Let us denote this projection by P^κ . The convolution is equivariant, i.e., $\text{Ind}_H^G \rho \circ P^\kappa = P^\kappa \circ \rho_{reg}^G$.

Therefore, the nonlinearity σ is the composition of lifting, elementwise nonlinearity and projection. We prove the




Figure 3. Activation layer: (a). Lifting inputs $\{f^{l_{in_i}} : G/H \rightarrow V_{l_{in_i}}\}$ of different orders to the Mackey functions $\{f^{l_{in_i} \uparrow^G} : G \rightarrow V_{l_{in_i}}\}$ and summing them up to find $l : G \rightarrow V$. (b). Element-wise activation on l . (c). Convolution with the designed kernel to find Mackey functions $\{f^{l_{out_i} \uparrow^G} : G \rightarrow V_{l_{out_i}}\}$ of different orders. (d). Projection to the homogeneous space to find fields $\{f^{l_{out_i}} : G/H \rightarrow V_{l_{out_i}}\}$ of different orders.

equivariance of σ :

$$\begin{aligned} \sigma \circ \bigoplus_i \text{Ind}_H^G \rho^{l_{in_i}} &= \left(\bigoplus_j P^{\kappa_{out_j}} \right) \circ \xi \circ \bar{\Lambda} \circ \bigoplus_i \text{Ind}_H^G \rho^{l_{in_i}} \\ &= \left(\bigoplus_j P^{\kappa_{out_j}} \right) \circ \xi \circ \rho_{reg}^G \circ \bar{\Lambda} = \left(\bigoplus_j P^{\kappa_{out_j}} \right) \circ \rho_{reg}^G \circ \xi \circ \bar{\Lambda} \\ &= \bigoplus_j \text{Ind}_H^G \rho^{l_{out_j}} \circ \left(\bigoplus_j P^{\kappa_{out_j}} \right) \circ \xi \circ \bar{\Lambda} \\ &= \bigoplus_j \text{Ind}_H^G \rho^{l_{out_j}} \circ \sigma, \end{aligned}$$

where κ_{out_j} is the constrained kernel corresponding to $\text{Ind}_H^G \rho^{l_{out_j}}$.

This is exactly the general form of non-linearity which treats the tensor fields on the homogeneous space as the Fourier coefficients. Taking $SE(3)$ as an example, suppose $G = SE(3)$ and $H = SO(3)$. We use the real form of the Wigner D -matrix. Lifting multiple fields $\{f^l : \mathbb{R}^3 \rightarrow V\}$ to the function on the group has the form

$$l_n(x, R) = \sum_{l=0}^{l_{\max}} \sum_{m=-l}^l D_{nm}^l(R^{-1}) f_m^l(x),$$

where D^l is the Wigner D -matrix, the irreducible representation for the field f^l .

When we take the 0-th element of $l : SE(3) \rightarrow V$, we find

$$\begin{aligned} l_0(x, R) &= \sum_{l=0}^{l_{\max}} \sum_{m=-l}^l D_{0m}^l(R^{-1}) f_m^l(x) \\ &= \sum_{l=0}^{l_{\max}} \sum_{m=-l}^l D_{m0}^l(R) f_m^l(x) = \sum_{l=0}^{l_{\max}} \sum_{m=-l}^l D_{m0}^l(r) f_m^l(x) \\ &= \sum_{l=0}^{l_{\max}} \sum_{m=-l}^l Y_m^l(r) f_m^l(x) = \mathcal{F}^+(f(x))(r), \end{aligned}$$

where $R = (\alpha, \beta, \gamma) \in G$ and $r = (\alpha, \beta) \in \mathbb{S}^2$. This is exactly the form of the signals on the sphere obtained in (Poulenard & Guibas, 2021) through the inverse Spherical Harmonics Transform (iSHT).

On the other hand, when we take $\kappa(x)_{ij} = \delta(i)\delta(j)\delta(x)$, the projection is equivalent to the Spherical Harmonics Transform (SHT) in (Poulenard & Guibas, 2021), as

$$\begin{aligned} f_m^l(x) &= \sum_{n,j} \int D_{mn}^l(h(g^{-1}s(x))^{-1}) \kappa_{nj}(g^{-1}x) \xi(l_j(g)) dg \\ &= C_1 \int D_{m0}^l(R) \xi(l_0(x, R)) dR \\ &= C_2 \int Y_m^l(r) \xi(\mathcal{F}^+(f(x))(r)) dr, \end{aligned}$$

where C_1 and C_2 are constants. The second equality holds because κ_{nj} can be nonzero only when $gH = x$, $n = 0$ and $j = 0$.

To keep the network simple, we do not use trainable weights for the kernel. Instead, we take $\kappa(x)_{ij} = \delta(x)\delta(i-j)$, where the first δ is the Dirac delta function and the second is the Kronecker delta function as used previously. Then the projection becomes $f_i(x) = \int_H \sum_j \rho_{ij}(h) l_j(s(x)h) dh$. Weiler & Cesa (2019) describe such a nonlinearity when the stabilizer subgroup H is $O(2)$. In the spectral domain, we can extract the corresponding column from the Fourier matrix of $l(g)$ for simplicity.

4 Implementation and Results

4.1 SO(3): Vector field prediction on the sphere

We experimentally study equivariant vector field prediction on the spherical vector field MNIST (SVMNIST), a dataset proposed in (Esteves et al., 2020). We build a U-Net structure which takes the grayscale spherical image as input, and outputs a vector field on the sphere. The prediction target corresponds to the image gradients of the MNIST characters

when they are mapped to a sphere as shown in Figure 4. We now explain the key components (convolution layer and activation layer) of our U-Net.



Figure 4. The input, output, and ground truth of the vector field prediction task.




Figure 5. Direction change of nonzero-order features in the activation layer. The intensity in the figure reflects the angle difference before and after the activation.

Convolution Layer: We implement convolution in the spectral domain and parameterize the kernel via its Fourier coefficients. As illustrated in Section 3.3.1, we implement a fast Fourier Transform that integrates over \mathbb{S}^2 . For any hidden layers in the U-Net, we use features of 0th and 1st order, both with the same number of channels, to increase expressivity. This is different from spherical convolution methods that only process features of order 0 (Cohen et al., 2018b), and computationally more efficient than convolution directly on $SO(3)$ (Cohen et al., 2018b). Esteves et al. (2020) define and implement convolution in the spectral domain through spin-weighted spherical harmonics, which is equivalent to the linear part in our method. In our paper, this method is a natural implication of the spectral sparsity, shown for all groups where a Fourier transform exists, and not just for $SO(3)$.

Activation Layer: We first lift the fields to the Mackey function through $f \uparrow^G(\alpha, \beta, \gamma) = e^{-im\gamma} f(\alpha, \beta)$. Then, we sum up the Mackey functions corresponding to different fields. These two steps reduce the computation compared to (Cohen et al., 2018b) due to the sparsity in the Fourier coefficients. The element-wise nonlinearity on features is implemented over $SO(3)$. The activated group function can then be projected back to the homogeneous space by convolution over the Fourier domain, as shown in the spectral part of figure 3, bringing expressivity. Alternatively, we can pick a specific column to output for the Fourier transform,

which reduces computation. In the U-Net, we always take the latter approach for better efficiency.

