

FAST, EXPRESSIVE $SE(n)$ EQUIVARIANT NETWORKS THROUGH WEIGHT-SHARING IN POSITION-ORIENTATION SPACE

Erik J. Bekkers^{1*}, Sharvaree Vadgama^{1*}

Rob D. Hesselink¹, Putri A. van der Linden¹, David W. Romero^{2,3}

¹ University of Amsterdam ² Vrije Universiteit Amsterdam ³ NVIDIA Research

*Equal contribution. Code available at <https://github.com/ebekkers/ponita>

ABSTRACT

Based on the theory of homogeneous spaces we derive *geometrically optimal edge attributes* to be used within the flexible message-passing framework. We formalize the notion of weight sharing in convolutional networks as the sharing of message functions over point-pairs that should be treated equally. We define equivalence classes of point-pairs that are identical up to a transformation in the group and derive attributes that uniquely identify these classes. Weight sharing is then obtained by conditioning message functions on these attributes. As an application of the theory, we develop an efficient equivariant group convolutional network for processing 3D point clouds. The theory of homogeneous spaces tells us how to do group convolutions with feature maps over the homogeneous space of positions \mathbb{R}^3 , position and orientations $\mathbb{R}^3 \times S^2$, and the group $SE(3)$ itself. Among these, $\mathbb{R}^3 \times S^2$ is an optimal choice due to the ability to represent directional information, which \mathbb{R}^3 methods cannot, and it significantly enhances computational efficiency compared to indexing features on the full $SE(3)$ group. We support this claim with state-of-the-art results—in accuracy and speed—on five different benchmarks in 2D and 3D, including interatomic potential energy prediction, trajectory forecasting in N-body systems, and generating molecules via equivariant diffusion models.

1 INTRODUCTION

Inspired by the foundational role of convolution operators in deep learning and the ‘convolution is all you need’ theorem (Cohen et al. (2019, Thm 3.1); Bekkers (2019, Thm 1)) -which asserts that any layer that is linear and equivariant must be a group convolution, we propose an efficient and expressive group convolutional approach for constructing neural networks equivariant to $SE(n)$: the group of n -dimensional translations and rotations. While this theorem is a theoretical result, several studies provide empirical truth to the statement as well. For example, ConvNeXt (Liu et al., 2022b) challenges the need for Transformers (Vaswani et al., 2017) in vision tasks, and Romero et al. (2021); Poli et al. (2023) show that convolutions are sufficient to model long context in sequences, e.g., language, without the need for transformers or recurrent networks.

In our work, we revisit the influential inductive bias of *weight sharing* in convolutions (LeCun et al., 1998), classically defined as the sharing of a convolution kernel (linear transformation) over all the neighborhoods in an image. In the discrete image setting, the kernel is given as a set of weights and it is appropriate to refer to the convolution kernel as *the weights*. However, in the continuous case, and more general processing frameworks such as in message passing networks (MPNs) (Gilmer et al., 2017a), this terminology no longer literally applies. Therefore, we formalize the notion of weight sharing to gain insights into how to use this inductive bias in a more general setting.

Traditionally, weight sharing refers to the idea of using the same linear transformation matrix over neighborhoods identical up to translation. We generalize this notion by the construction of equivalence classes of neighboring point-pairs, in which we say neighbor pairs are equivalent if they are identical up to a transformation in a group G . Weight sharing then becomes the notion of sharing message functions in MPNs over the equivalence classes, which is achieved by conditioning the functions on *attributes* that act as identifiers for the equivalence classes. In Sec. 3.1 we formalize this construction and present all the attributes one needs for $SE(n)$ equivariant weight sharing.

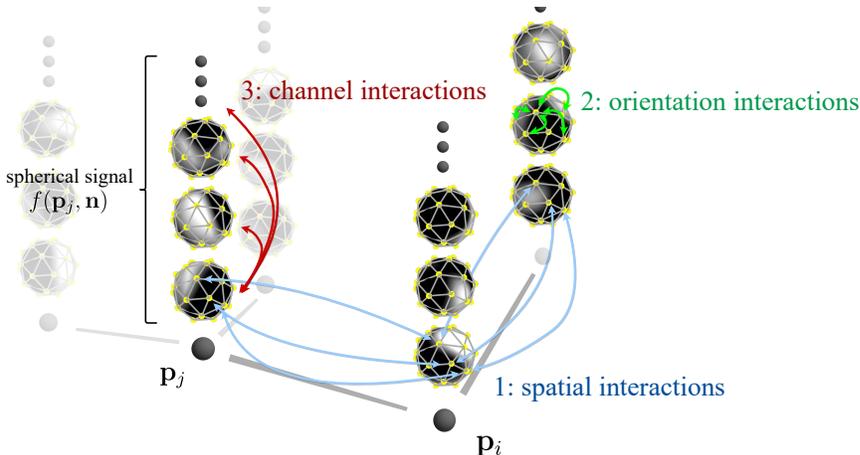


Figure 1: Separable group convolutions on position orientation space $\mathbb{R}^3 \times S^2$. Efficiency is obtained due to parallelizing the most expensive step, **step 1 (message passing)**, over orientations and channels. **Steps 2 and 3** are efficient as well as **spherical convolutions** are batched over positions and channels, and **channel mixing** is batched over positions and orientations.

As an application of the theory, we construct expressive and efficient $SE(n)$ equivariant group convolutional networks for the processing of 3D point clouds. Our simple, fully convolutional architecture achieves state-of-the-art results on three different benchmarks: interatomic potential energy prediction, trajectory forecasting in N-body systems, and generating molecules via equivariant diffusion models. The contributions of this paper are summarized as follows:

- We formalize the notion of weight sharing in Sec. 3.1.
- We derive optimal attributes (Thm. 1) for weight sharing over equivalence classes of point pairs and show that these *are all you need* to build equivariant universal approximators (Cor. 1.1).
- We present a fast, expressive equivariant architecture based purely on convolutions (Fig. 1). Our method reaches state-of-the-art performance on three equivariant benchmarks.

2 BACKGROUND

2.1 MATHEMATICAL PREREQUISITES

Groups and homogeneous spaces. A *group* is an algebraic construction defined by a set G and a binary operator $\cdot : G \times G \rightarrow G$, known as the *group product*. This structure must satisfy the following axioms: *closure*, where $\forall h, g \in G : hg \in G$; the existence of both an *identity* e and an *inverse* g^{-1} element such that $g^{-1} \cdot g = e$; and *associativity*, where $\forall g, h, i \in G : (g \cdot h) \cdot i = g \cdot (h \cdot i)$. We denote the group product between two elements $g, g' \in G$ by juxtaposition, i.e., as gg' .

It is useful to think of the elements of G as transformations. The group product then tells how a transformation g' followed by another transformation g can be represented by a single transformation gg' . We focus on the *Special Euclidean motion group* $SE(n)$ consisting of distance and orientation-preserving transformations. Elements $g = (\mathbf{x}, \mathbf{R}) \in SE(n)$ are parameterized by translation vectors $\mathbf{x} \in \mathbb{R}^n$ and rotation matrices $\mathbf{R} \in SO(n)$. Here, $SO(n)$ denotes the set of $n \times n$ matrices with determinant 1, which forms a group in itself with matrix multiplication as a group product. The $SE(n)$ group product between two roto-translations $g = (\mathbf{x}, \mathbf{R})$ and $g' = (\mathbf{x}', \mathbf{R}')$ is given by $(\mathbf{x}, \mathbf{R})(\mathbf{x}', \mathbf{R}') = (\mathbf{R}\mathbf{x}' + \mathbf{x}, \mathbf{R}\mathbf{R}')$, and its identity element is given by $e = (\mathbf{0}, \mathbf{I})$.

A group can act on spaces other than itself via a *group action* $\mathcal{T} : G \times X \rightarrow X$, where X is the space on which G acts. For simplicity, we denote the action of $g \in G$ on $x \in X$ as gx . Such a transformation is called a group action if it is homomorphic to G and its group product. That is, it follows the group structure: $(gg')x = g(g'x) \forall g, g' \in G, x \in X$, and $ex = x$. For example, consider the space of 3D positions $X = \mathbb{R}^3$, e.g., atomic coordinates, acted upon by the group $G = SE(3)$. A position $\mathbf{p} \in \mathbb{R}^3$ is roto-translated by the action of an element $(\mathbf{x}, \mathbf{R}) \in SE(3)$ as $(\mathbf{x}, \mathbf{R})\mathbf{p} = \mathbf{R}\mathbf{p} + \mathbf{x}$.

A group action is termed *transitive* if every element $x \in X$ can be reached from an arbitrary origin $x_0 \in X$ through the action of some $g \in G$, i.e., $x = gx_0$. A space X equipped with a transitive action of G is called a *homogeneous space* of G . Finally, the *orbit* $Gx := \{gx \mid g \in G\}$ of an element

x under the action of a group G represents the set of all possible transformations of x by G . For homogeneous spaces, $X=Gx_0$ for any arbitrary origin $x_0 \in X$.

Quotient spaces. The aforementioned space of 3D positions $X=\mathbb{R}^3$ serves as a homogeneous space of $G = \text{SE}(3)$, as every element \mathbf{p} can be reached by a roto-translation from $\mathbf{0}$, i.e., for every \mathbf{p} there exists a (\mathbf{x}, \mathbf{R}) such that $\mathbf{p}=(\mathbf{x}, \mathbf{R}) \mathbf{0}=\mathbf{R} \mathbf{0} + \mathbf{x}=\mathbf{x}$. Note that there are several elements in $\text{SE}(3)$ that transport the origin $\mathbf{0}$ to \mathbf{p} , as any action with a translation vector $\mathbf{x}=\mathbf{p}$ suffices regardless of the rotation \mathbf{R} . This is because any rotation $\mathbf{R}' \in \text{SO}(3)$ leaves the origin unaltered.

We denote the set of all elements in G that leave an origin $x_0 \in X$ unaltered the *stabilizer subgroup* $\text{Stab}_G(x_0)$. In subsequent analyses, we use the symbol H to denote the stabilizer subgroup of a chosen origin x_0 in a homogeneous space, i.e., $H=\text{Stab}_G(x_0)$. We further denote the *left coset* of H in G as $gH := \{gh \mid h \in H\}$. In the example of positions $\mathbf{p} \in X=\mathbb{R}^3$ we concluded that we can associate a point \mathbf{p} with many group elements $g \in \text{SE}(3)$ that satisfy $\mathbf{p}=g \mathbf{0}$. In general, letting g_x be any group element s.t. $x=g_x x_0$, then any group element in the left set $g_x H$ is also identified with the point \mathbf{p} . Hence, any $x \in X$ can be identified with a left coset $g_x H$ and vice versa.

Left cosets gH then establish an *equivalence relation* \sim among transformations in G . We say that two elements $g, g' \in G$ are equivalent, i.e., $g \sim g'$, if and only if $g x_0=g' x_0$. That is, if they belong to the same coset gH . The space of left cosets is commonly referred to as the *quotient space* G/H .

We consider *feature maps* $f : X \rightarrow \mathbb{R}^C$ as multi-channel signals over homogeneous spaces X . Such maps are of interest as they often form the hidden representations in various deep learning tasks. In this work, we treat point clouds as sparse feature maps, e.g., sampled only at atomic positions. In the general continuous setting, we denote the space of feature maps over X with \mathcal{X} . Such feature maps undergo group transformations through *regular group representations* $\rho^X(g) : \mathcal{X} \rightarrow \mathcal{X}$ parameterized by g , and which transform functions $f \in \mathcal{X}$ via $[\rho^X(g)f](x)=f(g^{-1}x)$.

2.2 MOTIVATION 1: GROUP CONVOLUTION IS ALL YOU NEED

In deep learning, we often employ learnable operators $\Phi : \mathcal{X} \rightarrow \mathcal{Y}$, such as self-attention or convolution layers, to iteratively transform feature maps. Such an operator is termed *G-equivariant* if it commutes with group representations on the input and output feature maps: $\rho^Y(g) \circ \Phi=\Phi \circ \rho^X(g)$. Group equivariance ensures that operators preserve the geometric structure of the data, meaning that derived features follow the same transformation laws as the input.

An important result in the field of equivariant deep learning and signal processing is that if we want Φ to be linear and group equivariant, then it *must be a group convolution*:

$$[\Phi f](y)=\int_X k(g_y^{-1}x)f(x)dx. \quad (1)$$

Essentially, the group convolution performs template matching of a kernel k against patterns in f by taking \mathbb{L}_2 -inner products of the shifted kernel $k(g_y^{-1}\cdot)$ and f . Recall that g_y denotes any group element such that $y=g_y y_0$, and any other group element in the set $g_y H$ is also valid. This implies that Eq. 1 is only valid if the kernel is invariant to left actions of H , i.e. if $\forall_{h \in H} : k(h^{-1}x)=k(x)$. These findings are summarized in various seminal equivariant deep learning works (Cohen et al., 2019, Thm. 3.1 *convolution is all you need*), (Bekkers, 2019, Thm 1), (Kondor & Trivedi, 2018).

In light of the ‘‘convolution is all you need’’ claim and the pivotal role of convolutions in areas like computer vision and signal processing, one might question the need for anything more complex than convolutions. We therefore explore the potential of a straightforward, yet theoretically grounded, fully convolutional approach to equivariant deep learning.

2.3 MOTIVATION 2: EFFICIENCY AND EXPRESSIVITY –THE HOMOGENEOUS SPACE $\mathbb{R}^3 \times S^2$

It is worth noting that the domain of the output signal Y in Eq. 1 does not have to match the domain of the input signal X . The aforementioned theorems show that $\text{SE}(3)$ equivariant convolutions on feature maps with domain $\mathbb{R}^3=\text{SE}(3)/\text{SO}(3)$ require isotropic, rotation invariant kernels, as seen in SchNet (Schütt et al., 2023). However, maximal expressivity is gained when the kernel has no constraints. This is achieved when the domain Y is the group itself, i.e., $\text{SE}(3)$ as $H=\{e\}$ is then the trivial group. Generating higher-dimensional feature maps, from $X=\mathbb{R}^3$ to $Y=\text{SE}(3)$, is typically referred to as *lifting*. It is done by performing template matching (Eq. 1) over both translations and rotations, as opposed to just translations. Despite enhanced expressivity, however, subsequent layers must compute integrals over the whole $\text{SE}(3)$ space, which can be computationally restrictive.