We report the vector prediction mean-squared errors (Esteves et al., 2020) weighted by the spherical map sampling area. Our U-Net uses a number of parameters similar to the baselines. The comparison is shown in Table 1. Our network outperforms the state-of-the-art equivariant networks, especially when the input has no rotation augmentation and the testing data is rotated (NR/R). In addition to being equivariant for different feature types, our model also has a more expressive nonlinearity layer. The novel nonlinearity entangles different fields and directional information, and can change the direction of the nonzero-order tensor in the activated output, making it more expressive than performing nonlinearity over the invariant norm (Esteves et al., 2020). Figure 5 shows the angle difference of the vector (order-one) feature maps before and after activation in a hidden layer. Please see App. D.1 for more details.

| Method | NR/NR | R/R | NR/R |
|-------------------------------|------------|------------|------------|
| Planar (Esteves et al., 2020) | 0.3 | 5.0 | 9.3 |
| SphCNN (Esteves et al., 2018) | 9.7 | 31.0 | 45.6 |
| SWSCNN (Esteves et al., 2020) | 2.9 | 3.4 | 4.3 |
| Ours | 2.9 | 3.2 | 3.8 |

Table 1. Results of spherical scalar to vector prediction. We report the mean-squared error $\times 10^3$ (lower is better). The baseline methods are planar (convolutional 2D CNNs), SphCNN (Esteves et al., 2018) and SWSCNN (Esteves et al., 2020). NR/NR is nonrotated train and test set; NR/R is nonrotated train set and rotated test set; R/R is rotated train and test set.

4.2 SE(3) Prediction

Designing $SE(3)$ equivariant networks for point sets is an important problem in many application areas, like chemistry and computer vision. The standard method in the current literature is Tensor Field Networks (TFN) (Thomas et al., 2018). Following our theory, we enhance TFN with a novel non-linearity that can capture the directional information of higher order features. We verify its effectiveness on two tasks: (i) QM9 (Ramakrishnan et al., 2014) missing atom prediction and (ii) ModelNet40 (Wu et al., 2015) point cloud shape classification.

Convolution Layer: For $SE(3)$, the homogeneous space is \mathbb{R}^3 . Since $SE(3)$ has a semidirect product structure, and the Fourier Transform is computationally expensive, we analytically derive the kernel for the convolution in the spatial domain, and implement convolution on the homogeneous space \mathbb{R}^3 . As proved in the sparsity of the $SE(3)$ -spectrum, the derived kernel is equivalent to that from (Thomas et al., 2018) and (Weiler et al., 2018); thus, we use the convolution in TFN as implemented in (Geiger et al., 2020)).

Activation Layer: The main difference between our ex-

tended model and the original TFN (Thomas et al., 2018) is the nonlinearity. After lifting the field from \mathbb{R}^3 to $SE(3)$, we project it to \mathbb{S}^2 attached to every point (using bandwidth 8) to save memory, while keeping equivariance (see App. D.2 for details). To obtain more expressivity, we apply a small per-point MLP to the group function following (Poulenard & Guibas, 2021), which also maintains equivariance. Finally, our implementation of the projection from $SE(3)$ to \mathbb{R}^3 is equivalent to that in (Poulenard & Guibas, 2021), as shown in Sec. 3.4.

4.2.1 SE(3): Molecular Structure Completion

To show advantages beyond scalar field prediction, we conduct experiments to predict missing atoms in molecules. Here, positional prediction requires equivariant vector prediction. The setup of the experiment follows (Thomas et al., 2018). During training, we randomly remove an atom from the molecule, and use our model to predict the atom type (order-0) and position (order-1) of the atom. During testing, we iteratively remove all atoms one at a time for every molecule. We use two metrics for evaluation. The distance error is the average error of the predicted position from the ground truth atom position; the accuracy is the proportion of the instances correctly predicted and with a distance error less than 0.5\AA . More details are in App. D.3.

We report the accuracy and distance MAE in Table 2. Since rotations exist naturally in molecule structures, the result illustrates the equivariance of our model. Our model generalizes well to test datasets of molecules with different numbers of atoms, and outperforms tensor field networks on every test dataset, which shows the effectiveness of our activation compared to the norm activation in the TFN.

| Atoms | Accuracy↑(%) | | Distance↓(Å) | |
|-------|--------------|-------------|--------------|-------------|
| | TFN | Ours | TFN | Ours |
| 19 | 93.9 | 98.0 | 0.14 | 0.06 |
| 23 | 96.5 | 97.1 | 0.13 | 0.10 |
| 25-29 | 97.3 | 98.3 | 0.16 | 0.10 |

Table 2. Results for missing atom prediction. We have three test datasets with 5-18, 23 and 25-29 atoms in one molecule, respectively. Every dataset has 1000 molecules.

4.2.2 SE(3): ModelNet40 classification

Using our novel activation, we build a point cloud classification network based on the approaches in (Poulenard & Guibas, 2021; Thomas et al., 2018). We report the classification accuracy in Table 3. The dataset we use for the 3D object shape recognition is ModelNet40 (Wu et al., 2015), and it consists of 12311 3D shapes (9843 for training and 2468 for test) over 40 categories. Please see App. D.4 for more details. Our method has performance similarly to state-of-the-art methods.

| Methods | z/z | $z/SO(3)$ | $SO(3)/SO(3)$ |
|---------------|-------------|-------------|---------------|
| Spherical-CNN | 88.9 | 76.7 | 86.9 |
| a^3S -CNN | 89.6 | 87.9 | 88.7 |
| SFCNN | 91.4 | 84.8 | 90.1 |
| TFN | 88.5 | 85.3 | 87.6 |
| RI-Conv | 86.5 | 86.4 | 86.4 |
| SPHNet | 87.7 | 86.6 | 87.6 |
| ClusterNet | 87.1 | 87.1 | 87.1 |
| GC-Conv | 89.0 | 89.1 | 89.2 |
| RI-Framework | 89.4 | 89.4 | 89.3 |
| VN-PointNet | 77.5 | 77.5 | 77.2 |
| VN-DGCNN | 89.5 | 89.5 | 90.2 |
| TFN[mlp]-P | 89.7 | 89.7 | 89.7 |
| Ours | 89.7 | 89.7 | 89.3 |

Table 3. Classification accuracy in three train/test setups. Here z stands for aligned data augmented by random rotations around the vertical axis and $SO(3)$ indicates augmentation by random rotations. The quantitative results of previous methods are from (Deng et al., 2021). Please see Table 4 for the source of each method.

5 Limitation and Discussion

This paper provides a unified perspective and practical techniques for designing the linear and activation layers for equivariant networks. Several interesting directions remain for future exploration. Although our theory covers many groups of interest for a wide range of applications, groups with trivial stabilizers, among others, and groups without Fourier transforms, remain unexplored. The input and output homogeneous spaces of each layer in this paper are the same, and the setting with different homogeneous spaces (different stabilizer groups) needs further study. Another future direction is to use our unified architecture for other groups, for example, the Lorentz group and finite permutation groups.

6 Conclusions

This paper provides a Fourier perspective for group convolutional neural networks on homogeneous spaces. We discovered a form of spectral sparsity and used it in designing the kernel and the nonlinearity in a unified way. Our networks showed their effectiveness in several tasks.

Acknowledgements

We gratefully acknowledge financial support by the grants NSF TRIPODS 1934960, ARO MURI W911NF-20-1-0080, ONR N00014-17-1-2093, and ARL DCIST CRA W911NF-17-2-0181. We also thank Adrien Poulenard for his help on Tensor Field Networks implementation and all reviewers for their constructive feedback.

References

- Anderson, B., Hy, T.-S., and Kondor, R. Cormorant: Covariant molecular neural networks. *arXiv preprint arXiv:1906.04015*, 2019.
- Atzmon, M., Nagano, K., Fidler, S., Khamis, S., and Lipman, Y. Frame averaging for equivariant shape space learning. *arXiv preprint arXiv:2112.01741*, 2021.
- Bekkers, E. J. B-spline cnns on lie groups. *arXiv preprint arXiv:1909.12057*, 2019.
- Chen, C., Li, G., Xu, R., Chen, T., Wang, M., and Lin, L. Clusternet: Deep hierarchical cluster network with rigorously rotation-invariant representation for point cloud analysis. In *IEEE Comp. Vision Patt. Recog.*, pp. 4994–5002, 2019.
- Chen, H., Liu, S., Chen, W., Li, H., and Hill, R. Equivariant point network for 3d point cloud analysis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 14514–14523, 2021.
- Chirikjian, G. S., Kyatkin, A. B., and Buckingham, A. Engineering applications of noncommutative harmonic analysis: with emphasis on rotation and motion groups. *Appl. Mech. Rev.*, 54(6):B97–B98, 2001.
- Cobb, O. J., Wallis, C. G., Mavor-Parker, A. N., Marignier, A., Price, M. A., d’Avezac, M., and McEwen, J. D. Efficient generalized spherical cnns. *arXiv preprint arXiv:2010.11661*, 2020.
- Cohen, T. and Welling, M. Group equivariant convolutional networks. In *International conference on machine learning*, pp. 2990–2999. PMLR, 2016.
- Cohen, T., Geiger, M., and Weiler, M. A general theory of equivariant cnns on homogeneous spaces. *arXiv preprint arXiv:1811.02017*, 2018a.
- Cohen, T. S., Geiger, M., Köhler, J., and Welling, M. Spherical cnns. *arXiv preprint arXiv:1801.10130*, 2018b.
- De Haan, P., Weiler, M., Cohen, T., and Welling, M. Gauge equivariant mesh cnns: Anisotropic convolutions on geometric graphs. *arXiv preprint arXiv:2003.05425*, 2020.
- Deng, C., Litany, O., Duan, Y., Poulenard, A., Tagliasacchi, A., and Guibas, L. Vector neurons: A general framework for so(3)-equivariant networks. *arXiv preprint arXiv:2104.12229*, 2021.
- Dym, N. and Maron, H. On the universality of rotation equivariant point cloud networks. *arXiv preprint arXiv:2010.02449*, 2020.
- Esteves, C., Allen-Blanchette, C., Zhou, X., and Daniilidis, K. Polar transformer networks. *arXiv preprint arXiv:1709.01889*, 2017.
- Esteves, C., Allen-Blanchette, C., Makadia, A., and Daniilidis, K. Learning so(3) equivariant representations with spherical cnns. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pp. 52–68, 2018.
- Esteves, C., Xu, Y., Allen-Blanchette, C., and Daniilidis, K. Equivariant multi-view networks. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 1568–1577, 2019.
- Esteves, C., Makadia, A., and Daniilidis, K. Spin-weighted spherical cnns. *arXiv preprint arXiv:2006.10731*, 2020.
- Finzi, M., Stanton, S., Izmailov, P., and Wilson, A. G. Generalizing convolutional neural networks for equivariance to lie groups on arbitrary continuous data. In *International Conference on Machine Learning*, pp. 3165–3176. PMLR, 2020.
- Finzi, M., Welling, M., and Wilson, A. G. A practical method for constructing equivariant multilayer perceptrons for arbitrary matrix groups. *arXiv preprint arXiv:2104.09459*, 2021.
- Folland, G. B. *A course in abstract harmonic analysis*, volume 29. CRC press, 2016.
- Fuchs, F. B., Worrall, D. E., Fischer, V., and Welling, M. Se(3)-transformers: 3d roto-translation equivariant attention networks. *arXiv preprint arXiv:2006.10503*, 2020.
- Fukushima, K. Neocognitron: A self-organizing neural network model for a mechanism of pattern recognition unaffected by shift in position. *Biological cybernetics*, 36(4):193–202, 1980.
- Fulton, W. and Harris, J. *Representation theory: a first course*, volume 129. Springer Science & Business Media, 2013.
- Gauthier, J.-P., Bornard, G., and Silbermann, M. Motions and pattern analysis: harmonic analysis on motion groups and their homogeneous spaces. *IEEE transactions on systems, man, and cybernetics*, 21(1):159–172, 1991.
- Geiger, M., Smidt, T., M., A., Miller, B. K., Boomsma, W., Dice, B., Lapchevskyi, K., Weiler, M., Tyszkiewicz, M., Batzner, S., Uhrin, M., Frelsen, J., Jung, N., Sanborn, S., Rackers, J., and Bailey, M. *Euclidean neural networks: e3nn*. Zenodo, 2020.
- Gross, K. I. On the evolution of noncommutative harmonic analysis. *The American Mathematical Monthly*, 85(7): 525–548, 1978.

- Halmos, P. R. *Measure theory*, volume 18. Springer, 2013.
- Hutchinson, M. J., Le Lan, C., Zaidi, S., Dupont, E., Teh, Y. W., and Kim, H. Lietransformer: Equivariant self-attention for lie groups. In *International Conference on Machine Learning*, pp. 4533–4543. PMLR, 2021.
- Keriven, N. and Peyré, G. Universal invariant and equivariant graph neural networks. *Advances in Neural Information Processing Systems*, 32:7092–7101, 2019.
- Kondor, R. and Trivedi, S. On the generalization of equivariance and convolution in neural networks to the action of compact groups. In *International Conference on Machine Learning*, pp. 2747–2755. PMLR, 2018.
- Kondor, R., Lin, Z., and Trivedi, S. Clebsch–gordan nets: a fully fourier space spherical convolutional neural network. *Advances in Neural Information Processing Systems*, 31: 10117–10126, 2018.
- LeCun, Y., Boser, B., Denker, J. S., Henderson, D., Howard, R. E., Hubbard, W., and Jackel, L. D. Backpropagation applied to handwritten zip code recognition. *Neural computation*, 1(4):541–551, 1989.
- Li, X., Li, R., Chen, G., Fu, C.-W., Cohen-Or, D., and Heng, P.-A. A rotation-invariant framework for deep point cloud analysis. *arXiv preprint arXiv:2003.07238*, 2020.
- Liu, M., Yao, F., Choi, C., Sinha, A., and Ramani, K. Deep learning 3d shapes using alt-az anisotropic 2-sphere convolution. In *Int. Conf. Learning Representations*, 2018.
- Maron, H., Ben-Hamu, H., Shamir, N., and Lipman, Y. Invariant and equivariant graph networks. *arXiv preprint arXiv:1812.09902*, 2018.
- Maron, H., Litany, O., Chechik, G., and Fetaya, E. On learning sets of symmetric elements. In *International Conference on Machine Learning*, pp. 6734–6744. PMLR, 2020.
- Poulenard, A. and Guibas, L. J. A functional approach to rotation equivariant non-linearities for tensor field networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 13174–13183, 2021.
- Poulenard, A., Rakotosaona, M.-J., Ponty, Y., and Ovsjanikov, M. Effective rotation-invariant point cnn with spherical harmonics kernels. In *2019 International Conference on 3D Vision (3DV)*, pp. 47–56. IEEE, 2019.
- Puny, O., Atzmon, M., Ben-Hamu, H., Smith, E. J., Misra, I., Grover, A., and Lipman, Y. Frame averaging for invariant and equivariant network design. *arXiv preprint arXiv:2110.03336*, 2021.
- Ramakrishnan, R., Dral, P. O., Rupp, M., and Von Lilienfeld, O. A. Quantum chemistry structures and properties of 134 kilo molecules. *Scientific data*, 1(1):1–7, 2014.
- Rao, Y., Lu, J., and Zhou, J. Spherical fractal convolutional neural networks for point cloud recognition. In *IEEE Comp. Vision Patt. Recog.*, pp. 452–460, 2019.
- Romero, D. W. and Cordonnier, J.-B. Group equivariant stand-alone self-attention for vision. *arXiv preprint arXiv:2010.00977*, 2020.
- Satorras, V. G., Hoogeboom, E., and Welling, M. E (n) equivariant graph neural networks. In *International Conference on Machine Learning*, pp. 9323–9332. PMLR, 2021.
- Sosnovik, I., Szmaja, M., and Smeulders, A. Scale-equivariant steerable networks. *arXiv preprint arXiv:1910.11093*, 2019.
- Stroppel, M. *Locally compact groups*, volume 3. European Mathematical Society, 2006.
- Thomas, N., Smidt, T., Kearnes, S., Yang, L., Li, L., Kohlhoff, K., and Riley, P. Tensor field networks: Rotation-and translation-equivariant neural networks for 3d point clouds. *arXiv preprint arXiv:1802.08219*, 2018.
- Weiler, M. and Cesa, G. General $e(2)$ -equivariant steerable cnns. *arXiv preprint arXiv:1911.08251*, 2019.
- Weiler, M., Geiger, M., Welling, M., Boomsma, W., and Cohen, T. 3d steerable cnns: Learning rotationally equivariant features in volumetric data. *arXiv preprint arXiv:1807.02547*, 2018.
- Worrall, D. and Brostow, G. Cubenet: Equivariance to 3d rotation and translation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pp. 567–584, 2018.
- Worrall, D. E., Garbin, S. J., Turmukhambetov, D., and Brostow, G. J. Harmonic networks: Deep translation and rotation equivariance. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 5028–5037, 2017.
- Wu, Z., Song, S., Khosla, A., Yu, F., Zhang, L., Tang, X., and Xiao, J. 3d shapenets: A deep representation for volumetric shapes. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 1912–1920, 2015.
- Zhang, Z., Hua, B.-S., Rosen, D. W., and Yeung, S.-K. Rotation invariant convolutions for 3d point clouds deep learning. In *IEEE International Conference on 3D Vision*, pp. 204–213, 2019.