In this work, we opt for a middle ground and define feature maps on the **homogeneous space of positions and orientations** $X = \mathbb{R}^3 \times S^2 \cong \text{SE}(3)/\text{SO}(2)$. We denote elements in $\mathbb{R}^3 \times S^2$ as tuples (\mathbf{p}, \mathbf{o}) of positions $\mathbf{p} \in \mathbb{R}^3$ and orientations $\mathbf{o} \in S^2$. We set the origin as $x_0 = (\mathbf{0}, \mathbf{e}_z)$, where $\mathbf{e}_z = (0, 0, 1)^\top$. Every orientation can be obtained from a rotation $\mathbf{R}_\mathbf{o} \in \text{SO}(3)$ that maps the unit vector \mathbf{e}_z to \mathbf{o} . Hence, S^2 is a homogeneous space of $\text{SO}(3)$ with stabilizer subgroup $H = \text{Stab}_{\text{SO}(3)}(\mathbf{e}_z) \cong \text{SO}(2)$: the group of rotations around the z -axis. Therefore, each position-orientation tuple corresponds to an equivalence class of transformations in $\text{SE}(3)$, formalized as $(\mathbf{p}, \mathbf{o}) \equiv (\mathbf{p}, \mathbf{R}_\mathbf{o}) H$ in $\text{SE}(3)/\text{SO}(2)$.

While methods based on \mathbb{R}^3 positions are computationally efficient, intermediate feature maps may fail to capture directional information due to kernel constraints. On the other hand, full $\text{SE}(3)$ group convolutions are computationally demanding but excel at capturing directional features. Compared to the latter, our method offers a significant improvement in computational efficiency *without compromising the capability to represent directional features*. This claim is proven in (Gasteiger et al., 2021, Thm. 3) by relying on the fact that equivariant predictions can be obtained by combining an equivariant basis with coefficients obtained from an invariant network (Villar et al., 2021). In DimeNet and GemNet (Gasteiger et al., 2019; 2021), the basis is determined by the direction of edges in an atomic point cloud, and coefficients are derived through *invariant message passing*. Rather than assigning a sparse basis of reference directions in S^2 to each point position, our method employs a *dense basis* and assign a *spherical signal* to each point in \mathbb{R}^3 , see Fig. 1.

2.4 MOTIVATION 3: CONDITIONAL MESSAGE PASSING AS GENERALIZED CONVOLUTION

Message passing. Our primary focus lies in the processing of point clouds on homogeneous spaces. Point cloud methods are conventionally discussed within the broader context of message passing (Gilmer et al., 2017a). We intentionally adopt this framework as it serves as a familiar and versatile paradigm suitable for the description of various deep learning layers, including convolutions.

We model point-clouds as graphs $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, with nodes $i \in \mathcal{V}$ and edges $(i, j) \in \mathcal{E}$. Each node i has an associated coordinate $x_i \in X$ in a homogeneous space X . We consider features over such graphs as sparse discretizations of dense feature maps $f : X \rightarrow \mathbb{R}^C$ with node features $(x_i, f_i) \Leftrightarrow f(x_i)$. Considering optional pair-wise edge **attributes** a_{ij} , message passing layers update feature nodes as:

1. compute the message m_{ij} from node j to i :
$$m_{ij} = \phi_m(f_i, f_j, a_{ij}), \quad (2)$$

2. Perform a permutation invariant message aggregation:
$$m_i = \sum_{j \in \mathcal{N}(i)} m_{ij}, \quad (3)$$

3. Update the node features:
$$f_i^{\text{out}} = \phi_f(f_i, m_i). \quad (4)$$

Here, $\mathcal{N}(i) = \{j \mid (i, j) \in \mathcal{E}\}$ denotes the set of neighbors of the node i .

Convolutional message passing. Convolutions can be described in the message passing framework as approximations over a sparse set of points on which continuous signals are sampled:

$$\int_X k(g_x^{-1}x')f(x')dx' \approx \sum_{j \in \mathcal{N}(i)} k(g_{x_i}^{-1}x_j)f_j. \quad (5)$$

Here, we can recognize the message passing form, in which $\phi_m(f_j, g_{x_i}^{-1}x_j) = k(g_{x_i}^{-1}x_j)f_j$ is a linear transformation performed by matrix-vector multiplication with the kernel $k(g_{x_i}^{-1}x_j) \in \mathbb{R}^{C_{\text{out}} \times C}$ that depends on $g_{x_i}^{-1}x_j$. Following Brandstetter et al. (2021), we interpret this as using a message function that is *conditioned on the attribute* a_{ij} , denoted as $\phi_m(f_i, f_j; a_{ij})$ to emphasize this dependency. We could directly consider the use of the **invariant attribute** $a_{ij} = g_{x_j}^{-1}x_i$ in conditional message passing as an extension of convolution. However, it's crucial to recognize that, despite its invariance under global group actions, the value of this attribute depends on the selected representative g_{x_i} . Specifically, any $g_{x_i} h$, where h belongs to the stabilizing group H , can result in a different $a_{ij} = h^{-1}g_{x_i}^{-1}x_j$. In essence, any $a_{ij} \in H g_{x_i} x_j$ within the orbit of H should be treated equivalently.

3 FAST, EXPRESSIVE $\text{SE}(n)$ EQUIVARIANT NEURAL NETWORKS

3.1 WEIGHT-SHARING AND OPTIMAL INVARIANT ATTRIBUTES

Problem statement. Our objective is to find an **invariant attribute** a_{ij} that can be associated with two points (x_i, x_j) in a homogeneous space X of a group G , such that it satisfies these criteria:

- (i) *Invariance to the global action of G* . Any pair in the equivalence class $[x_i, x_j] := \{g x_i, g x_j \mid g \in G\}$ must be mapped to same attribute a_{ij} .
- (ii) *Uniqueness*. Each attribute a_{ij} should be unique for the given equivalence class.

This problem boils down to finding a *bijective* map $[x_i, x_j] \mapsto a_{ij}$. Note that invariance alone is insufficient to build an expressive network. For example, a mapping $[x_i, x_j] \mapsto 0$ is invariant but trivial, and thus not useful. Bijectivity ensures that each attribute is unique and it fully characterizes the space of all possible equivalence classes of point pairs. In essence, a bijective attribute mapping is all you need to enable weight-sharing over equivalent point pairs and obtain full expressiveness. The notion of weight sharing is formalized by the following three definitions:

Definition 3.1 (*Equivalent point pairs*). Two point pairs $(x_i, x_j), (x'_i, x'_j) \in X \times X$ in a homogeneous space X of a group G are called *equivalent* iff they can be transformed into each other through left multiplication by an element $g \in G$. We define this equivalence relation as:

$$(x_i, x_j) \sim (x'_i, x'_j) \iff \exists g \in G : (x'_i, x'_j) = (g x_i, g x_j).$$

Definition 3.2 (*Equivalence class of point pairs*). The equivalence relation \sim of definition 3.1 defines an *equivalence class* of point pairs, denoted as:

$$[x_i, x_j] = \{(x'_i, x'_j) \in X \times X \mid (x'_i, x'_j) \sim (x_i, x_j)\}.$$

Here, $[x_i, x_j]$ represents a set of point pairs that should be treated as equivalent, with (x_i, x_j) serving as its *representative*. The *space of equivalence classes* of point pairs is denoted as $X \times X / \sim$.

Definition 3.3 (*Weight-sharing in message passing*). A message passing layer (Eq. 2-4) is said to *share weights* over equivalent point pairs if it processes equivalent point pairs in the same way. That is, if its message function $\phi_m(\mathbf{f}_i, \mathbf{f}_j; [x_i, x_j])$ is conditioned on the equivalence class $[x_i, x_j]$.

See (Koishekenov & Bekkers, 2023) for options for conditioning message functions. Our goal now is to parameterize the space of equivalence classes $X \times X / \sim$ with concrete attributes that condition message functions. To that end, we reduce the problem of identifying equivalence classes of point pairs to the task of identifying orbits of single points X , akin to LieConv (Finzi et al., 2020a). We show that equivalence classes $[x_i, x_j]$ correspond to orbits in $H \backslash X$ using the following lemma:

Lemma 1 (*Equivalence class correspond to H -orbits in X*). For any chosen representatives $g_i \in G$ of $x_i \in X \equiv G/H$ such that $x_i = g_i x_0$, and any $x_j \in X$, the following mapping from the space of equivalence classes of point pairs $X \times X / \sim$ to the space $H \backslash X$ of orbits of H in X is a bijection:

$$[x_i, x_j] \mapsto H g_i^{-1} x_j, \quad [x_i, x_j] \in X \times X / \sim, \quad g_i^{-1} x_j \in H \backslash X. \quad (6)$$

Proof. See the appendix, section B.1. □

The subsequent theorem characterizes bijective mappings from equivalence classes $[x_i, x_j]$ to concrete attributes a_{ij} that we can compute for homogeneous spaces of $\text{SE}(n)$. Since these mappings are bijective, the attributes serve as unique identifiers of the equivalence classes. Therefore, they are sufficient for the construction of expressive yet efficient equivariant message passing networks.

Theorem 1 (*Bijective attributes for homogeneous spaces of $\text{SE}(n)$*). Consider three pairs of homogeneous spaces $X \equiv \text{SE}(n)/H$ and stabilizer subgroups H : (i) *position space*: $X = \mathbb{R}^n$, $H = \text{SO}(n)$; (ii) *position-orientation space*: $X = \mathbb{R}^d \times S^{n-1}$, $H = \text{SO}(n-1)$; and (iii) *the whole group*: $X = \text{SE}(n)$, $H = \{e\}$. For $n=2, 3$, the following mappings from equivalence classes $[x_i, x_j] \in X \times X / \sim$, as defined in definitions 3.1 and 3.2, to explicit attributes are bijective:

$$\mathbb{R}^2 \text{ and } \mathbb{R}^3 : \quad [(\mathbf{p}_i, \mathbf{p}_j)] \mapsto a_{ij} = \|\mathbf{p}_j - \mathbf{p}_i\|, \quad (7)$$

$$\mathbb{R}^2 \times S^1 \text{ and } \text{SE}(2) : \quad [(\mathbf{p}_i, \mathbf{o}_i), (\mathbf{p}_j, \mathbf{o}_j)] \mapsto a_{ij} = (\mathbf{R}_{\mathbf{o}_i}^{-1}(\mathbf{p}_j - \mathbf{p}_i), \arccos \mathbf{o}_i^\top \mathbf{o}_j), \quad (8)$$

$$\mathbb{R}^3 \times S^2 : \quad [(\mathbf{p}_i, \mathbf{o}_i), (\mathbf{p}_j, \mathbf{o}_j)] \mapsto a_{ij} = \left(\begin{array}{c} \mathbf{o}_i^\top (\mathbf{p}_j - \mathbf{p}_i) \\ \|\mathbf{p}_j - \mathbf{p}_i\| - \mathbf{o}_i^\top (\mathbf{p}_j - \mathbf{p}_i) \mathbf{o}_i \\ \arccos \mathbf{o}_i^\top \mathbf{o}_j \end{array} \right), \quad (9)$$

$$\text{SE}(3) : \quad [(\mathbf{p}_i, \mathbf{R}_i), (\mathbf{p}_j, \mathbf{R}_j)] \mapsto a_{ij} = (\mathbf{R}_i^{-1}(\mathbf{p}_j - \mathbf{p}_i), \mathbf{R}_i^{-1} \mathbf{R}_j). \quad (10)$$

Proof. A proof is provided in the appendix, section B.2. □

The three components of a_{ij} in Eqs. 8, 9 share the same interpretation: the first two decompose displacement $\mathbf{p}_j - \mathbf{p}_i$ into shifts parallel and orthogonal to \mathbf{o}_i , the third represents the angle between \mathbf{o}_i and \mathbf{o}_j . As for any coordinate system, these mappings are not unique: we can rewrite Eq. 9

in polar coordinates as $a_{ij} = (\|\mathbf{p}_j - \mathbf{p}_i\|, \arccos \mathbf{o}_i^\top (\mathbf{p}_j - \mathbf{p}_i), \arccos \mathbf{o}_i^\top \mathbf{o}_j)^\top$: the invariants used in DimeNet (Gasteiger et al., 2019). Eqs. 8-10 implement $g_i^{-1}g_j$. The following is proven in Appx. B.4

Corollary 1.1. *Message passing networks in geometric graphs over $\mathbb{R}^n \times S^{n-1}$ and $SE(n)$, with message functions conditioned on the attributes of Thm. 1, are equivariant universal approximators.*

3.2 SEPARABLE GROUP CONVOLUTION IN POSITION-ORIENTATION SPACE

Based on the bijective equivalence class embedding from Sec. 3.1, we define a group convolution as

$$f^{out}(\mathbf{p}, \mathbf{o}) = \int_{\mathbb{R}^3} \int_{S^2} k([\mathbf{p}, \mathbf{o}], [\mathbf{p}', \mathbf{o}']) f(\mathbf{p}', \mathbf{o}') d\mathbf{p}' d\mathbf{o}'. \quad (11)$$

Here, $f : \mathbb{R}^3 \times S^2 \rightarrow \mathbb{R}^{C_{in}}$ represents the input feature map with C_{in} channels. the kernel $k([\mathbf{p}, \mathbf{o}], [\mathbf{p}', \mathbf{o}']) \in \mathbb{R}^{C_{out} \times C_{in}}$ returns a linear transformation matrix for each equivalence class of point pairs, resulting in an output signal $f^{out} : \mathbb{R}^3 \times S^2 \rightarrow \mathbb{R}^{C_{out}}$ with C_{out} channels.