Zhang, Z., Hua, B.-S., Chen, W., Tian, Y., and Yeung, S.-K.
Global context aware convolutions for 3d point cloud
understanding. *arXiv preprint arXiv:2008.02986*, 2020.

A Preliminary

A.1 Group actions and homogeneous spaces

Consider a group G acting on the homogeneous space X , and let x_0 be the origin of X inducing the set of group elements $H = \{h \in G | hx_0 = x_0\}$ that leaves x_0 unchanged, i.e., the stabilizer subgroup H of x_0 in G . The set of left cosets $gH := \{gh | h \in H\}$ is called the left quotient space G/H , and is isomorphic to the homogeneous space X . Take $G = SE(2)$ as the acting group and $X = \mathbb{R}^2$ as its homogeneous space. Any element in $g \in SE(2)$ can be denoted as $g = (t_x, r_\theta)$, where t_x is the translation by the vector $x \in \mathbb{R}^2$, and r_θ is the rotation by the angle $\theta \in [0, 2\pi)$. Any rotation r_θ leaves the point $x_0 = (0, 0) \in X$ unchanged, and these rotations compose the stabilizer subgroup $H = SO(2)$. For any element $g = (t_x, r_\theta) \in SE(2)$, $gH = \{(t_x, r_\theta r_{\theta'}) | r_{\theta'} \in SO(2)\} = \{(t_x, \star)\}$ and all elements in the left coset gH map x_0 to x . The left quotient space $G/H = \{x \in \mathbb{R}^2 : (t_x, \star)\}$ is isomorphic to the homogeneous space \mathbb{R}^2 .

The group G can be viewed as a principal bundle through the partition of the group into cosets. The base space is G/H , and the canonical fiber is H , with the projection map $p : G \rightarrow G/H$, $p(g) = gH = x$ and the section $s : G/H \rightarrow G$ such that $p \circ s = id_{G/H}$, the identity map on G/H . The action of G induces a twist of the fibers as $gs(x) = s(gx)h(g, x)$ where $h : G \times G/H \rightarrow H$ is the twist function. For simplicity, we denote $h(g, eH)$ as $h(g)$. When the group G is a semidirect product group $G/H \rtimes H$, h does not depend on the choice of x in the homogeneous space and can be simplified to $h(g, x) = h(g)$. We use $SE(2)$ as an example to illustrate the bundle structure where the base space is \mathbb{R}^2 and the fiber is $SO(2)$. Then the projection map $p : SE(2) \rightarrow \mathbb{R}^2$ is $p((t_x, r_\theta)) = x$ and the section $s : \mathbb{R}^2 \rightarrow SE(2)$ is $s(x) = (t_x, r_0)$. For any $g = (t_x, r_\theta) \in SE(2)$ and any $x' \in \mathbb{R}^2$, we have $(t_x, r_\theta)(t_{x'}, r_0) = (t_{x+r_\theta x'}, r_0)(0, r_\theta)$, therefore the twist function $h : SE(2) \times \mathbb{R}^2 \rightarrow SO(2)$ is $h((t_x, r_\theta), x') = h((t_x, r_\theta)) = r_\theta$.

A.2 Irreducible Representations

Let V be a vector space over a field K . A representation of a group G on V is a homomorphism $\rho : G \rightarrow GL(V)$, i.e., for any $g_1, g_2 \in G$, $\rho(g_1 g_2) = \rho(g_1) \rho(g_2)$, where $GL(V)$ is the general linear group over V .

If the subspace W of V is invariant under the action of all group elements, that is, for any $g \in G$ and any $w \in W$, we have $\rho(g)w \in W$, we call it a sub-representation of ρ . If ρ has only two sub-representations, the whole space V and $\{0\} \subset V$, then ρ is called an irreducible representation. If for any $g \in G$, $\rho(g^{-1}) = \overline{\rho(g)}^\top$, the representation ρ is called a unitary representation. Every locally compact group has a unitary representation.

Let us denote by U a unitary irreducible representation. In any particular basis, we can view this as a matrix, and denote by U_{ij} its entry in row i and column j . Irreducible representations satisfy the group orthogonality theorem; any two unitary irreducible representations U^{l_1} and U^{l_2} satisfy $\langle U_{m_1 n_1}^{l_1}, U_{m_2 n_2}^{l_2} \rangle = \delta(l_1 - l_2) \delta(m_1 - m_2) \delta(n_1 - n_2)$, where $\langle U_{m_1 n_1}^{l_1}, U_{m_2 n_2}^{l_2} \rangle$ is the inner product: $\int_G \overline{U_{m_1 n_1}^{l_1}(g)} U_{m_2 n_2}^{l_2}(g) dg$.

Two representations ρ_1 and ρ_2 are equivalent when there exists an invertible matrix Q such that for any $g \in G$, $Q^{-1} \rho_1(g) Q = \rho_2(g)$. For a compact group or a semisimple Lie group, any representation U on a Hilbert space (thus any finite-dimensional representation) is equivalent to the direct sum of unitary irreducible representations, that is, there exists an invertible matrix Q such that $U = Q^{-1} \bigoplus_i U^i Q$, where $(U^i)_{i \in I}$ are irreducible representations of G (indexed by some set I that is suppressed for brevity) and $\bigoplus_i U^i$ is a block diagonal matrix with blocks U^i , $i \in I$. For details, we refer to (Folland (2016), Ch. 3).

For a group G and a subgroup H , and for any representation $\rho : G \rightarrow GL(V)$ of G , a restricted representation $\rho|_H : H \rightarrow GL(V)$ is the restriction of ρ to H , namely $\rho|_H(h) = \rho(h)$. Even when ρ is irreducible, $\rho|_H$ may still be reducible.