We use the fact that the third component of the invariant attribute in Eq. 9 only depends on orientations to separate the $\mathbb{R}^3 \times S^2$ convolution into three parts: spatial convolution, spherical convolution, and channel mixing, as illustrated in figure 1. This approach combines the efficiency strategies of Knigge et al. (2022); Chollet (2017), to separate group convolutions over different group parts, and channel interactions from spatial interactions –generally homogeneous space interactions–, respectively. Following this strategy, we factorize the convolution kernel into three parts:

$$k([\mathbf{p}, \mathbf{o}], [\mathbf{p}', \mathbf{o}']) = K^{(3)} k^{(2)}(\mathbf{o}^\top \mathbf{o}') k^{(1)}(\mathbf{o}^\top (\mathbf{p}' - \mathbf{p}), \|\mathbf{o} \perp (\mathbf{p}' - \mathbf{p})\|), \quad (12)$$

with $k^{(1)} : \mathbb{R} \times \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}^{C_{in}}$ the channel-wise spatial mixing kernel, $k^{(2)} : [0, \pi] \rightarrow \mathbb{R}^{C_{in}}$ the channel-wise orientation mixing kernel, and $K^{(3)} \in \mathbb{R}^{C_{out} \times C_{in}}$ the channel mixing kernel. Substituting the factorized kernel in Eq. 11 allows us to decompose it into three distinct modules:

$$f \longrightarrow \boxed{\text{SpatialGConv}} \longrightarrow \boxed{\text{SphericalGConv}} \longrightarrow \boxed{\text{Linear}} \longrightarrow f^{out}, \quad (13)$$

with $\boxed{\text{Linear}}$ the usual linear layer –implementing $K^{(3)}$ – and

$$\boxed{\text{SpatialGConv}} : f_{i,o}^{(1)} = \sum_{j \in \mathcal{N}(i)} k^{(1)}(\mathbf{o}_o^\top (\mathbf{p}'_j - \mathbf{p}_i), \|\mathbf{o}_o \perp (\mathbf{p}_j - \mathbf{p}_i)\|) \odot f_{j,o} \quad (14)$$

$$\boxed{\text{SphericalGConv}} : f_{i,o}^{(2)} = \sum_{o'=0}^{N-1} k^{(2)}(\mathbf{o}_o^\top \mathbf{o}'_o) \odot f_{i,o'}^{(1)}. \quad (15)$$

Here \odot denotes element-wise multiplication and we turned the integrals of Eq. 11 into discrete sums via Eq. 5. In the discrete implementations, the position-orientation space feature maps f are of dimension $[P, N, C]$, with P the number of points in the point cloud, N the number of orientations on a precomputed spherical grid that is shared over all positions, and with C the number of channels. We use $f_{j,o}$ to denote the feature vector at position \mathbf{p}_j and orientation \mathbf{o}_o . Unlike traditional point-cloud methods, we introduce an additional axis to index features over different orientations.

Each of the operations in Eq. 13 is highly efficient. The most expensive module is SpatialGConv, as it requires aggregating over spatial neighborhoods. We minimize the computation cost of this operation by avoiding channel mixing in this step, essentially performing a depth-wise separable convolution (Chollet, 2017). SphericalGConv is efficient as it can be batched over positions and channels. Moreover, since the spherical grids are shared over all positions, the kernel $k^{(2)}$ can be precomputed, resulting in a kernel of shape $[N, N', C]$. Eq. 15 is then implemented as an einsum over the N' axis of the kernel and input with shape $[P, N', C]$.

While the feature maps in Eq. 14, 15 are discretized, the convolution kernels are still continuous. We parameterize these using Multilayer Perceptrons (MLPs) that take attributes as input. Following standard practices, we first embed the attributes, which we choose to do using a Polynomial Embedding (PE) (cf. Appx. C). That is, we sample the kernel via $a \rightarrow \boxed{\text{PE}} \rightarrow \boxed{\text{MLP}} \rightarrow k$.

3.3 THE PÖNITA ARCHITECTURE

We use the separable $\mathbb{R}^3 \times S^2$ group convolution in a ConvNeXt (Liu et al., 2022b) layer structure:

$$f \downarrow \boxed{\text{SpatialGConv}} \rightarrow \boxed{\text{SphericalGConv}} \rightarrow \boxed{\text{LayerNorm}} \rightarrow \boxed{\text{Linear}} \rightarrow \boxed{\text{Activation}} \rightarrow \boxed{\text{Linear}} \uparrow f^{out}$$

and use it to construct a simple fully convolutional neural network given by:

$$f \rightarrow \boxed{\text{NodeEmbed}} \rightarrow \boxed{\text{ConvNext} \times L} \rightarrow \boxed{\text{Readout}} \rightarrow f^{out} . \quad (16)$$

We call this position-orientation space network based on invariant attributes PΘNITA. Since PΘNITA uses position-orientation space feature maps, we require a method to embed inputs onto spherical signals, and readout outputs from it. We do so via the following node-wise modules for embedding and predicting scalar fields $s : \mathcal{V} \rightarrow \mathbb{R}$ and vector fields $\mathbf{v} : \mathcal{V} \rightarrow \mathbb{R}^3$:

$$\begin{aligned} \boxed{\text{SphereToScalar}} : s_i &= \sum_{o=0}^{N-1} f_{i,o}, & \boxed{\text{ScalarToSphere}} : f_{i,o} &= s_i, \\ \boxed{\text{SphereToVec}} : \mathbf{v}_i &= \sum_{o=0}^{N-1} f_{i,o} \mathbf{o}_o, & \boxed{\text{VecToSphere}} : f_{i,o} &= \mathbf{v}_i^\top \mathbf{o}_o. \end{aligned}$$

The VecToSphere module is essentially a spherical harmonic embedding of frequency one, followed by an inverse spherical Fourier transform: $f_i(\mathbf{o}) = \mathcal{F}_{S^2}^{-1}[Y^{(l=1)}(\mathbf{v})](\mathbf{o})$. This perspective offers a generalized approach for embedding higher-order tensors onto the sphere by using them as coefficients in an inverse Fourier transform, e.g., using the e3nn library (Geiger & Smidt, 2022).

4 RELATED WORK

The theoretical part of our work is inspired by other works on the construction of equivariant neural networks on homogeneous spaces, such as (Cohen et al., 2019; Kondor & Trivedi, 2018; Bekkers, 2019). Each of these papers contains a version of the ‘convolution is all you need’ theorem, which we take as a motivating starting point in this paper (cf. section 2.2). For a comprehensive unifying coverage of equivariant methods, we recommend (Weiler et al., 2023).

Group convolutions are typically defined in *regular* or *steerable / tensor field* form. The regular group convolution viewpoint -which we adopt in this paper- naturally extends convolution as template matching of a kernel over an underlying signal, but now matching all transformation in a group G instead of translation only (Cohen & Welling, 2016; Bekkers, 2019). On gridded data, such as images, regular group convolutions either achieve *exact equivariance* by considering symmetry groups of the grid (Cohen & Welling, 2016; Worrall & Brostow, 2018), or *numerically approximate equivariance* by relying on interpolation (Bekkers et al., 2018; Kuipers & Bekkers, 2023), basis functions (Weiler et al., 2018; Sosnovik et al., 2019), or continuous MLP-based kernels (Finzi et al., 2020a; Knigge et al., 2022). The steerable approach bypasses the need for discretizations by working fields of vectors that transform under (irreducible) representations of the $SO(n)$ group Weiler et al. (2021). Under the ‘convolution is all you need’ paradigm, regular and steerable methods are equivalent and they relate to each other via $SO(n)$ Fourier transforms (Kondor et al., 2018; Cesa et al., 2021). In Appx. A we make precise how to interpret PΘNITA in either the steerable or regular form.

In the field of deep learning-based molecular modeling, in which physical modeling based on Spherical Harmonics, Clebsch-Gordan tensor products, and representation theory of $SO(n)$ is more common, there is a tendency to rely on the tensor field paradigm. Influential contributions in this field are *Tensor Field Network* (Thomas et al., 2018) and *Cormorant* (Anderson et al., 2019). Limitations of such methods are the high computational overhead due to the specialized Clebsch-Gordan tensor products and the need for a solid understanding of representation theory. An advantage of the tensor field approach is that one obtains exact equivariance, whereas the regular group convolution approach requires discretizations of the group, which leads to equivariance only up to the discretization resolution. In this work, however, we show that these numerical errors are not detrimental to performance and opt for simplicity and familiarity with the standard convolutional paradigm.

Another class of equivariant methods is based on message passing with invariant geometric attributes, such as DimeNet (Gasteiger et al., 2019), GemNet (Gasteiger et al., 2021), and SphereNet (Liu et al., 2022a). These are closest to the regular group convolution paradigm and could be framed as non-linear group convolutions (Brandstetter et al., 2021) over the *homogeneous space of positions and orientations*. Our method is different in that it is based on the standard convolution paradigm instead of intricate wiring of interaction modules, and we use dense grids over the sphere instead of sparse equivariant grids defined by edge directions. A notable result from (Gasteiger et al., 2021) is that networks over position-orientation space are *universal equivariant approximators* and thus as expressive as equivariant methods that use feature representations over the full group $SE(n)$.

Our method connects to a line of papers on theory and algorithms for the processing of functions over position-orientation space via *orientation scores*. See e.g., (Duits & Franken, 2010a;b) and (Janssen et al., 2018) for respectively 2D and 3D applications (Duits et al., 2021). Noteworthy, Portegies et al. (2015) derive invariant attributes for position-orientation pairs derived via the logarithmic map on $SE(3)$. In particular, these works address the processing of Diffusion-weighted MRI data -which are signals over $\mathbb{R}^3 \times S^2$, for which our work provides a simple equivariant learning recipe.

5 EXPERIMENTS

In this section, we evaluate our approach. Comprehensive implementation details, including architecture specifications and optimization techniques, can be found in Appx. C and D.

Benchmark 1: Predicting interatomic potential energy and forces on rMD17. rMD17 (Christensen & Von Lilienfeld, 2020) is a dataset comprising molecular dynamics trajectories of ten small molecules. Each molecule is represented as an atomic point cloud consisting of 3D atom positions and atomic numbers. The objective is to predict the energy-conserving force field for each molecule. The regression targets are the total energy of each molecule and the force on each atom. Following common practice and the laws of physics, we predict the energy with an invariant model and compute the forces at each atom as the gradients of the predicted energy with respect to the positions of atoms. We utilize the PΘNITA model as given in Eq. 16, with a linear embedding layer that takes one-hot encodings of the atom numbers as input and maps them onto the sphere via the ScalarToSphere module. Following MACE (Batatia et al., 2022b), DimeNet and GemNet (Gasteiger et al., 2019; 2021), we predict a node-wise energy E_{io}^l after every layer l using a linear readout layer and obtain the total energy as the sum of all predicted energies $E = \sum_{i,o,l} E_{io}^l$.

Our results are reported in Tab. 1. As a baseline, we construct the same architecture as PΘNITA, except that internal feature maps live solely in position space \mathbb{R}^3 . We label it PNITA. In this case, the pairwise attribute is simply the distance (Eq. 7), which makes PNITA similar to SchNet (Schütt et al., 2023). The inability of such models to learn direction-sensitive representations explains the large performance gap relative to PΘNITA and other methods that capture directional information.

Tab. 1 also shows that PΘNITA either matches or outperforms the state-of-the-art on several molecules while being remarkably faster. PΘNITA is about 3.5 times faster than the established NequIP method (Batzner et al., 2022): a type of steerable tensor field network (Thomas et al., 2018). We benchmarked their optimized codebase within our training framework and recorded a training time per epoch of 20.7, compared to 5.7 seconds for PNITA and 6.1 for PΘNITA. Additionally, PΘNITA stands out for its *simplicity*. It is a general purpose convolutional network, whose design does not require in-depth knowledge of representation theory and molecular modeling like UNiTE (Qiao et al., 2022), MACE (Batatia et al., 2022b), NequIP (Batzner et al., 2022), Allegro (Musaelian et al., 2023) and BOTNet (Batatia et al., 2022a), nor does it involve an intricate wiring of interaction layers as in GemNet (Gasteiger et al., 2021) and ViSNet (Wang et al., 2022).

Benchmark 2: Generating molecules with equivariant diffusion models trained on QM9. Diffusion models have proven to be very effective for the generation of data such as images (Song et al., 2021) and molecules (Hoogetboom et al., 2022). Prior work Gebauer et al. (2019); Simonovsky & Komodakis (2018); Simm et al. (2021) has demonstrated the importance of leveraging molecular symmetries for generalization. Through weight-sharing, PΘNITA is designed to represent such symmetries, making it a promising candidate for this application. Recently advancements like EDMs (Hoogetboom et al., 2022), and MiDi (Vignac et al., 2023) have employed denoising diffusion for molecular data generation. These models utilize EGNN (Satorras et al., 2021) as an equivariant denoising diffusion model, which operates on both atomic positions and types.

We extend EDMs by building a similar architecture for molecular generation, incorporating a joint diffusion process for both continuous coordinates and discrete features. However, we substitute the EGNN backbone with our proposed method PΘNITA. For the readout layer in Eq. 16, we use the Sphere2Vec module to predict denoising displacement vectors. We train on the QM9 dataset (Ramakrishnan R., 2014), a standard dataset containing molecular properties, one-hot representations of atom types (H, C, N, O, F) and 3D coordinates for 130K molecules with up to 9 heavy atoms.

Tab. 2 reveals that PΘNITA outperforms several baselines in terms of atom stability –determined by atoms with the right valency–, and molecule stability –determined by the proportion of molecules for which all atoms are stable–. We attribute this significant improvement to PΘNITA’s ability to

Table 1: Mean absolute errors (MAE) of energy (E) (kcal/mol) and force (F) (kcal/mol/Å) for 10 small organic molecules on rMD17 compared with the state-of-the-art.

Molecule		UNiTE	GemNet	NequIP (20.7 sec/epoch)	BOTNet	Allegro	MACE	ViSNet	PNITA (5.7 sec/epoch)	PONITA (6.1 sec/epoch)
Aspirin	E	2.4	-	2.3	2.3	2.3	2.2	1.9	4.7	1.7 \pm 0.03
	F	7.6	9.5	8.2	8.5	7.3	6.6	6.6	16.3	5.8 \pm 0.18
Azobenzene	E	1.1	-	0.7	0.7	1.2	1.2	0.7	3.2	0.7 \pm 0.01
	F	4.2	-	2.9	3.3	2.6	3.0	2.5	12.2	2.3 \pm 0.11
Benzene	E	0.07	-	0.04	0.03	0.3	0.4	0.03	0.2	0.17 \pm 0.01
	F	0.73	0.5	0.3	0.3	0.2	0.3	0.2	0.4	0.3 \pm 0.00
Ethanol	E	0.62	-	0.4	0.4	0.4	0.4	0.3	0.7	0.4 \pm 0.02
	F	3.7	3.6	2.8	3.2	2.1	2.1	2.3	4.1	2.5 \pm 0.09
Malonaldehyde	E	1.1	-	0.8	0.8	0.6	0.8	0.6	0.9	0.6 \pm 0.05
	F	6.6	6.6	5.1	5.8	3.6	4.1	3.9	5.1	4.0 \pm 0.20
Naphthalene	E	0.46	-	0.2	0.2	0.2	0.4	0.2	1.1	0.3 \pm 0.00
	F	2.6	1.9	1.3	1.8	0.9	1.6	1.3	5.6	1.3 \pm 0.00
Paracetamol	E	1.9	-	1.4	1.3	1.5	1.3	1.1	2.8	1.1 \pm 0.03
	F	7.1	-	5.9	5.8	4.9	4.8	4.5	11.4	4.3 \pm 0.22
Salicylic acid	E	0.73	-	0.7	0.8	0.9	0.9	0.7	1.7	0.7 \pm 0.00
	F	3.8	5.3	4.0	4.3	2.9	3.1	3.4	8.6	3.3 \pm 0.07
Toluene	E	0.45	-	0.3	0.3	0.4	0.5	0.3	0.6	0.3 \pm 0.01
	F	2.5	2.2	1.6	1.9	1.8	1.5	1.1	3.4	1.3 \pm 0.07
Uracil	E	0.58	-	0.4	0.4	0.6	0.5	0.3	0.9	0.4 \pm 0.04
	F	3.8	3.8	3.1	3.2	1.8	2.1	2.1	5.6	2.4 \pm 0.09

Table 2: Molecule generation via denoising diffusion models trained on QM9. Negative Log Likelihood (NLL), atom, and molecule stability.

Models	NLL	Atom stability	Mol stability
E-NF*	-59.7	85.0	4.9
G-Schnet*	N.A.	95.7	68.1
GDM*	-94.7	97.0	63.2
GDM-aug*	-92.5	97.6	71.6
EDM*	-110.7 \pm 1.5	98.7 \pm 0.1	82.0 \pm 0.4
PONITA	-137.4	98.9	87.8

Table 3: Mean squared error on N-body trajectory prediction.

Method	MSE	sec/epoch
SE(3)-Tr.	.0244	
TFN	.0155	
NMP	.0107	
Radial Field	.0104	
EGNN	.0070 \pm .00022	
SEGNN _{G+P}	.0043 \pm .00015	1.59
CGENN	.0039 \pm .0001	
PONITA	.0043 \pm .0001	0.66

learn orientation-sensitive representations: a feature absent in other equivariant diffusion models such as G-Schnet (Gebauer et al., 2019), E-NF (Garcia Satorras et al., 2021) and therein introduced non-equivariant baseline Graph Diffusion models (GDM) trained with and without augmentation.

Benchmark 3: Predicting trajectories in N-body systems. Finally, we benchmark PONITA on the charged N-body particle system experiment proposed in Kipf et al. (2018). It involves five particles that carry either a positive or negative charge, each having initial position and velocity in a 3D space. The objective is to predict all particle positions after a fixed number of time steps. We adapt the experimental setup from (Satorras et al., 2021; Brandstetter et al., 2021) and use PONITA for trajectory prediction. Following the baseline methods, PONITA takes two scalars: velocity magnitude and charge, and two vectors: initial velocity and direction to the center of the system, as input, and outputs a displacement vector. We measure the performance in Mean Squared Error (MSE) and compare against SE(3)-Transformers (Fuchs et al., 2020), Tensor Field Networks (Thomas et al., 2018), Neural Message Passing (Gilmer et al., 2017b), Radial Field (Köhler et al., 2020), EGNN (Satorras et al., 2021), SEGNN (Brandstetter et al., 2021) and CGENN (Ruhe et al., 2023). Tab. 3 shows the results and the reproduced train-time/epoch for SEGNN, showing that PONITA is ± 2.5 times faster. Once again, our versatile equivariant architecture, PONITA, delivers outstanding performance.

Benchmark 4: Superpixel MNIST and Benchmark 5: QM9 regression. In App. E.1 and E.2 we further show the potency of PONITA with state-of-the-art results on benchmarks in 2D and 3D and further compare an edge-index-induced $\mathbb{R}^3 \times S^2$ point cloud vs a grid-based (fiber bundle) approach.

6 CONCLUSION

In summary, our method simplifies the construction of equivariant neural networks by following the paradigm of weight-sharing, achieving state-of-the-art results in various benchmarks. In Thm. 1, we derived geometric attributes that fully characterize the space of equivalence classes of point-pairs in a homogeneous space of $SE(n)$. As such, *these attributes are all you need* to facilitate $SE(n)$ equivariant weight-sharing. As an application of the theory, we propose PONITA: an efficient convolutional architecture for equivariant processing of 3D point clouds. We show that our general-purpose equivariant approach serves as a compelling alternative to prevailing tensor field methods.

REFERENCES

- Brandon Anderson, Truong Son Hy, and Risi Kondor. Cormorant: Covariant molecular neural networks. *Advances in Neural Information Processing Systems*, 32, 2019.
- Jimmy Aronsson. *Mathematical Foundations of Equivariant Neural Networks*. Chalmers University of Technology, PhD Thesis, 2023.
- Pedro HC Avelar, Anderson R Tavares, Thiago LT da Silveira, Cláudio R Jung, and Luís C Lamb. Superpixel image classification with graph attention networks. In *2020 33rd SIBGRAPI Conference on Graphics, Patterns and Images (SIBGRAPI)*, pp. 203–209. IEEE, 2020.
- Ilyes Batatia, Simon Batzner, Dávid Péter Kovács, Albert Musaelian, Gregor NC Simm, Ralf Drautz, Christoph Ortner, Boris Kozinsky, and Gábor Csányi. The design space of $e(3)$ -equivariant atom-centered interatomic potentials. *arXiv preprint arXiv:2205.06643*, 2022a.
- Ilyes Batatia, David P Kovacs, Gregor Simm, Christoph Ortner, and Gábor Csányi. Mace: Higher order equivariant message passing neural networks for fast and accurate force fields. *Advances in Neural Information Processing Systems*, 35:11423–11436, 2022b.
- Simon Batzner, Albert Musaelian, Lixin Sun, Mario Geiger, Jonathan P Mailoa, Mordechai Kornbluth, Nicola Molinari, Tess E Smidt, and Boris Kozinsky. $E(3)$ -equivariant graph neural networks for data-efficient and accurate interatomic potentials. *Nature communications*, 13(1): 2453, 2022.
- Erik J Bekkers. B-spline cnns on lie groups. In *International Conference on Learning Representations*, 2019.
- Erik J Bekkers, Maxime W Lafarge, Mitko Veta, Koen AJ Eppenhof, Josien PW Pluim, and Remco Duits. Roto-translation covariant convolutional networks for medical image analysis. In *Medical Image Computing and Computer Assisted Intervention—MICCAI 2018: 21st International Conference, Granada, Spain, September 16-20, 2018, Proceedings, Part I*, pp. 440–448. Springer, 2018.
- Lukas Biewald. Experiment tracking with weights and biases, 2020. URL <https://www.wandb.com/>. Software available from wandb.com.
- Johannes Brandstetter, Rob Hesselink, Elise van der Pol, Erik J Bekkers, and Max Welling. Geometric and physical quantities improve $e(3)$ equivariant message passing. In *International Conference on Learning Representations*, 2021.
- Gabriele Cesa, Leon Lang, and Maurice Weiler. A program to build $e(n)$ -equivariant steerable cnns. In *International Conference on Learning Representations*, 2021.
- François Chollet. Xception: Deep learning with depthwise separable convolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 1251–1258, 2017.
- Anders S Christensen and O Anatole Von Lilienfeld. On the role of gradients for machine learning of molecular energies and forces. *Machine Learning: Science and Technology*, 1(4):045018, 2020.
- Taco Cohen and Max Welling. Group equivariant convolutional networks. In *International conference on machine learning*, pp. 2990–2999. PMLR, 2016.
- Taco S Cohen, Mario Geiger, and Maurice Weiler. A general theory of equivariant cnns on homogeneous spaces. *Advances in neural information processing systems*, 32, 2019.
- Pim De Haan, Maurice Weiler, Taco Cohen, and Max Welling. Gauge equivariant mesh cnns: Anisotropic convolutions on geometric graphs. In *International Conference on Learning Representations*, 2020.
- Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. *Advances in neural information processing systems*, 34:8780–8794, 2021.

- R Duits, BMN Smets, AJ Wemmenhove, JW Portegies, and EJ Bekkers. Recent geometric flows in multi-orientation image processing via a cartan connection. Handbook of Mathematical Models and Algorithms in Computer Vision and Imaging: Mathematical Imaging and Vision, pp. 1–60, 2021.
- Remco Duits and Erik Franken. Left-invariant parabolic evolutions on $se(2)$ and contour enhancement via invertible orientation scores part i: Linear left-invariant diffusion equations on $se(2)$. Quarterly of Applied Mathematics, 68(2):255–292, 2010a.
- Remco Duits and Erik Franken. Left-invariant parabolic evolutions on $se(2)$ and contour enhancement via invertible orientation scores part ii: Nonlinear left-invariant diffusions on invertible orientation scores. Quarterly of applied mathematics, 68(2):293–331, 2010b.
- Nadav Dym and Haggai Maron. On the universality of rotation equivariant point cloud networks. In International Conference on Learning Representations, 2020.
- Matthias Fey and Jan E. Lenssen. Fast graph representation learning with PyTorch Geometric. In ICLR Workshop on Representation Learning on Graphs and Manifolds, 2019.
- 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.
- Marc Finzi, Samuel Stanton, Pavel Izmailov, and Andrew Gordon Wilson. Generalizing convolutional neural networks for equivariance to lie groups on arbitrary continuous data. In International Conference on Machine Learning, pp. 3165–3176. PMLR, 2020a.
- Marc Anton Finzi, Roberto Bondesan, and Max Welling. Probabilistic numeric convolutional neural networks. In International Conference on Learning Representations, 2020b.
- Fabian Fuchs, Daniel Worrall, Volker Fischer, and Max Welling. Se (3)-transformers: 3d rotation equivariant attention networks. Advances in neural information processing systems, 33:1970–1981, 2020.
- Victor Garcia Satorras, Emiel Hoogeboom, Fabian Fuchs, Ingmar Posner, and Max Welling. E (n) equivariant normalizing flows. Advances in Neural Information Processing Systems, 34:4181–4192, 2021.
- Johannes Gasteiger, Janek Groß, and Stephan Günnemann. Directional message passing for molecular graphs. In International Conference on Learning Representations, 2019.
- Johannes Gasteiger, Shankari Giri, Johannes T Margraf, and Stephan Günnemann. Fast and uncertainty-aware directional message passing for non-equilibrium molecules. arXiv preprint arXiv:2011.14115, 2020.
- Johannes Gasteiger, Florian Becker, and Stephan Günnemann. Gemnet: Universal directional graph neural networks for molecules. In A. Beygelzimer, Y. Dauphin, P. Liang, and J. Wortman Vaughan (eds.), Advances in Neural Information Processing Systems, 2021. URL https://openreview.net/forum?id=HS_sOaxS9K-.
- Niklas Gebauer, Michael Gastegger, and Kristof Schütt. Symmetry-adapted generation of 3d point sets for the targeted discovery of molecules. Advances in neural information processing systems, 32, 2019.
- Mario Geiger and Tess Smidt. e3nn: Euclidean neural networks, 2022. URL <https://arxiv.org/abs/2207.09453>.
- Justin Gilmer, Samuel S Schoenholz, Patrick F Riley, Oriol Vinyals, and George E Dahl. Neural message passing for quantum chemistry. In International conference on machine learning, pp. 1263–1272. PMLR, 2017a.

- Justin Gilmer, Samuel S. Schoenholz, Patrick F. Riley, Oriol Vinyals, and George E. Dahl. Neural message passing for quantum chemistry. In Doina Precup and Yee Whye Teh (eds.), Proceedings of the 34th International Conference on Machine Learning, volume 70 of Proceedings of Machine Learning Research, pp. 1263–1272. PMLR, 06–11 Aug 2017b. URL <https://proceedings.mlr.press/v70/gilmer17a.html>.
- Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. Advances in neural information processing systems, 33:6840–6851, 2020.
- Emiel Hoogetboom, Victor Garcia Satorras, Clément Vignac, and Max Welling. Equivariant diffusion for molecule generation in 3d. In International conference on machine learning, pp. 8867–8887. PMLR, 2022.
- Michiel HJ Janssen, Augustus JEM Janssen, Erik J Bekkers, J Oliván Bescós, and Remco Duits. Design and processing of invertible orientation scores of 3d images. Journal of mathematical imaging and vision, 60:1427–1458, 2018.
- Diederik Kingma, Tim Salimans, Ben Poole, and Jonathan Ho. Variational diffusion models. Advances in neural information processing systems, 34:21696–21707, 2021.
- Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980, 2014.
- Thomas Kipf, Ethan Fetaya, Kuan-Chieh Wang, Max Welling, and Richard Zemel. Neural relational inference for interacting systems. In International conference on machine learning, pp. 2688–2697. PMLR, 2018.
- David M Knigge, David W Romero, and Erik J Bekkers. Exploiting redundancy: Separable group convolutional networks on lie groups. In International Conference on Machine Learning, pp. 11359–11386. PMLR, 2022.
- Jonas Köhler, Leon Klein, and Frank Noé. Equivariant flows: exact likelihood generative learning for symmetric densities. In International conference on machine learning, pp. 5361–5370. PMLR, 2020.
- Yeskendir Koishekenov and Erik J Bekkers. An exploration of conditioning methods in graph neural networks. In ICLR 2023-Machine Learning for Drug Discovery workshop, 2023.
- Risi Kondor and Shubhendu Trivedi. 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.
- Risi Kondor, Zhen Lin, and Shubhendu Trivedi. Clebsch–gordan nets: a fully fourier space spherical convolutional neural network. Advances in Neural Information Processing Systems, 31, 2018.
- Thijs P Kuipers and Erik J Bekkers. Regular se (3) group convolutions for volumetric medical image analysis. In Medical Image Computing and Computer Assisted Intervention–MICCAI 2023. arXiv preprint arXiv:2306.13960, 2023.
- Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. Proceedings of the IEEE, 86(11):2278–2324, 1998.
- Yi Liu, Limei Wang, Meng Liu, Yuchao Lin, Xuan Zhang, Bora Oztekin, and Shuiwang Ji. Spherical message passing for 3d molecular graphs. In International Conference on Learning Representations, 2022a. URL <https://openreview.net/forum?id=givsRXsOt9r>.
- Zhuang Liu, Hanzi Mao, Chao-Yuan Wu, Christoph Feichtenhofer, Trevor Darrell, and Saining Xie. A convnet for the 2020s. In Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pp. 11976–11986, 2022b.
- Federico Monti, Davide Boscaini, Jonathan Masci, Emanuele Rodola, Jan Svoboda, and Michael M Bronstein. Geometric deep learning on graphs and manifolds using mixture model cnns. In Proceedings of the IEEE conference on computer vision and pattern recognition, pp. 5115–5124, 2017.