Given a representation ρ of the group G , the dual representation $\bar{\rho}$ is defined by $\bar{\rho}(g) = (\rho(g^{-1}))^\top$. When ρ is a unitary representation, $\bar{\rho}$ is the complex conjugate of ρ . The dual representation $\bar{\rho}$ may not be equivalent to the representation ρ . For example, for $SO(2)$ the representation $\theta \mapsto e^{im\theta}$ is not equivalent to $\theta \mapsto e^{-im\theta}$ unless $m = 0$. However, the irreducible representations are self-dual for some groups, including $SO(3)$ and the special unitary group of order two, $SU(2)$ (consisting of two-by-two complex-valued unitary matrices having unit determinant, with the multiplication operation) (Fulton & Harris (2013), Ch. 8).

A.3 Unimodular Separable Locally Compact Groups of Type I

A group G is a unimodular group if its left Haar measure is also a right Haar measure (Folland (2016) Ch. 2). A topological group is separable (second countable) if its topology has a countable base (Halmos (2013), Ch. 0). If the underlying topology of the group is locally compact and Hausdorff, we call the group locally compact (Stroppel (2006), Ch. 2). A group is said to be of Type I if each of its primary representations (factor representations) is a direct sum of copies of some irreducible representation (Folland (2016), Ch. 7). Following (Gross (1978), Folland (2016), Ch. 7), the Fourier Transform is well defined for any unimodular separable locally compact group of Type I.

A.4 Clebsch-Gordan Decomposition

The dual group \hat{G} is finite or countable for finite or compact G , respectively. For such groups, since $U_{m_1 n_1}(g, p_1)U_{m_2 n_2}(g, p_2) \in \mathcal{L}^2(G)$ for any $p_1, p_2 \in \hat{G}$, we can apply the Inverse Fourier Transform to express

$$U_{m_1 n_1}(g, p_1)U_{m_2 n_2}(g, p_2) = \sum_{p \in \hat{G}, m, n} \mathcal{C}_{p_1, p_2, m_1, m_2, n_1, n_2}^{p, m, n} U_{mn}(g, p)$$

where m_1, m_2, n_1, n_2, m, n are the appropriate row and column indices. The Clebsch-Gordan coefficients $\mathcal{C}_{p_1, p_2, m_1, m_2, n_1, n_2}^{p, m, n}$ are the Fourier coefficients $\mathcal{F}(U_{m_1 n_1}(\cdot, p_1)U_{m_2 n_2}(\cdot, p_2))_{mn}(p)$. This so-called Clebsch-Gordan decomposition describes the decomposition of the tensor product of two irreducible representations. For $SO(3)$ and $SU(2)$, it can be shown (Chirikjian et al. (2001), Ch. 10), that $\mathcal{C}_{p_1, p_2, m_1, m_2, n_1, n_2}^{p, m, n} = \mathcal{C}_{p_1, m_1, p_2, m_2}^{p, m} \mathcal{C}_{p_1, n_1, p_2, n_2}^{p, n}$, where $\mathcal{C}_{p_1, m_1, p_2, m_2}^{p, m}$ and $\mathcal{C}_{p_1, n_1, p_2, n_2}^{p, n}$ are the Clebsch-Gordan coefficients as they appear in (Kondor et al., 2018).

B Details and Proofs

B.1 Proof of Lemma 3.1

Since H is the subgroup of G , $U(\cdot, p)|_H$ is a representation of H , but not necessarily an irreducible one. Because H is a compact Lie group, there is an invertible Q such that $Q^{-1}U(h, p)Q = \oplus_{i \in \mathcal{Q}(p)} \rho^i(h)$, where $\{\rho^i(h)\}_{i \in \mathcal{Q}(p)}$ are the irreducible representations of H . Since $Q^{-1}U(\cdot, p)Q$ is a unitary irreducible representation of G , for simplicity, we will use $U(\cdot, p)$ to denote $Q^{-1}U(\cdot, p)Q$. A unitary irreducible representation of G can be expressed as $U(g, p) = U(s(gH)h(g), p)$ where s is the section map and h is the twist function defined in App. A.1. Since $U(\cdot, p)$ is a representation, based on the previous properties, $U(g, p) = U(s(gH), p) \cdot U(h(g), p)$. Thus, the unitary irreducible representation $U(\cdot, p)$ can be decomposed as:

$$U(g, p) = U(s(gH), p) \cdot \oplus_{i \in \mathcal{Q}(p)} \rho^i(h(g)). \quad (4)$$

B.2 Proof of Proposition 3.2

We know that ρ is an unitary irreducible representation of the stabilizer subgroup H . Denote its type as i , writing $\rho = \rho^i$. Writing the definition of a Mackey function in a block form, we have $(f \uparrow^G)_k(g) = \sum_t \rho_{kt}^i(h(g)^{-1}) f_t(gH)$. Using also (4), and decomposing the integral over G into integrals over G/H and H , the Fourier Transform of the k -th element of the vector function $f \uparrow^G$ can be calculated as follows:

$$\begin{aligned} ((f \uparrow^G)_k)_{mn}(p) &= \int_G (f \uparrow^G)_k(g) \overline{U_{mn}(g, p)} dg = \sum_{t, j} \int \rho_{kt}^i(h(g)^{-1}) f_t(gH) \overline{U_{mj}(s(gH), p) U_{jn}(h(g), p)} dg \\ &= \sum_{t, j} \int_{G/H} f_t(x) \overline{U_{mj}(s(x), p)} dx \int_H \rho_{kt}^i(h^{-1}) \overline{U_{jn}(h, p)} dh. \end{aligned}$$

We refer Theorem 2.49 in (Folland, 2016) for the measures used in the above equation. Since $U(h, p) = \oplus_{o \in \mathcal{Q}(p)} \rho^o(h(g))$, we know that U_{jn} is either zero or belongs to a nonzero block of $U(h, p)$. When U_{jn} is zero, then $\rho_{kt}^i(h^{-1}) \overline{U_{jn}(h, p)} dh$ is zero. Else, suppose that U_{jn} is an element of the irreducible representation $\rho_{f_1^p}^{f_1^p(j, n)}$, where f_1^p is a function parametrized by p —i.e., the block U_{jn} belongs to can depend on j, n and p . For the same reason, we have $U_{jn}(h, p) = \rho_{f_2^p(j) f_3^p(n)}^{f_1^p(j, n)}(h)$.

The second integral thus becomes $\int_H \rho_{tk}^i(h) \overline{\rho_{f_2^p(j) f_3^p(n)}^{f_1^p(j, n)}(h)} dh$.

As mentioned in App. A.4, the compact group H has the Clebsch-Gordan decomposition

$$\rho_{m_1 n_1}^{l_1} \rho_{m_2 n_2}^{l_2} = \sum_{l, m, n} C_{l_1, m_1, n_1, l_2, m_2, n_2}^{l, m, n} \rho_{mn}^l.$$

Since $\langle \rho_{mn}^l, \rho_{00}^0 \rangle = \delta(l)\delta(m)\delta(n)$, the integral $\int_H \rho_{tk}^l(h) \rho_{f_2^p(j) f_3^p(n)}^{f_1^p(j, n)}(h) dh$ can only be nonzero when the decomposition of $\rho_{tk}^l(h) \rho_{f_2^p(j) f_3^p(n)}^{f_1^p(j, n)}$ includes the trivial representation ρ_{00}^0 .