- Albert Musaelian, Simon Batzner, Anders Johansson, Lixin Sun, Cameron J Owen, Mordechai Kornbluth, and Boris Kozinsky. Learning local equivariant representations for large-scale atomistic dynamics. *Nature Communications*, 14(1):579, 2023.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32, 2019.
- Michael Poli, Stefano Massaroli, Eric Nguyen, Daniel Y Fu, Tri Dao, Stephen Baccus, Yoshua Bengio, Stefano Ermon, and Christopher Ré. Hyena hierarchy: Towards larger convolutional language models. *arXiv preprint arXiv:2302.10866*, 2023.
- Jorg Portegies, Gonzalo Sanguinetti, Stephan Meesters, and Remco Duits. New approximation of a scale space kernel on $se(3)$ and applications in neuroimaging. In *International Conference on Scale Space and Variational Methods in Computer Vision*, pp. 40–52. Springer, 2015.
- Zhuoran Qiao, Anders S Christensen, Matthew Welborn, Frederick R Manby, Anima Anandkumar, and Thomas F Miller III. Informing geometric deep learning with electronic interactions to accelerate quantum chemistry. *Proceedings of the National Academy of Sciences*, 119(31): e2205221119, 2022.
- Rupp M. and Von Lilienfeld O.A. Ramakrishnan R., Dral P.O. Quantum chemistry structures and properties of 134 kilo molecules, 2014.
- David W Romero, Anna Kuzina, Erik J Bekkers, Jakub Mikolaj Tomczak, and Mark Hoogenboom. Ckconv: Continuous kernel convolution for sequential data. In *International Conference on Learning Representations*, 2021.
- David Ruhe, Johannes Brandstetter, and Patrick Forré. Clifford group equivariant neural networks. *arXiv preprint arXiv:2305.11141*, 2023.
- Victor Garcia Satorras, Emiel Hooeboom, and Max Welling. E(n) equivariant graph neural networks. In *International conference on machine learning*, pp. 9323–9332. PMLR, 2021.
- Kristof T Schütt, Stefaan SP Hessmann, Niklas WA Gebauer, Jonas Lederer, and Michael Gastegger. Schnetpack 2.0: A neural network toolbox for atomistic machine learning. *The Journal of Chemical Physics*, 158(14), 2023.
- Gregor N. C. Simm, Robert Pinsler, Gábor Csányi, and José Miguel Hernández-Lobato. Symmetry-aware actor-critic for 3d molecular design. In *International Conference on Learning Representations*, 2021. URL <https://openreview.net/forum?id=jEYKjPE1xYN>.
- Martin Simonovsky and Nikos Komodakis. Graphvae: Towards generation of small graphs using variational autoencoders. In *Artificial Neural Networks and Machine Learning—ICANN 2018: 27th International Conference on Artificial Neural Networks, Rhodes, Greece, October 4-7, 2018, Proceedings, Part I 27*, pp. 412–422. Springer, 2018.
- Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *International conference on machine learning*, pp. 2256–2265. PMLR, 2015.
- Yang Song, Conor Durkan, Iain Murray, and Stefano Ermon. Maximum likelihood training of score-based diffusion models. *Advances in Neural Information Processing Systems*, 34:1415–1428, 2021.
- Ivan Sosnovik, Michał Szmaja, and Arnold Smeulders. Scale-equivariant steerable networks. In *International Conference on Learning Representations*, 2019.
- Matthew Tancik, Pratul Srinivasan, Ben Mildenhall, Sara Fridovich-Keil, Nithin Raghavan, Utkarsh Singhal, Ravi Ramamoorthi, Jonathan Barron, and Ren Ng. Fourier features let networks learn high frequency functions in low dimensional domains. *Advances in Neural Information Processing Systems*, 33:7537–7547, 2020.

- Nathaniel Thomas, Tess Smidt, Steven Kearnes, Lusann Yang, Li Li, Kai Kohlhoff, and Patrick Riley. Tensor field networks: Rotation-and translation-equivariant neural networks for 3d point clouds. [arXiv preprint arXiv:1802.08219](#), 2018.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. [Advances in neural information processing systems](#), 30, 2017.
- Clement Vignac, Nagham Osman, Laura Toni, and Pascal Frossard. Midi: Mixed graph and 3d denoising diffusion for molecule generation. [arXiv preprint arXiv:2302.09048](#), 2023.
- Soledad Villar, David W Hogg, Kate Storey-Fisher, Weichi Yao, and Ben Blum-Smith. Scalars are universal: Equivariant machine learning, structured like classical physics. [Advances in Neural Information Processing Systems](#), 34:28848–28863, 2021.
- Ian Walker and Ben Glocker. Graph convolutional gaussian processes. In [International Conference on Machine Learning](#), pp. 6495–6504. PMLR, 2019.
- Yusong Wang, Shaoning Li, Xinheng He, Mingyu Li, Zun Wang, Nanning Zheng, Bin Shao, Tong Wang, and Tie-Yan Liu. Visnet: a scalable and accurate geometric deep learning potential for molecular dynamics simulation. [arXiv preprint arXiv:2210.16518](#), 2022.
- Maurice Weiler and Gabriele Cesa. General e (2)-equivariant steerable cnns. [Advances in neural information processing systems](#), 32, 2019.
- Maurice Weiler, Fred A Hamprecht, and Martin Storath. Learning steerable filters for rotation equivariant cnns. In [Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition](#), pp. 849–858, 2018.
- Maurice Weiler, Patrick Forré, Erik Verlinde, and Max Welling. Coordinate independent convolutional networks—*isometry and gauge equivariant convolutions on riemannian manifolds*. [arXiv preprint arXiv:2106.06020](#), 2021.
- Maurice Weiler, Patrick Forré, Erik Verlinde, and Max Welling. [Equivariant and Coordinate Independent Convolutional Networks](#). 2023.
URL https://maurice-weiler.gitlab.io/cnn_book/EquivariantAndCoordinateIndependentCNNs.pdf.
- Daniel Worrall and Gabriel Brostow. Cubenet: Equivariance to 3d rotation and translation. In [Proceedings of the European Conference on Computer Vision \(ECCV\)](#), pp. 567–584, 2018.
- Minkai Xu, Lantao Yu, Yang Song, Chence Shi, Stefano Ermon, and Jian Tang. Geodiff: A geometric diffusion model for molecular conformation generation. [arXiv preprint arXiv:2203.02923](#), 2022.

A THE FIBER BUNDLE VS POINT CLOUD VIEWPOINT

A.1 REGULAR VS STEERABLE GROUP CONVOLUTIONS

In the field of group equivariant deep learning, one often adopts a steerable or regular convolutional viewpoint on equivariant neural networks (Weiler & Cesa, 2019). Let us consider the case of $SE(n)$ equivariance, and note that the group $SE(n)$ is a semi-direct product between the translation group and the rotation group $SO(n)$. We then consider functions (feature fields) over domains X that are homogeneous spaces of $SE(n)$, namely the case $X = \mathbb{R}^n$, $X = \mathbb{R}^n \times S^{n-1}$, and $X = SE(n)$. The difference between steerable and regular group convolutions then lies in the type of feature fields that are considered.

With *regular group convolutions* one considers multi-channel scalar fields $f : X \rightarrow \mathbb{R}^c$ of the homogeneous space X . And convolutions are of the form

$$f^{out}(y) = \int_X k(g_y^{-1}x) f^{in}(x) dx, \quad (\text{A1})$$

with g_y a representative of the point y such that $y = g_y x_0$ with x_0 some arbitrary origin in X . Regular feature fields, and convolution kernels, transform via the regular representation of $SE(n)$ via

$$(\rho^L(g)k)(x) := k(g^{-1}x).$$

Given this, regular group convolutions can be simply considered as template matching between a roto-translated convolution kernel (by transforming it with the regular representation) and the underlying signal, where similarity is measured by the \mathbb{L}_2 inner product. When $X \equiv SE(n)/H$ with H a non-trivial group, such as $H = SO(3)$ when $X = \mathbb{R}^n$ or $H = SO(2)$ when $X = \mathbb{R}^3 \times S^2$, then the kernel k should satisfy a kernel constraint $\forall_{h \in H} : k(x) = k(hx)$, see (Bekkers, 2019). As discussed in the main body of this paper, this constraint is automatically satisfied if we parametrize the kernel by the invariant attributes a_{ij} of Theorem 1. The advantage of regular group convolutions is that the codomain of the feature fields are just scalars and one can use any point-wise activation function without breaking equivariance. A drawback of such methods though is that one needs to discretize the domain X , including the fibers S^{n-1} or $SO(n)$, which leads to numerical inexactness of the equivariance property¹.

With *steerable group convolutions* one considers *feature fields* $f : \mathbb{R}^n \rightarrow V_\rho$ over a base space \mathbb{R}^n with a codomain which is a vector space V_ρ on which a representation ρ of $SO(n)$ can act. Such feature fields transform via the so-called *induced representation* of $SE(n)$, defined as

$$([\text{Ind}_{SO(n)}^{SE(n)} \rho](g)f)(\mathbf{x}) = \rho(\mathbf{R}) f(g^{-1}\mathbf{x}).$$

I.e., the induced representation $[\text{Ind}_{SO(n)}^{SE(n)} \rho]$ does not only transform the feature map’s domain but also its codomain via the representation ρ . For example, if one roto-translates a vector field $\mathbf{v} : \mathbb{R}^3 \rightarrow \mathbb{R}^3$, the vectors $\mathbf{v}(\mathbf{x})$ in the field at location \mathbf{x} should be moved to the new locations $g^{-1}\mathbf{x}$, but also their values should be rotated via $\mathbf{R}\mathbf{v}(g^{-1}\mathbf{x})$. The steerable convolutions are then simply given by

$$f^{out}(\mathbf{y}) = \int_{\mathbb{R}^n} k(\mathbf{x} - \mathbf{y}) f^{in}(\mathbf{x}) d\mathbf{x}. \quad (\text{A2})$$

Now, however, the kernel needs to satisfy the kernel constraint

$$k(\mathbf{R}\mathbf{x}) = \rho_{out}(\mathbf{R})k(\mathbf{x})\rho_{in}(\mathbf{R}^{-1}),$$

with ρ_{in} and ρ_{out} the representations that define the vectors spaces of the codomains of the input and output feature fields respectively, see e.g. (Weiler & Cesa, 2019; Cesa et al., 2021). An advantage of the steerable approach is that one can obtain exact equivariance by choosing to let the vector spaces V_ρ transform via representations $\rho = \oplus_l \rho_l$ that are given as a direct sum of irreducible representations. Such irreducible representations do not require a grid and rotations can be computed exactly. A drawback of such an approach is, however, that one can no longer apply arbitrary point-wise activation functions but one needs specialized operations on V_ρ in order not to break equivariance. This limits expressivity (Weiler & Cesa, 2019). Another drawback, arguably, is that this approach requires considerable understanding of representation theory for proper use in practice.

¹Our experiments, however, show that this approximate equivariance is not detrimental to performance

A.2 REGULAR GROUP CONVOLUTIONS IMPLEMENTED AS STEERABLE GROUP CONVOLUTIONS

It is important to note that in the regular group convolution case for $SE(n)$ equivariance, one recognizes a semi-direct product structure of the group $SE(n) = (\mathbb{R}^n, +) \rtimes SO(n)$, with $(\mathbb{R}^n, +)$ the translation group, which as a space can be identified with position space \mathbb{R}^n . When one then considers regular group convolutions over $X = \mathbb{R}^n \times S^{n-1}$ or $X = \mathbb{R}^n \times SO(n)$, one could think of the feature fields $f : \mathbb{R}^n \times Y \rightarrow \mathbb{R}^c$ as assigning to every position $\mathbf{x} \in \mathbb{R}^n$ a function $f_{\mathbf{x}} := f(\mathbf{x}, \cdot)$ over the space Y , with here $Y = S^{n-1}$ or $Y = SO(n)$. These functions $f_{\mathbf{x}}$ are to be considered vectors in the infinite-dimensional vector space of square-integrable functions $\mathbb{L}_2(Y)$. For this vector space we know the left-regular representations $\rho_{\mathbb{L}_2(Y)}^L$, which are simply given as

$$(\rho_{\mathbb{L}_2(S^{n-1})}^L(\mathbf{R})f_{\mathbf{x}})(\mathbf{o}) = f_{\mathbf{x}}(\mathbf{R}^T \mathbf{o}), \quad \text{and} \quad (\rho_{\mathbb{L}_2(SO(n))}^L(\mathbf{R})f_{\mathbf{x}})(\mathbf{R}') = f_{\mathbf{x}}(\mathbf{R}^T \mathbf{R}').$$

Thus, one can think of regular group convolutions as working with feature fields $f : \mathbb{R}^n \rightarrow V_{\rho_{\mathbb{L}_2(Y)}^L}$.