Since H is a compact Lie group, by calculating the character of the tensor product of the two irreducible representations ρ^i and $\rho^{f_1^p(j, n)}$ (see App. B.3), $C_{l_1, m_1, n_1, l_2, m_2, n_2}^{00}$ is nonzero only when $\rho^{f_1^p(j, n)}$ is equivalent to the dual representation of ρ^i . This finishes the proof.

B.3 Character of the Tensor Product of Two Irreducible Representations

Suppose ρ' and ρ'' are two unitary irreducible representations (irreps) of a compact group G . The character of the tensor product of these two irreps is $\mathcal{X}_{\rho' \otimes \rho''} = \text{tr}(\rho' \otimes \rho'') = \text{tr}(\rho')\text{tr}(\rho'') = \mathcal{X}_{\rho'} \mathcal{X}_{\rho''}$.

From the Inverse Fourier Transform and the properties of compact groups, we have $\rho' \otimes \rho'' = Q^{-1}(\bigoplus \rho)Q$, where Q is an invertible matrix. Then $\mathcal{X}_{\rho'} \mathcal{X}_{\rho''} = \sum_{\rho} c_{\rho', \rho''}^{\rho} \mathcal{X}_{\rho}$, where $c_{\rho', \rho''}^{\rho}$ are positive integers.

The trivial representation is contained in $\rho \otimes \rho''$ only when ρ and ρ'' are dual representations. This is because $c_{\rho', \rho''}^{\rho} = \int \mathcal{X}_{\rho'} \mathcal{X}_{\rho''} = \int \mathcal{X}_{\rho'} \overline{\mathcal{X}_{\rho''}} = \langle \mathcal{X}_{\rho'}, \mathcal{X}_{\rho''} \rangle$. The integral can be nonzero only when $\overline{\rho''}$ and ρ are equivalent, where $\overline{\rho''}$ is the dual representation of the unitary irreducible representation ρ'' .

B.4 Proof of Corollary 3.3

Due to the convolution theorem (Chirikjian et al. (2001), Ch.8), we have $\widehat{f_2 \uparrow^G}(p) = \widehat{f_1 \uparrow^G}(p) \widehat{\kappa}(p)$. The sparsity pattern exists in $\widehat{f_1 \uparrow^G}(p)$ and $\widehat{f_2 \uparrow^G}(p)$ due to Proposition 3.2. Suppose $\widehat{f_1 \uparrow^G}(p)$ has zeroes in columns A_1 , and $\widehat{f_2 \uparrow^G}(p)$ has zeroes in columns A_2 . Then the equation $\widehat{f_2 \uparrow^G}(p) = \widehat{f_1 \uparrow^G}(p) \widehat{\kappa}(p)$ can be reduced by taking the subsets of rows in the complement of A_1 and columns in the complement of A_2 . It is easy to see that the entries of $\widehat{\kappa}$ outside of these indices do not enter the calculation. Thus, they can be taken to be anything, and in particular as zeroes. This proves the desired claim.

B.5 Discussion of completeness

Revisiting $\int_H \rho_{m_1 n_1}^{l_1}(h) \rho_{m_2 n_2}^{l_2}(h) dh$, if the integral is nonzero, then ρ^{l_2} is equivalent to $\overline{\rho^{l_1}}$. Supposing $\rho^{l_2} = Q \overline{\rho^{l_1}} Q^{-1}$, we have

$$\begin{aligned} \int_H \rho_{m_1 n_1}^{l_1}(h) \rho_{m_2 n_2}^{l_2}(h) dh &= \int \sum_{a, b} \rho_{m_1 n_1}^{l_1} Q_{m_2 a} \overline{\rho^{l_1}}_{ab} Q_{b n_2}^{-1} dh \\ &= \sum_{a, b} Q_{m_2 a} Q_{n_2 b} \int_H \rho_{m_1 n_1}^{l_1} \overline{\rho^{l_1}}_{ab} dh = Q_{m_2 m_1} Q_{n_2 n_1}. \end{aligned}$$

Therefore, the value of the integral is related to Q . We can choose a unitary representation of G , and a basis of the underlying vector space, such that ρ^{l_2} is the dual representation of ρ^{l_1} . This holds because the representation of compact Lie groups can split into an orthogonal direct sum of irreducible finite-dimensional unitary representations (according to the Peter-Weyl Theorem) and the dual representation of a finite-dimensional representation is irreducible.

In this case, when $\widehat{f \uparrow^G}$ is a Mackey function lifted via $f \uparrow^G(g) = \rho(\mathfrak{h}(g)^{-1})f(gH)$ from a field f , $\widehat{f \uparrow^G}_k$ is nonzero only on the k -th column in the corresponding block and $\widehat{f \uparrow^G}_{kmn} = \widehat{f \uparrow^G}_{k+1m(n+1)}$ for any k, m, n . The converse also stands.

For the convolution $(f_{out} \uparrow^G)_a(g) = \int_G \sum_b \kappa(\nu^{-1}g)_{ab} (f_{in} \uparrow^G)_b(\nu)$, in the Fourier domain we have $\widehat{(f_{out} \uparrow^G)}_{amn} = \sum_{b, t} \widehat{f_{in} \uparrow^G}_{bmt} \widehat{\kappa}_{abtn}$. Now, f_{out} obeys the above property and f_{in} is arbitrary, therefore, we have $\widehat{\kappa}_{abmn} = \widehat{\kappa}_{(a+1)b_m(n+1)}$

for any a, b, m, n . Meanwhile, the nonzero columns in $\widehat{\kappa}_{ab}$ and $(\widehat{f_{out} \uparrow^G})_a$ have the same indices, which implies that $\kappa(gh) = \rho_{out}(h^{-1})\kappa(g)$.

On the other hand, $(\widehat{f_{in} \uparrow^G})$ is also a Mackey function, and the relation between its Fourier coefficients can help us simplify $\widehat{\kappa}$. For any κ with the above constraint, we can design a kernel κ^* that satisfies the following constraints:

1. $\widehat{\kappa_{abmn}^*} = \widehat{\kappa_{a(b+1)(m+1)n}^*}$ for any a, b, m, n ;
2. the nonzero rows in $\widehat{\kappa_{ab}^*}$ and the nonzero columns in $(\widehat{f_{in} \uparrow^G})_b$ have the same indices;
3. the nonzero columns in $\widehat{\kappa_{ab}^*}$ and $(\widehat{f_{out} \uparrow^G})_a$ have the same indices, and $\widehat{\kappa_{abmn}^*} = \widehat{\kappa_{(a+1)b_m(n+1)}^*}$ for any a, b, m, n .

such that $\int_G \kappa(\nu^{-1}g)(f_{in} \uparrow^G)(\nu) = \int_G \kappa^*(\nu^{-1}g)(f_{in} \uparrow^G)(\nu)$. Meanwhile, from the Fourier domain, we know that when $\kappa_1^* \neq \kappa_2^*$, then $\int_G \kappa_1^*(\nu^{-1}g)(f_{in} \uparrow^G)(\nu) \neq \int_G \kappa_2^*(\nu^{-1}g)(f_{in} \uparrow^G)(\nu)$. Therefore we find a complete characterization of the equivariant linear map $\{\kappa^*\}$.

In spatial domain, since $\int \kappa_{st}(g)\overline{U_{mn}(g)}dg = \int \kappa_{ts}^\top(g^{-1})\overline{U_{nm}(g)}dg$, we know that $f(g) = \kappa^\top(g^{-1})$ is also a Mackey function, satisfying $f(gh) = \rho_{in}(h^{-1})f(g)$. Therefore, we have $\kappa^*(hg) = \kappa^*(g)\rho_{in}(h^{-1})$. This is equivalent to the space of kernels such that $\kappa^*(h_1gh_2) = \rho(h_2^{-1})\kappa^*\rho(h_1^{-1})$, from (Weiler & Cesa, 2019). (Weiler & Cesa, 2019) proves that these are all equivariant kernels, which finishes our proof

C Details and Proofs for the Examples

C.1 $SE(3)$

Consider the group $SE(3)$, with the stabilizer subgroup $SO(3)$ and the homogeneous space \mathbb{R}^3 . For any element $(x, R) \in SE(3)$, where $x \in \mathbb{R}^3$ and $R \in SO(3)$, the unitary irreducible representations of $SE(3)$ have the form, for $p \in \mathbb{R}^+$ and an integer s ,

$$U_{l', m'; l, m}^s(x, R; p) = \sum_{j=-l}^l [l', m' | p, s | l, j](x) \tilde{U}_{jm}^l(R), \quad (5)$$

where

$$[l', m' | p, s | l, j](x) = (4\pi)^{\frac{1}{2}} \sum_{k=|l'-l|}^{l'+l} i^k \sqrt{\frac{(2l'+1)(2k+1)}{(2l+1)}} J_k(pa),$$

$$\mathcal{C}(k, 0; l', s | l, s) \mathcal{C}(k, m - m'; l', m' | l, m) Y_k^{m-m'}(\theta, \phi),$$

and $\tilde{U}_{mn}^l = (-1)^{n-m} D_{mn}^l$, where D_{mn}^l is the l -th order Wigner D -matrix—an unitary irreducible representation of $SO(3)$, \mathcal{C} is a Clebsch-Gordan coefficient, Y_k is a k -th order spherical harmonic, J_k is the k -th order spherical Bessel function, a is the length of x , and θ, ϕ are the spherical coordinates of the unit vector \hat{x} .

The lifting process takes the form

$$f \uparrow^G (g) = f \uparrow^G (x, R) = D^l(R^\top) f(x).$$

The Fourier Transform $(f \uparrow^G (x, R))_{l, m; l', m'}^s(p)$ (see App. C.3 for details) is

$$\sum_j \int_{\mathbb{R}^3} (-1)^{-k-j} f_{-j}(x, I) \overline{U_{l, m; l', j}^s(x, I; p)} dx \delta(l' - t) \delta(k + m'),$$

where I is the identity matrix, the identity element in $SO(3)$.

The convolution on $SE(3)$ takes the form

$$(\kappa * l)_i(g) = \sum_j \int_{\nu \in G} \kappa_{ij}(\nu^{-1}g) l_j(\nu) d\nu.$$

The Fourier transform $\mathcal{F}((\kappa * l)_i)_{l',m';l,m}^s(p)$ becomes

$$\sum_j \sum_{a=|s|}^{\infty} \sum_{b=-a}^a (\widehat{l}_j)_{l',m';a,b}^s(p) (\widehat{\kappa}_{ij})_{a,b;l,m}^s(p).$$

Assume that the order of the input field is l_1 , and the order of the output field is l_2 . Then κ_{ij} only has nonzero Fourier coefficients at $\widehat{\kappa}_{ij}^s_{l_1,-j;l_2,-i}(p)$. Meanwhile, we can prove that $\widehat{\kappa}_{ij}^s_{l_1,-j;l_2,-i}(p)$ has the form $(-1)^{j-i}c$, where c is a constant over i, j , for direct convolution on \mathbb{R}^3 (see AppC.4). Then we can project κ to the homogeneous space \mathbb{R}^3 , obtaining $\kappa'_{ij}(x) = \kappa_{ij}(x, I)$ as

$$\frac{1}{2\pi^2} \sum_{s=-\min(l_1,l_2)}^{\min(l_1,l_2)} \int_0^{\infty} (-1)^{j-i} c^s(p) U_{l_1,-j;l_2,-i}^s(x, I; p) p^2 dp = \sum_{m=-t}^t \sum_{t=|l_1-l_2|}^{l_1+l_2} C_t(\|x\|) C_{t,m,l_1,j}^{l_2,i} \overline{Y}_t^m.$$

We see that κ is in the span of a fixed basis, and the weight function C_t of $\|x\|$ is arbitrary and thus learnable. This result is consistent with the form of tensor field networks in Thomas et al. (2018) and 3D steerable CNNs in Weiler et al. (2018).




Figure 6. Fourier matrix multiplication for SE(3)

C.2 Details for $SO(3)$

The Fourier Transform becomes:

$$\begin{aligned} \widehat{f}_{mn}^{\uparrow G} &= \int_{R \in SO(3)} f_{\uparrow G}^R (R) \overline{D}_{mn}^l(R) dR = \int_{(\alpha,\beta,\gamma) \in SO(3)} e^{-ik\gamma} f(\alpha, \beta) e^{im\alpha} d_{mn}^l(\beta) e^{in\gamma} d\alpha \sin(\beta) d\beta d\gamma \\ &= \int_{(\alpha,\beta) \in \mathbb{S}^2} f(\alpha, \beta) e^{im\alpha} d_{mn}^l(\beta) d\alpha \sin(\beta) d\beta \int_{\gamma \in SO(2)} e^{-ik\gamma} e^{in\gamma} d\gamma. \end{aligned}$$

This further equals

$$\int_{(\alpha,\beta) \in \mathbb{S}^2} f(\alpha, \beta) e^{im\alpha} d_{mn}^l(\beta) d\alpha \sin(\beta) d\beta \cdot 2\pi \delta(k - n).$$

Projecting $f_{out}^{\uparrow G}$ to f_{out} takes the form

$$\begin{aligned} f_{out}(\alpha, \beta) &= f_{out}^{\uparrow G}(\alpha, \beta, 0) = \int_G \kappa(R_g^{-1} R_{\alpha,\beta,0}) f_{in}^{\uparrow G}(g) dg \\ &= \int_G \kappa(R_{\alpha',\beta',\gamma'}^{-1} R_{\alpha,\beta,0}) f_{in}^{\uparrow G}(\alpha', \beta', \gamma') d\alpha' \sin(\beta') d\beta' d\gamma' \\ &= \int_G \kappa(R_z(-\gamma') R_y(-\beta') R_z(-\alpha') R_z(\alpha) R_y(\beta)) e^{-im\gamma'} f_{in}(\alpha', \beta') d\alpha' \sin(\beta') d\beta' d\gamma' \\ &= \int_G e^{im_1\gamma'} \kappa(R_y(-\beta') R_z(-\alpha') R_z(\alpha) R_y(\beta)) e^{-im_1\gamma'} f_{in}(\alpha', \beta') d\alpha' \sin(\beta') d\beta' d\gamma' \\ &= C \int_{\mathbb{S}^2} \kappa(R_y(-\beta') R_z(-\alpha') R_z(\alpha) R_y(\beta)) f_{in}(\alpha', \beta') d\alpha' \sin(\beta') d\beta' \\ &= C \int_{\mathbb{S}^2} e^{-im_2 h(R^{-1}(\alpha', \beta', 0) R(\alpha, \beta, 0))} \kappa'(R_y(-\beta') R_z(-\alpha') x_{\alpha,\beta}) f_{in}(\alpha', \beta') d\alpha' \sin(\beta') d\beta'. \end{aligned}$$

This completes the details for $SO(3)$.