Further, note that any $SO(n)$ representation can be decomposed into a direct sum of irreducible representations via a change of basis. I.e., there is a linear change of basis such that $V_{\rho_{\mathbb{L}_2(Y)}^L} \equiv V_{\rho}$ for some $\rho = \oplus_l \rho_l$, with ρ_l an irreducible representation. This change of basis is given by the Fourier transform on Y . Namely, a function $f_{\mathbf{x}} \in V_{\rho_{\mathbb{L}_2(Y)}^L}$ can be transformed to a vector of Fourier coefficients via $\hat{\mathbf{f}}_{\mathbf{x}} = \mathcal{F}_Y[f_{\mathbf{x}}] \in V_{\rho}$, and back via the inverse Fourier transform $f_{\mathbf{x}} = \mathcal{F}_Y^{-1}[\hat{\mathbf{f}}_{\mathbf{x}}]$.

Regular group convolutions can thus either be implemented as steerable convolutions with fields of signals over Y , i.e., $f : \mathbb{R}^n \rightarrow V_{\rho_{\mathbb{L}_2(Y)}^L}$, or via fields of irreducible representations $f : \mathbb{R}^n \rightarrow V_{\rho}$ that can be thought of as fields of Fourier coefficients. In this latter case, the kernel constraint is satisfied by parametrizing k using the Clebsch-Gordan tensor product (Cesa et al., 2021; Brandstetter et al., 2021). This approach is in fact how group equivariant neural networks are often implemented (Weiler & Cesa, 2019; Thomas et al., 2018), however, it has the drawback that one either has to use specialized activation functions or use activation functions of the form `InverseFourier` \rightarrow `ActivationFunction` \rightarrow `FourierTransform`, which leads to computational overhead and leads to discretization artifacts that break exact equivariance. In this paper, we present a fully *regular* group convolution viewpoint that does not require Fourier transforms or specialized activation functions.

A.3 FIBER BUNDLE VIEWPOINT FOR PΘNITA

From a fiber bundle-theoretic viewpoint, the feature fields $f : \mathbb{R}^n \rightarrow V_{\rho}$ can be regarded as sections of vector bundles associated with the principal fiber bundle $\mathbb{R}^n \times SO(n) \rightarrow \mathbb{R}^n$ (Weiler et al., 2023; Aronsson, 2023). As such, without going into further technical details behind this theory, we merely use this connection to refer to the vector spaces V_{ρ} as fibers. The PΘNITA model processes feature fields $f : \mathbb{R}^n \rightarrow V_{\rho_{\mathbb{L}_2(S^{n-1})}^L}$ with vector spaces of spherical signals $V_{\rho_{\mathbb{L}_2(S^{n-1})}^L}$ as fibers. Instead of describing these signals in a basis of irreducible representations (the Fourier approach), we sample them on a grid $\mathcal{S} \subset S^{n-1}$ of orientations. See Figure 1 in which the grid is depicted as a set of yellow dots, and the actual spherical signal is depicted as a gray-scale density. Our model makes efficient use of the fiber bundle structure by separating convolutional operations over the field into a spatial interaction part (convolving over \mathbb{R}^n) followed by a within-fiber orientation interaction part (convolving over S^{n-1}).

A.4 A POINT CLOUD VIEWPOINT FOR PΘNITA

The efficient implementation of PΘNITA is made possible by the fact that each node shares the same grid (fiber) of orientations \mathcal{S} . Let us denote with $\mathcal{S}_{\mathbf{x}} \subset S^{n-1}$ the grid of orientations available at position \mathbf{x} . If we share the same grid over all positions, i.e., $\mathcal{S}_{\mathbf{x}} = \mathcal{S}$, we can decide to spatially convolve per orientation in the grid as follows

$$f^{out}(\mathbf{x}, \mathbf{o}) = \int_{\mathbb{R}^n} k([\mathbf{x}, \mathbf{o}], [\mathbf{x}', \mathbf{o}]) f^{in}(\mathbf{x}', \mathbf{o}) d\mathbf{x}',$$

which is the continuous formulation of the discretized spatial convolution of Eq. (14). This purely spatial convolution is made possible by the fact that at each position \mathbf{x}' we can sample the same

orientation $\mathbf{o} \in \mathcal{S}$. However, if we do not share the same grid \mathcal{S} over all positions we can no longer separate the convolution. In this case, we adopt a point-cloud viewpoint.

There are natural instances in which one has to compute with point clouds in position-orientation space. Examples include triangulated shape meshes, which can be represented by the centroids of the triangles \mathbf{p} and their normal vectors \mathbf{o} . See e.g. (De Haan et al., 2020). In contrast to (De Haan et al., 2020) our method would take the curvature of the surface mesh into account (via the invariant $\mathbf{o}_i^T \mathbf{o}_j$), as well as out of plane distance, however, the method by De Haan et al. (2020) takes anisotropic features within the plane into account whereas the second invariant in Eq. (9)) restricts the kernel to be isotropic within the tangent plane.

Another instance of a position-orientation space point cloud would be in a message passing over the edges in an atomic point cloud, as in DimeNet/GemNet (Gasteiger et al., 2019; 2021). In those works, a grid is assigned to each atom location \mathbf{x}_i based on the direction towards neighboring atom locations \mathbf{x}_j that are connected by covalent bonds. I.e., each atom obtains a grid $\mathcal{S}_{\mathbf{x}_i} = \{ \frac{\mathbf{x}_j - \mathbf{x}_i}{\|\mathbf{x}_j - \mathbf{x}_i\|} \mid j \in \mathcal{N}(i) \}$, with $\mathcal{N}(i)$ the set of neighbor indices of node i .

Convolutions in these case can no longer be separable and include spatial and orientation interactions simultaneously:

$$f^{out}(\mathbf{x}_i, \mathbf{o}_i) = \sum_{j \in \mathcal{N}(i)} \sum_{\mathbf{o} \in \mathcal{S}_{\mathbf{x}_j}} k([\mathbf{x}_i, \mathbf{o}_i], [\mathbf{x}_j, \mathbf{o}]) f^{in}(\mathbf{x}_j, \mathbf{o}).$$

This is the discretized version of the full continuous convolution of Eq. (11).

A.5 IMPLEMENTATION

Both versions of position-orientation space convolutions –the fiber bundle approach as well as the point cloud approach– are made available in the public repository <https://github.com/ebekkers/ponita>.

B PROOFS

B.1 PROOF OF LEMMA 1

Lemma 1 shows that equivalence classes of point pairs $[x_i, x_j] = \{(g x_i, g x_j) \mid g \in G\}$ correspond to orbits $H g_i^{-1} x_j$, for any chosen g_i such that $x_i = g_i x_0$, with $x_i, x_j, x_0 \in X$. In the following proof, one might recognize that the space of equivalence classes $X \times X / \sim$ can also be identified with the space (H, H) -double cosets in G . That is $X \times X / \sim \equiv H \backslash G / H \equiv H \backslash X$. The proof of Lemma 1 is as follows:

Proof. To prove the bijection given by equation 6 we rewrite the equivalence class into the form of a double coset as follows. For any representatives (g_i, g_j) we have

$$\begin{aligned} [x_i, x_j] &= \{(g x_i, g x_j) \mid g \in G\} \\ &\stackrel{\forall h \in H: h x_0 = x_0}{=} \{(g g_i h_i m_0, g g_j h_j m_0) \mid g \in G, h_i, h_j \in H\} \\ &\stackrel{g \rightarrow g h_i^{-1} g_i^{-1}}{=} \{(h_i^{-1} g_i^{-1} g g_j h_j m_0, g m_0) \mid g \in G, h_i, h_j \in H\} \\ &\stackrel{g^{-1} g = e}{=} \{(h_i^{-1} g_i^{-1} g_j h_j m_0, G m_0) \mid h_i, h_j \in H\} \\ &\stackrel{x_j = g_j m_0, h_j x_0 = m_0}{=} \{(h_i^{-1} g_i^{-1} x_j, G m_0) \mid h_i, h_j \in H\} \\ &\stackrel{G m_0 = M}{=} \{h_i g_i^{-1} x_j \mid h_i \in H\} \\ &= H g_i^{-1} x_j, \end{aligned}$$

where in the second to last step we note that the second item in the tuple, $G x_0$ is constant (it represents the entire homogeneous space $X = G m_0$), and thus we can represent the tuple simply with the first element $H g_i^{-1} x_j$. Thus any equivalence class $[x_i, x_j]$ can be represented by the orbit $H g_i^{-1} x_j$, with g_i any such that $x_i = g_i x_0$. \square

B.2 PROOF OF THEOREM 1

Theorem 1 lists four cases of equivalence classes and provides a bijective map for each of them:

$$\mathbb{R}^2 \text{ and } \mathbb{R}^3 : \quad [\mathbf{p}_i, \mathbf{p}_j] \mapsto a_{ij} = \|\mathbf{p}_j - \mathbf{p}_i\|, \quad (\text{A3})$$

$$\mathbb{R}^2 \times S^1 \text{ and } SE(2) : \quad [(\mathbf{p}_i, \mathbf{o}_i), (\mathbf{p}_j, \mathbf{o}_j)] \mapsto a_{ij} = (\mathbf{R}_{\mathbf{o}_i}^{-1}(\mathbf{p}_j - \mathbf{p}_i), \arccos \mathbf{o}_i^T \mathbf{o}_j), \quad (\text{A4})$$

$$\mathbb{R}^3 \times S^2 : \quad [(\mathbf{p}_i, \mathbf{o}_i), (\mathbf{p}_j, \mathbf{o}_j)] \mapsto a_{ij} = \left(\begin{array}{c} \mathbf{o}_i^T (\mathbf{p}_j - \mathbf{p}_i) \\ \|\mathbf{p}_j - \mathbf{p}_i\| - \mathbf{o}_i^T (\mathbf{p}_j - \mathbf{p}_i) \mathbf{o}_i \\ \arccos \mathbf{o}_i^T \mathbf{o}_j \end{array} \right), \quad (\text{A5})$$

$$SE(3) : \quad [(\mathbf{p}_i, \mathbf{R}_i), (\mathbf{p}_j, \mathbf{R}_j)] \mapsto a_{ij} = (\mathbf{R}_i^{-1}(\mathbf{p}_j - \mathbf{p}_i), \mathbf{R}_i^{-1} \mathbf{R}_j). \quad (\text{A6})$$

We provide a constructive proof for each of these cases by providing an inverse mapping $a_{ij} \mapsto [\mathbf{p}_i, \mathbf{p}_j]$. This boils down to defining representative x_{ij} of the orbit of $H g_i^{-1} x_j$ for a given a_{ij} . The proofs are as follows.

Proof. For the \mathbb{R}^n case the orbit representative can be given by

$$x_{ij} = a_{ij} \mathbf{e}_z, \quad (\text{A7})$$

with $\mathbf{e}_z = \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix}$. The equivalence class is represented by the orbit $H g_i^{-1} x_j = H \mathbf{R}_i^{-1}(\mathbf{p}_j - \mathbf{p}_i) = H \mathbf{p}_{ij}$, with $\mathbf{p}_{ij} = \mathbf{p}_j - \mathbf{p}_i$. Thus \mathbf{p}_{ij} is a valid representative. Since H does not alter the norm of \mathbf{p}_{ij} and it acts transitively on spheres of a given radius, we can say that x_{ij} lies in the orbit since $\|x_{ij}\| = \|a_{ij} \mathbf{e}_z\| = \|\mathbf{p}_{ij}\|$. This proves that the mapping in (A3) is bijective, with (A7) as inverse mapping.

To prove the $\mathbb{R}^2 \times S^1$, $SE(2)$, and $SE(3)$ cases we first remark that $\mathbb{R}^2 \times S^1$ is equivalent to the $SE(2)$. We can uniquely identify any $(\mathbf{p}, \mathbf{o}) \in \mathbb{R}^2 \times S^1$ with a roto-translation $(\mathbf{p}, \mathbf{R}_{\mathbf{o}}) \in SO(2)$ via $(\mathbf{p}, \mathbf{o}) = (\mathbf{p}, \mathbf{R}_{\mathbf{o}}) (\mathbf{0}, \mathbf{e}_x)$ with $\mathbf{e}_x = \begin{pmatrix} 1 \\ 0 \end{pmatrix}$, noting that the stabilizer group $\text{Stab}_{SO(2)}(\mathbf{0}, \mathbf{e}_x) = \{(\mathbf{0}, \mathbf{I})\}$ is trivial. Thus the spaces are equivalent. Since in these three cases the stabilizer group $H = \{e\}$ is trivial, there is only one unique group element associated with each point $x_i \equiv g_i$ and $x_j \equiv g_j$, with g_i, g_j such that $x_i = g_i x_0$ and $x_j = g_j x_0$. And thus the orbit that represents the equivalence class $[x_i, x_j] \equiv H g_i^{-1} x_j = \{g_i^{-1} g_j\}$ consist of a single element. We can write its representative

$$x_{ij} = g_i^{-1} g_j, \quad (\text{A8})$$

which is precisely the relative group action for $SE(2)$ in (A4) and for $SE(3)$ in (A6). This proves that (A4) and (A6) are bijective with (A8) as inverse mapping.

The representative for the $\mathbb{R}^3 \times S^2$ case is given by

$$x_{ij} = \left(\begin{array}{c} b \\ 0 \\ a \end{array} \right), \left(\begin{array}{c} \sin c \\ 0 \\ \cos c \end{array} \right), \quad (\text{A9})$$

with a, b, c the components of the attribute where $a_{ij} = (a, b, c)^T$.

Note that the orbits are given by $H (\mathbf{R}_{\mathbf{o}_i}^{-1} \mathbf{p}_{ij}, \mathbf{R}_{\mathbf{o}_i}^{-1} \mathbf{o}_j)$. Further note that rotations in H around the \mathbf{e}_z axis act on three independent subspaces of $\mathbb{R}^3 \times S^1$. The spatial part of which is given by $V_a = \text{span}\{\mathbf{e}_z\}$, and $V_b = \text{span}\{\mathbf{e}_x, \mathbf{e}_y\}$, with $\mathbf{e}_x = \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix}$ and $\mathbf{e}_y = \begin{pmatrix} 0 \\ 1 \\ 0 \end{pmatrix}$. We project the chosen representative onto this basis, such that $\mathbf{R}_{\mathbf{o}_i}^{-1} \mathbf{p}_{ij} = x \mathbf{e}_x + y \mathbf{e}_y + z \mathbf{e}_z$. First compute z as

$$z = (\mathbf{R}_{\mathbf{o}_i}^{-1} \mathbf{p}_{ij})^T \mathbf{e}_z = \mathbf{p}_{ij}^T \mathbf{R}_{\mathbf{o}_i}^T = \mathbf{p}_{ij}^T \mathbf{o}_i = a.$$

The vector $a \mathbf{e}_z$ is invariant to the action of H . The z component of the unit vector in S^2 is also invariant, and thus uniquely determined as

$$(\mathbf{R}_{\mathbf{o}_i}^{-1} \mathbf{n}_j)^T \mathbf{e}_z = \mathbf{n}_j^T \mathbf{n}_i = \cos c,$$

and thus $c = \arccos \mathbf{n}_i^T \mathbf{n}_j$. Since both quantities a and c are obtained independent of a chosen rotation in H , we can pick any b as long as the vector $\begin{pmatrix} b \\ 0 \\ a \end{pmatrix}$ lies in the orbit $H \mathbf{R}_{\mathbf{o}_i}^{-1} \mathbf{p}_{ij}$. Since H is norm preserving, we only have to check that for the chosen b the norm $\|(b, 0)^T\| = b$ coincides with the norm of the V_b vector component of $\mathbf{R}_{\mathbf{o}_i}^{-1} \mathbf{p}_{ij}$. The component of $\mathbf{R}_{\mathbf{o}_i}^{-1} \mathbf{p}_{ij}$ in V_b is given by

$(\mathbf{R}_{\mathbf{o}_i}^{-1} \mathbf{p}_{ij}) - (\mathbf{p}_{ij}^T \mathbf{o}_i) \mathbf{e}_z$, and its norm by

$$\|(\mathbf{R}_{\mathbf{o}_i}^{-1} \mathbf{p}_{ij}) - (\mathbf{p}_{ij}^T \mathbf{o}_i) \mathbf{e}_z\| = \|(\mathbf{R}_{\mathbf{o}_i} \mathbf{R}_{\mathbf{o}_i}^{-1} \mathbf{p}_{ij}) - (\mathbf{p}_{ij}^T \mathbf{o}_i) \mathbf{R}_{\mathbf{o}_i} \mathbf{e}_z\| = \|(\mathbf{p}_{ij}) - (\mathbf{p}_{ij}^T \mathbf{o}_i) \mathbf{o}_i\| = b,$$

where in the second step we used that the norm is invariant under any rotation in $SO(3)$, and thus we are free to multiply the vector in the norm with $\mathbf{R}_{\mathbf{o}_i}$. In conclusion, the components a and c were uniquely determined the relative group action $(\mathbf{R}_{\mathbf{o}_i}^{-1} \mathbf{p}_{ij}, \mathbf{R}_{\mathbf{o}_i}^{-1} \mathbf{o}_j)$ and b had to be chosen to correspond to the radius of the orbit in the xy plane. This proves that (A5) is bijective, with (??) as inverse mapping.

For all cases, we have now provided valid representatives, and thus for each provided a mapping from attribute to orbit. And thus the provided equivalence class to attribute mappings $[x_i, x_j] \mapsto a_{ij}$ are bijective. \square

B.3 PROOF OF COROLLARY 1.1

Corollary 1 implies that the PΘNITA architectures are equivariant universal approximators. In order to prove this we first show that each of the steps in the architecture is indeed equivariant, and then make use of the results in Dym & Maron (2020); Villar et al. (2021); Gasteiger et al. (2021) to prove universality.

B.3.1 EQUIVARIANCE OF MESSAGE PASSING LAYERS WITH INVARIANT ATTRIBUTES

First, we make precise the notion of equivariance in the point cloud setting (see Sec. A for the bundle vs point cloud viewpoint). Let us consider a point cloud in a space X which we denote with $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathcal{X}, \mathcal{F})$, in which $\mathcal{X} \subset X = \{x_i \mid i \in \mathcal{V}\}$ denotes the set of points associated with the nodes i , and $\mathcal{F} = \{f_i \in V \mid i \in \mathcal{V}\}$ denotes the set node features that live in a vector space V_ρ , with ρ the representation of $SO(n)$ that acts on V_ρ . Let the action of $g = (\mathbf{x}, \mathbf{R}) \in SE(n)$ on the graph be denoted with

$$g\mathcal{G} = (\mathcal{V}, \mathcal{E}, g\mathcal{X}, g\mathcal{F}),$$

with

$$\begin{aligned} g\mathcal{X} &= \{gx_i \mid i \in \mathcal{V}_\downarrow\}, \\ g\mathcal{F} &= \{\rho(\mathbf{R})f_i \mid i \in \mathcal{V}\}. \end{aligned}$$

Then a graph message passing layer MPN is said to be equivariant to the group $SE(n)$ if

$$MPN(g\mathcal{G}) = gMPN(\mathcal{G}).$$

Proposition 1.1. *Message passing layers as described by Eqs 2-4 as conditioned on the attributes in Theorem 1 are $SE(n)$ equivariant.*

Proof. The layers perform a regular group convolution (cf Sec. A) in which case the feature spaces are considered to be vectors of scalars in \mathbb{R}^c on which rotations act trivially, i.e., $\rho(\mathbf{R})f_i = f_i$. Since the positions of the point clouds are un-altered in the message passing steps the output point cloud \mathcal{X}' equals the input point cloud \mathbf{X} , and since the attributes $[gx_i, gx_j] \mapsto a_{ij}$ are invariant for all $g \in SE(n)$, all steps in the message passing scheme are invariant to the group actions. Let us denote the input and output graphs as $G^{in} = (\mathcal{V}, \mathcal{E}, \mathcal{X}, \mathcal{F}^{in})$ and $G^{out} = (\mathcal{V}, \mathcal{E}, \mathcal{X}, \mathcal{F}^{out})$. Then we specifically have that

$$gG^{in} \mapsto gG^{out} \quad \Leftrightarrow \quad (\mathcal{V}, \mathcal{E}, g\mathcal{X}, \mathcal{F}^{in}) \mapsto (\mathcal{V}, \mathcal{E}, g\mathcal{X}, \mathcal{F}^{out}).$$

So, the node features are invariant and the graph itself (point cloud) is not updated by the message passing layer and thus is trivially equivariant. \square

B.3.2 EQUIVARIANCE OF LIFTING THE GRAPH FROM \mathbb{R}^n TO $\mathbb{R}^n \times S^{n-1}$

Now consider a lifting transform that assigns an (approximately) uniform grid of elements in S^{n-1} to each based point in a point cloud in $X = \mathbb{R}^n$. Let such a grid be denoted with $\mathcal{S} \subset S^{n-1}$. For now let us consider the availability of an infinite dimensional grid $\mathcal{S} = S^{n-1}$, as in the fiber bundle viewpoint of Sec. A, and that we are thus able to assign to every point $\mathbf{p} \in \mathcal{X}$ and every $\mathbf{o} \in S^{n-1}$ a feature vector through the point-wise spherical signals $f_{\mathbf{x}} : S^{n-1}$.

We then consider lifting of a graph $G = (\mathcal{V}, \mathcal{E}, \mathcal{X}, \mathcal{F})$ with $\mathcal{X} \subset \mathbb{R}^n$ to a graph $G^\dagger = (\mathcal{V}, \mathcal{E}, \mathcal{X}^\dagger, \mathcal{F}^\dagger)$ via the `ScalarToSphere` and `VecToSphere` modules. The lifted graph G^\dagger can either be seen as a point cloud in $X^\dagger \subset \mathbb{R}^n \times S^{n-1}$ with scalar features $f_i \in \mathbb{R}^c$, or as a (steerable) feature field over the positions in $\mathcal{X} \subset \mathbb{R}^n$ and with spherical signals $f_i \in V_{\rho_{\mathbb{L}_2(S^{n-1})}}^L$ as node features. We denote the latter with $G^\rho = (\mathcal{V}, \mathcal{E}, \mathcal{X}, \mathcal{F}^\rho)$.

We adopt the bundle viewpoint and consider lifting to the graph G^ρ . Scalar features s_i at node $i \in \mathcal{V}$ will be lifted via

$$f_i(\mathbf{o}) = \text{ScalarToSphere}[s_i](\mathbf{o}) = s_i,$$

i.e., it generates constant signals over S^{n-1} with value s_i . Vector features \mathbf{v}_i at node i are lifted to a spherical signal via

$$f_i(\mathbf{o}) = \text{VecToSphere}[\mathbf{v}_i](\mathbf{o}) = \mathbf{v}_i^T \mathbf{o}.$$

Proposition 1.2. *Lifting a graph G to a graph G^ρ via the `VecToSphere` and `ScalarToSphere` operations, is equivariant via*

$$\begin{aligned} (\mathcal{V}, \mathcal{E}, \mathcal{X}, \mathcal{F}) &\mapsto (\mathcal{V}, \mathcal{E}, \mathcal{X}, \mathcal{F}^\rho) \Rightarrow \\ (\mathcal{V}, \mathcal{E}, g\mathcal{X}, g\mathcal{F}) &\mapsto (\mathcal{V}, \mathcal{E}, g\mathcal{X}, g\mathcal{F}^\rho). \end{aligned}$$

That is, the nodes in the lifted graph are equally transformed by the action of $g \in SE(n)$ and the spherical signals $f_i \in \mathbb{L}_2(S^{n-1})$ are permuted via the left-regular representation ρ via $g\mathcal{F}^\rho = \{\rho(g)f_i \mid i \in \mathcal{V}\}$.

Proof. For the `ScalarToSphere` operation we have that the spherical signals $f_i(\mathbf{o}) = s_i$ are invariant since the scalars in the original graph are unaffected by the group action. For the `VecToSphere` operation we have that

$$\mathbf{R}\mathbf{v}_i \mapsto \text{VecToSphere}[\mathbf{R}\mathbf{v}_i](\mathbf{o}) = \mathbf{v}_i^T \mathbf{R}^T \mathbf{o} = (\rho_{\mathbb{L}_2(S^{n-1})}^L(\mathbf{R}) \text{VecToSphere}[\mathbf{v}_i])(\mathbf{o}),$$

and thus a rotation of the input vectors leads to a permutation by $\rho_{\mathbb{L}_2(S^{n-1})}^L(\mathbf{R})$ of the spherical signals. \square

B.3.3 EQUIVARIANCE OF PREDICTING SCALARS AND VECTORS

From the lifted graph G^ρ we can predict at each position $\mathbf{x}_i \in \mathcal{X}$ a scalar or a vector using the `SphereToScalar` and `SphereToVec` modules. In their continuous forms they are given by

$$s_i = \text{SphereToScalar}[f_i] = \int_{S^{n-1}} f_i(\mathbf{o}), \quad (\text{A10})$$

$$\mathbf{v}_i = \text{SphereToVec}[f_i] = \int_{S^{n-1}} f_i(\mathbf{o}) \mathbf{o} \, d\mathbf{o}. \quad (\text{A11})$$

Proposition 1.3. *The process of predicting scalars via Eq. (A10) is invariant and predicting vectors via Eq. (A11) is $SE(n)$ equivariant via*

$$\text{SphereToVec}[\rho(\mathbf{R})f_i] = \mathbf{R} \text{SphereToVec}[f_i].$$

Proof. Since the scalars are invariant to rotations, the `SphereToScalar` module is trivially invariant. For the `SphereToVec` module we have that if the input spherical signals were to be transformed via $\rho(\mathbf{R})$, then we have

$$\begin{aligned} \rho(\mathbf{R})f_i &\mapsto \text{SphereToVec}[\rho(\mathbf{R})f_i] = \int_{S^{n-1}} \rho(\mathbf{R})f_i(\mathbf{o}) \mathbf{o} \, d\mathbf{o} \\ &= \int_{S^{n-1}} f_i(\mathbf{R}^T \mathbf{o}) \mathbf{o} \, d\mathbf{o} \\ &\stackrel{*}{=} \int_{S^{n-1}} f_i(\mathbf{o}) \mathbf{R}\mathbf{o} \, d\mathbf{o} \\ &= \mathbf{R} \int_{S^{n-1}} f_i(\mathbf{o}) \mathbf{o} \, d\mathbf{o} \\ &= \mathbf{R} \text{SphereToVec}[f_i], \end{aligned}$$

where in the step $\stackrel{*}{=}$ we used the change of variables $\mathbf{o} \rightarrow \mathbf{R}\mathbf{o}$. \square

B.4 PROOF OF COROLLARY 1.1

Corollary 1.1 states the following: Message passing networks in geometric graphs over $\mathbb{R}^n \times S^{n-1}$ and $SE(n)$, with message functions conditioned on the attributes of Thm. 1, are equivariant universal approximators. The proof is as follows.

Proof. In the above, we have shown that all the steps in the PΘNITA architecture are equivariant, and it allows us to predict scalar or vector fields. As discussed in Sec. A, the PΘNITA architecture and its layers fit the steerable tensor-field network class of equivariant graph neural networks. Dym & Maron (2020) have proven such types of equivariant graph neural networks to be equivariant universal approximators. As such, (Dym & Maron, 2020, Theorem 2) proves that the bundle implementation of PΘNITA, as described in the main body of this paper, is indeed an equivariant universal approximator. In the case of the point cloud approach (Sec. A.4), universality is given by (Gasteiger et al., 2021, Theorem 3), which is based on the results in (Villar et al., 2021). □

C IMPLEMENTATION

C.1 USED LIBRARY

Our implementations are done in Pytorch (Paszke et al., 2019), using Pytorch-Geometric’s message passing and graph operations modules (Fey & Lenssen, 2019), and made use of WandB (Biewald, 2020) for logging. Our code is available at <https://github.com/ebekkers/ponita>.

C.2 SPHERICAL GRIDS

Uniform grids on S^2 . As described in the main body, we assign a spherical grid to each node in the graph. The grid consists of N points which cover the sphere as uniformly as possible. Note that an exact uniform grid can only be achieved in five cases, corresponding to the five platonic solids with $N = 4, 6, 8, 12, 20$. To achieve an (approximate) uniform grid for any possible N , and thus have full flexibility in specifying angular resolution, we generate grids via a repulsion method. This method randomly initializes points on the sphere and then via gradient descent iteratively pushes them away from each other until convergence. We used the repulsion model code of Kuipers & Bekkers (2023), which was used in their work to generate uniform grids on $SO(3)$.