C.3 Fourier Transform over $SE(3)$

The Fourier Transform has the form

$$\begin{aligned}
 (f \uparrow^G \widehat{(x, R)_k})_{l, m; l', m'}^s(p) &= \int_{SE(3)} f \uparrow^G (x, R)_k(x, R) \overline{U_{l, m; l', m'}^s(x, R; p)} dR dx \\
 &= \sum_{g, l', j} \int_{SE(3)} D_{kg}^\top(R^\top) f_g(x, I) \overline{U_{l, m; l', j}^s(x, I, p) \tilde{U}_{j m'}^{l'}(R)} dR dx \\
 &= \sum_{g, l', j} \int_{\mathbb{R}^3} f_g(x, I) \overline{U_{l, m; l', j}^s(x, I; p)} dx \int_{SO(3)} (-1)^{m' - j} \overline{D_{gk}^\top(R) D_{j m'}^{l'}(R)} dR.
 \end{aligned}$$

Due to the fact that Clebsch–Gordan coefficients $C_{l_1, m_1, l_2, m_2}^{00} = \delta(l_1 - l_2) \delta(m_1 + m_2)$, the decomposition of $D_{gk}^\top(R) D_{j m'}^{l'}(R)$ includes the trivial representation only when $t = l'$, $g = -j$ and $k = -m'$. Therefore, the Fourier matrix has sparsity pattern

$$(f \uparrow^G \widehat{(x, R)_k})_{l, m; l', m'}^s(p) = \sum_j \int_{\mathbb{R}^3} (-1)^{-k-j} f_{-j}(x, I) \overline{U_{l, m; l', j}^s(x, I; p)} dx \delta(l' - t) \delta(k + m').$$

C.4 Kernel Property for $SE(3)$

From the Fourier transform, we know that $\widehat{f_{out} \uparrow_i^G}_{l', m'; l_1, -i}(p)$ has the form $(-1)^{-i} C$, where C does not depend on i , and $\widehat{f_{in} \uparrow_j^G}_{l', m'; l_1, -j}(p)$ has the form $(-1)^{-j} C'$, where C' does not depend on j . Thus, $\sum_j \widehat{\kappa_{ij}}_{l_1, -j; l_2, -i}^s(p)$ has the form $(-1)^{-i} C_j$, where C_j does not depend on i . Therefore, we find the kernel κ_{ij} whose Fourier matrix is shown in Figure 6, and $\kappa_{ij}(g)$ is

$$\frac{1}{2\pi^2} \sum_{s=-\min(l_1, l_2)}^{\min(l_1, l_2)} \int_0^\infty (-1)^{-i} c_{l_1, -j; l_2, -i}^s(p) U_{l_1, -j; l_2, -i}^s(x, R; p)$$

where $\sum_j c_{l_1, -j; l_2, -i}^s$ does not depend on i .

To directly apply the convolution on a homogeneous space acted on by a semidirect product, $c_{l_1, -j; l_2, -i}^s(p)$ should have the form $(-1)^{j-i} c$, where c does not depend on i or j . This is because κ and the function $p : G \rightarrow V$, where $p(g^{-1}) = \kappa^\top(g)$ for any $g \in G$, should both be Mackey functions.

D Experiment Details

D.1 Details of Spherical U-Net in Sec. 4.1

The input field to the U-Net has one order-0 feature and the output of the U-Net is one order-1 feature. The overall network has six layers of widths [32, 16, 16, 16, 16, 32]. The hidden features have [8, 12, 16, 12, 8] channels, respectively, for each type of field. The types in hidden features are order-0 and order-1. The loss is the same as (Estevés et al., 2020), the mean-squared error weighted by the spherical map sampling area. The whole network is trained end-to-end from scratch using the Adam optimizer with an initial learning rate of 1×10^{-3} . The learning rate decays by a factor of 0.2 at epochs 10 and 15.

D.2 Nonlinearity in the $SE(3)$ experiment (Sec. 4.2)

When we only take the -1 -st, 0 -th and 1 -st element of the vector-valued function l on $SE(3)$, the 0 -th element $l_0(*, x)$, of l is a Mackey function on $SO(3)$ at every point x . This is shown in the section 3.4. We observe that $x \mapsto \begin{pmatrix} l_{-1}(*, x) \\ l_1(*, x) \end{pmatrix}$ is a Mackey function on $SO(3)$ corresponding to the irreducible representation $\theta \mapsto \begin{pmatrix} \cos(\theta) & -\sin(\theta) \\ \sin(\theta) & \cos(\theta) \end{pmatrix}$ of $SO(2)$, using

the real form of the Wigner-D matrix. Therefore we can project $x \mapsto \begin{pmatrix} l_{-1}(R, x) \\ l_1(R, x) \end{pmatrix}$ to the sphere by taking $\begin{pmatrix} l_{-1}(r, x) \\ l_1(r, x) \end{pmatrix}$ and using its norm. For $l_0(R, x)$, we only need to take $l_0(r, x)$. Through the above method, we can project the signal from $SO(3)$ to the sphere and keep the equivariance.

D.3 Details of Molecular Completion in Sec. 4.2.1

Our network contains five convolution activation layers using 0, 1, 2-order fields as hidden features with corresponding hidden channel dimensions [32, 32, 32, 32]. In the last layer, we output six scalars for every existing atom (one for probability and five for one-hot atom type prediction for the missing atom) and one vector (relative position to the missing atom). The losses are the same as described in (Thomas et al., 2018) and the network is trained using the Adam optimizer with the initial learning rate 1×10^{-3} . The learning rate is decreased by a factor 0.3 at epoch 2500.

D.4 Details of Shape Classification in Sec. 4.2.2

Our network follows (Poulenard & Guibas, 2021; Poulenard et al., 2019) to perform PCA aligned KD-Tree pooling. It contains six convolution activation layers before global pooling, with KD-Tree pooling depth factors of [0, 2, 0, 2, 0, 2], leading to [1024, 1024, 256, 256, 64, 64, 16] points starting from the input point cloud. Global pooling is applied to the function lifted to the group as in (Poulenard & Guibas, 2021). Since in this task the translation equivariance can be trivially addressed via subtracting the center of mass, following (Poulenard & Guibas, 2021; Dym & Maron, 2020) we concatenate the global xyz coordinates as an order-1 feature to each layer’s input. We use the standard cross entropy loss. The network is trained via the Adam optimizer with a starting learning rate of 1×10^{-3} . The learning rate is decayed by a factor of 0.3 at epochs [100, 150, 200, 250, 300].

D.5 Table 3 with references

| Methods | z/z | $z/SO(3)$ | $SO(3)/SO(3)$ |
|---------------------------------------|-------------|-------------|---------------|
| Spherical-CNN (Esteves et al., 2018) | 88.9 | 76.7 | 86.9 |
| a^3 S-CNN (Liu et al., 2018) | 89.6 | 87.9 | 88.7 |
| SFCNN (Rao et al., 2019) | 91.4 | 84.8 | 90.1 |
| TFN (Thomas et al., 2018) | 88.5 | 85.3 | 87.6 |
| RI-Conv (Zhang et al., 2019) | 86.5 | 86.4 | 86.4 |
| SPHNet (Poulenard et al., 2019) | 87.7 | 86.6 | 87.6 |
| ClusterNet (Chen et al., 2019) | 87.1 | 87.1 | 87.1 |
| GC-Conv (Zhang et al., 2020) | 89.0 | 89.1 | 89.2 |
| RI-Framework (Li et al., 2020) | 89.4 | 89.4 | 89.3 |
| VN-PointNet (Deng et al., 2021) | 77.5 | 77.5 | 77.2 |
| VN-DGCNN (Deng et al., 2021) | 89.5 | 89.5 | 90.2 |
| TFN[mlp]-P (Poulenard & Guibas, 2021) | 89.7 | 89.7 | 89.7 |
| Ours | 89.7 | 89.7 | 89.3 |

Table 4. Due to the page limit, we provide the references for the methods in Table 3 here.