Random grids. To minimize a potential bias to particular directions in the grid, we randomly perturb the grids with a random rotation matrix in each forward pass of the method. Herein we follow the approach of Knigge et al. (2022) for achieving equivariance to continuous groups, sampled on discrete grids. Every graph in the batch obtains its own randomly rotated spherical grid. This is akin to rotational data augmentation.

Grid resolution We found that a grid resolution of $N = 12$ was generally sufficient for all tasks. However, we got slightly better results with $N = 20$, and thus ran all results with this setting. Increasing beyond $N = 20$ did not seem to improve results.

C.2.1 ATTRIBUTE EMBEDDING

The continuous kernels used in PΘNITA can be considered as Neural Fields. For Neural Fields, it is common to first embed/vectorize the input coordinates. A common approach for this is the Random Fourier Feature embedding, which samples random Fourier basis functions on the provided input coordinates (Tancik et al., 2020). We experimented with this approach as well but found that a simple Polynomial Embedding worked slightly better and was less sensitive to chosen hyperparameters (polynomial degree). The polynomial embedding modules take as input a vector of values (x, y, z, \dots) and generate polynomial combinations of the inputs up to a maximum degree. I.e. $\text{PolynomialEmbed}(x, y, z) = (x, y, z, xy, xz, yz, xxy, xxz, \dots)$.

In all experiments, the kernel used a shared basis which was computed at the start of the forward pass

$$a_{ij} \rightarrow \boxed{\text{PE}} \rightarrow \boxed{\text{Linear}} \rightarrow \boxed{\text{ActFn}} \rightarrow \boxed{\text{Linear}} \rightarrow \boxed{\text{ActFn}} \rightarrow b_{ij},$$

and each group convolution layer then obtains its layer-specific kernel from this shared basis via

$$b_{ij} \rightarrow \boxed{\text{Linear}} \rightarrow k_{ij}^l.$$

In all layers, we used the Gaussian Error Linear Unit (GeLU) as an activation function. In all experiments, we used a 256-dimensional basis.

D EXPERIMENTAL DETAILS

In this section, we provide the hyperparameters used for training the models and provide additional information regarding the benchmarks and the results.

D.1 MD17 EXPERIMENTS

Training settings. The rMD17 results were obtained with PϑNITA and PNITA with $L = 5$ layers, $C = 128$ hidden features. The polynomial degree was set to 3. The models were trained for 5000 epochs, with a batch size of 5. We used the Adam optimizer Kingma & Ba (2014), with a learning rate of $5e-4$, and with a CosineAnnealing learning rate schedule with a warmup period of 50 epochs. We used $N = 20$ grid points on the sphere. The networks optimized the following loss

$$\mathcal{L} = \lambda_E \|\hat{E} - E\|^2 + \lambda_F \frac{1}{3M} \sum_{i=1}^M \sum_{\alpha=1}^3 \left\| -\frac{\partial \hat{E}}{\partial \mathbf{p}_{i,\alpha}} - F_{i,\alpha} \right\|^2 \quad (\text{A12})$$

with $\mathbf{p}_i \in \mathbb{R}^3$. The loss is a weighted sum of energy and force loss terms. Here M is the number of atoms in the system, \hat{E} is the predicted potential energy and λ_E and λ_F are the energy- and force-weightings, respectively. We set $\lambda_F = 500$. The embedding layer is a linear embedding of the one-hot encoded atom numbers, followed by the `ScalarToSphere` module. As explained in the main body, the readout layer was applied after every ConvNeXt block and consisted of a single Linear layer.

D.2 DENOISING DIFFUSION MODELS

D.2.1 SUMMARY OF THE DENOISING DIFFUSION PROBLEM

In diffusion processes, distributions are learned through a reverse diffusion process, i.e. a denoising process. Consider point clouds $\mathbf{x} = (\mathbf{x}_1, \dots, \mathbf{x}_M) \in \mathbb{R}^{M \times 3}$ with corresponding features $\mathbf{h} = (\mathbf{h}_1, \dots, \mathbf{h}_M) \in \mathbb{R}^{M \times \text{nf}}$, the diffusion process that adds noise z_t for time $t = 0, \dots, T$ is given by a multivariate normal distribution:

$$q(z_t | \mathbf{x}) = \mathcal{N}(z_t | \alpha_t \mathbf{x}_t, \sigma_t^2 \mathbf{I}),$$

where $\alpha_t \in \mathbb{R}^+$ controls signal retention and σ_t controls noise added. Sohl-Dickstein et al. (2015); Ho et al. (2020) shows a special case of variance-preserving noising process where $\alpha_t = \sqrt{1 - \sigma_t^2}$. Kingma et al. (2021); Hoogetboom et al. (2022) define signal to noise ratio as $SNR(t) = \frac{\alpha_t^2}{\sigma_t^2}$. Taking into account that the diffusion process is Markov, the entire noising process is then written as:

$$q(z_0, z_1, \dots, z_T | \mathbf{x}) = q(z_0 | \mathbf{x}) \prod_{t=1}^T q(z_t | z_{t-1}).$$

The generative process in a diffusion model is defined with respect to the true denoising process, and the variable x is approximated using a neural network. Following the findings of Ho et al. (2020); Kingma et al. (2021), the variational lower bound on the log-likelihood term of x is given by

$$\mathcal{L}_t = \mathbb{E}_{\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})} \left[\frac{1}{2} (1 - SNR(t-1)/SNR(t)) \|\epsilon - \hat{\epsilon}\|^2 \right]$$

where we predict a gaussian noise and use the parameterization $z_t = \alpha_t \mathbf{x} + \sigma_t \epsilon$, then the neural network ϕ outputs $\hat{\epsilon} = \phi(z_t, t)$.

Köhler et al. (2020) showed that invariant distributions composed with an equivariant invertible function result in an invariant distribution. Additionally, Xu et al. (2022) proved that for $x \sim p(x)$ that is invariant to a group G , and the transition probabilities of a Markov chain defined as $y \sim p(y | x)$ are equivariant, then the marginal distribution of y at any time step is invariant to group transformations. In the case of the denoising diffusion model, we need an invariant distribution and the neural network parameterizing the denoising diffusing process to be equivariant. This results in the marginal distribution of the denoising model to be an invariant distribution. For detailed proof, we refer to Xu et al. (2022).

In order to define the diffusion process, we take $\alpha_t = \sqrt{1 - \sigma_t^2}$ as per Sohl-Dickstein et al. (2015) and let $\alpha_t = (1 - 2s).f(t) + s$ where $f(t) = (1 - (t/T)^2)$, such that values decrease monotonically, starting $\alpha_0 \approx 1$ and $\alpha_T \approx 0$, similar to EDMs for a fair comparison. To avoid numerical instabilities, during sampling, we follow a clipping procedure similar to Dhariwal & Nichol (2021) and compute $\alpha_{t|t-1} = \alpha_t / \alpha_{t-1}$, where $\alpha_{-1} = 1$ and values of $\alpha_{t|t-1}^2$ are clipped from below by .001 similar to Hooeboom et al. (2022).

D.2.2 EXPERIMENTAL DETAILS

Training settings. All hyperparameters were set the same as in the original EDM implementation Hooeboom et al. (2022). That is, we used $L = 9$ layers, $C = 256$ hidden features. The models were trained for 1000 epochs (deviating from the 3000 as the EDM baseline does), with batch size of 64. The polynomial degree of the basis was set to 3. We used the Adam optimizer Kingma & Ba (2014), with a learning rate of $5e - 4$, and without any learning rate scheduler. We used $N = 20$ grid points on the sphere.

The embedding layer is a linear embedding of the one-hot encoded atom numbers, followed by the `ScalarToSphere` module. The readout was a single linear layer, applied to the output of the last `ConvNext` block, and produced a single channel velocity signal on the sphere. This signal was converted to a vector via the `SphereToVec` module.

D.3 N-BODY EXPERIMENT

Training settings. We trained PΘNITA with $L = 5$ layers, $C = 128$ hidden features. The polynomial degree was set to 3. The models were trained for 10,000 epochs, with batch size of 100. We used the Adam optimizer Kingma & Ba (2014), with a learning rate of $5e - 4$, and with a CosineAnnealing learning rate schedule with a warmup period of 100 epochs. We used $N = 20$ grid points on the sphere.

The embedding layer took two vector features as inputs initial velocity and a direction vector pointing from the particle position to the average position of all particles. These two vectors were embedded as two scalar fields over the sphere via the `VecToSphere` module. The embedding also took two scalar features as input: charge (-1 or +1) and the norm of the initial velocity. Both were embedded on the sphere via the `ScalarToSphere` module. The 4 features were mapped to a vector of size C (the hidden dimension size). The readout layer was applied after every `ConvNext` block, and consisted of a Linear layer followed by the `SphereToVec` module. The final predicted velocity was the average of the velocity predictions after each layer.

In addition to the invariant geometric attributes, the basis functions also took the product the charges of the sending and receiving node as input.

E ADDITIONAL EXPERIMENTS

E.1 SUPER-PIXEL MNIST

To verify the impact of lifting point clouds to position orientation space in 2D, we benchmark PΘNITA on superpixel mnist. Superpixel mnist is a dataset of 2D point clouds of MNIST digits, in which each of the original images is segmented into 75 superpixels (clusters of similar pixels). We use the data loader as provided by the library `torch-geometric` Fey & Lenssen (2019). We utilize the exact same architecture as in all the other experiments presented so far, except that now we utilize the three 2D invariants as given in (8).

We used the following hyper parameters. We trained for 50 epochs with a batch size of 96, a cosine learning rate decay, and initial learning rate of $5e - 4$. For P Θ NITA we used $N = 10$ grid points to sample the circle.

We compare our results with PNITA and P Θ NITA in Table 4 against the recent state-of-the-art method of Finzi et al. (2020b) called Probabilistic Numeric Convolutional Neural Networks (PNCNN), as well as classic graph NN approaches MONET (Monti et al., 2017), SplineCNN (Fey et al., 2018), GCCP (Walker & Glocker, 2019), and GAT (Avelar et al., 2020). Our results are averaged over 3 runs with different seeds and are reported with the standard deviation in test error over these runs.

Apart from showing that with P Θ NITA state-of-the-art results are obtained with the simple group convolutional architecture P Θ NITA, the results in Table 4 also show the significant benefit of utilizing group convolutions (P Θ NITA) over planar convolutions with isotropic kernels (PNITA). The results further show that a simple fully convolutional architecture is sufficient to obtain outstanding results.

Table 4: Classification Error on the 75-SuperPixel MNIST problem.

Method	Error rate
MONET	8.89
SplineCNN	4.78
GCCP	4.2
GAT	3.81
PNCNN	$1.24 \pm .12$
PNITA	$3.04 \pm .09$
P Θ NITA	$1.17 \pm .11$

Table 5: QM9 regression results.

Target	Unit	DimeNet++	PNITA	P Θ NITA (point cloud)	P Θ NITA (fiber bundle)
μ	D	0.0286	0.0207	0.0115	0.0121
α	a_0^3	0.0469	0.0602	0.0446	0.0375
ϵ_{HOMO}	meV	27.8	26.1	18.6	16.0
ϵ_{LUMO}	meV	19.7	21.9	15.3	14.5
$\Delta\epsilon$	meV	34.8	43.4	33.5	30.4
$\langle R^2 \rangle$	a_0^2	0.331	0.149	0.227	0.235
ZPVE	meV	1.29	1.53	1.29	1.29
U_0	meV	8.02	10.71	9.20	8.31
U	meV	7.89	10.63	9.00	8.67
H	meV	8.11	11.00	8.54	8.04
G	meV	0.0249	0.0112	0.0095	0.00863
c_v	$\frac{\text{cal}}{\text{molK}}$	0.0249	0.0307	0.0250	0.0242

E.2 QM9 REGRESSION

To further demonstrate the potency of P Θ NITA on 3D point cloud tasks we validate it on the QM9 benchmark (Ramakrishnan R., 2014) on the task of predicting molecular properties from input atomic point clouds and their covalent bonds.

Since QM9 provides a connectivity derived from the covalent bonds, we have an option to treat the molecules as point clouds in position-orientation space. Namely, we can treat every edge, which connects neighboring point \mathbf{x}_j to central node \mathbf{x}_i , as a point in position-orientation space $\mathbb{R}^3 \times S^2$ by treating taking the starting point of the edge as position and the normalized direction vector as orientation. That is, for every edge $e = (j, i) \in \mathcal{E}$ we have identify a point $x_e = (\mathbf{p}_i, \frac{\mathbf{p}_j - \mathbf{p}_i}{\|\mathbf{p}_j - \mathbf{p}_i\|}) \in \mathbb{R}^3 \times S^2$. The covalent bonds thus naturally generate a point-cloud in position-orientation space.

Since this point cloud no-longer has a fiber bundle structure (see App. A) we can no longer separate the convolution over positions and orientations separately and so we do the "domain-mixing" in one single convolution, followed by a channel mixing linear layer. So now the convolution is separated over 2 steps, instead of 3 as for the other experiments, as described in Section 3.2. Otherwise, the PΘNITA architecture is held exactly the same. In this experiment we compare convolutions using either the atomic point cloud as input (PNITA) versus using the covalent-bond cloud (edges) as input (PΘNITA). The latter is close to the approach of DimeNet and GemNet (Gasteiger et al., 2019; 2021) which also perform message passing operations between edges. We therefore compare our approach against DimeNet++ (Gasteiger et al., 2020).

We did a hyperparameter search for C and L for PNITA and PΘNITA separately and found that for PNITA best results were obtained with $C = 128$ and $L = 5$ and for PΘNITA with $C = 256$ and $L = 9$. We further trained for 1000 epochs with a batch size of 96, with a cosine learning rate decay scheduler and initial learning rate of $5e - 4$. For all atoms we constructed a fully connected graph. For reference, we also trained a fiber bundle PΘNITA with $C = 256$, $L = 9$, and $N = 16$.

The results are presented in Table 5.