
Equivariant Maps for Hierarchical Structures

Renhao Wang, Marjan Albooyeh*
University of British Columbia,
{renhaow, albooyeh}@cs.ubc.ca

Siamak Ravanbakhsh
McGill University & Mila
siamak@cs.mcgill.ca

Abstract

While using invariant and equivariant maps, it is possible to apply deep learning to a range of primitive data structures, a formalism for dealing with hierarchy is lacking. This is a significant issue because many practical structures are hierarchies of simple building blocks; some examples include sequences of sets, graphs of graphs, or multiresolution images. Observing that the symmetry of a hierarchical structure is the “wreath product” of symmetries of the building blocks, we express the equivariant map for the hierarchy using an intuitive combination of the equivariant linear layers of the building blocks. More generally, we show that *any* equivariant map for the hierarchy has this form. To demonstrate the effectiveness of this approach to model design, we consider its application in the semantic segmentation of point-cloud data. By voxelizing the point cloud, we impose a hierarchy of translation and permutation symmetries on the data and report state-of-the-art on SEMANTIC3D, S3DIS, and VKITTI, that include some of the largest real-world point-cloud benchmarks.

1 Introduction

In designing deep models for structured data, equivariance (invariance) of the model to transformation groups has proven to be a powerful inductive bias, which enables sample efficient learning. A widely used family of equivariant deep models constrain the feed-forward layer so that specific transformations of the input lead to the corresponding transformations of the output. A canonical example is the convolution layer, in which the constrained MLP is equivariant to translation operations. Many recent works have extended this idea to design equivariant networks for more exotic structures such as sets, exchangeable tensors and graphs, as well as relational and geometric structures.

This paper considers a nested hierarchy of such structures, or more generally, any hierarchical composition of transformation symmetries. These hierarchies naturally appear in many settings: for example, the interaction between nodes in a social graph may be a sequence or a set of events. Or in diffusion tensor imaging of the brain, each subject may be modeled as a set of sequences, where each sequence is a fibre bundle in the brain. The application we consider in this paper models point clouds as 3D images, where each voxel is a set of points with coordinates relative to the center of that voxel.

To get an intuition for a hierarchy of symmetry transformations, consider the example of a sequence of sequences – *e.g.*, a text document can be viewed as a sequence of sentences, where each sentence is itself a sequence of words. Here, each inner sequence as well as the outer sequence is assumed to possess an “independent” translation symmetry. Contrast this with symmetries of an image (2D translation), where all inner sequences (say row pixels) translate together, so we have a total of two translations. This is the key difference between the *wreath product* of two translation groups (former) and their *direct product* (latter). It is the wreath product that often appears in nested structures. As is evident from this example, the wreath product results in a significantly larger set of transformations, and as we elaborate later, it provides a stronger inductive bias.

*Currently at Borealis AI. Work done while at UBC.

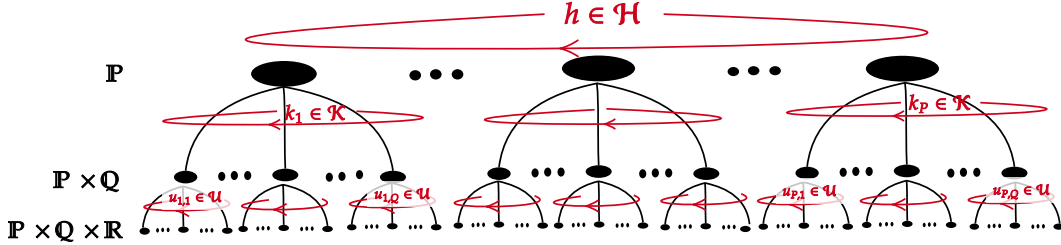


Figure 1: Wreath product can express the symmetries of hierarchical structures: wreath product of three groups $\mathcal{U} \wr \mathcal{K} \wr \mathcal{H}$ acting on the set of elements $\mathbb{P} \times \mathbb{Q} \times \mathbb{R}$, can be seen as *independent copies* of groups, \mathcal{H} , \mathcal{K} and \mathcal{U} at different level of hierarchy acting \wr on copies of \mathbb{P} , \mathbb{Q} and \mathbb{R} . Intuitively, a linear map $\mathbf{W}_{\mathcal{U} \wr \mathcal{K} \wr \mathcal{H}} : \mathbb{R}^{PQR} \rightarrow \mathbb{R}^{PQR}$ equivariant to $\mathcal{U} \wr \mathcal{K} \wr \mathcal{H}$, performs pooling over leaves under each inner node \bullet , applies equivariant map for each inner structure (*i.e.*, $\mathbf{W}_{\mathcal{U}}$, $\mathbf{W}_{\mathcal{K}}$ and $\mathbf{W}_{\mathcal{H}}$ respectively), and broadcasting the output back to the leaves. A $\mathcal{U} \wr \mathcal{K} \wr \mathcal{H}$ -equivariant map can be constructed as the sum of these three contributions, from equivariant maps at each level of hierarchy.

We are interested in application of equivariant/invariant deep learning to this type of nested structure. The building blocks of equivariant and invariant MLPs are *equivariant linear maps* of the feedforward layer. We show that any equivariant linear map for the hierarchical structure is built using equivariant maps for the individual symmetry group at different levels of the hierarchy. Our construction only uses additional pooling and broadcasting operations; see Fig. 1.

In the following, after discussing related works in Section 2, we give a short background on equivariant MLPs in Section 3. Section 4 starts by giving the closed form of equivariant maps for direct product of groups before moving to the more difficult case of wreath product in Section 4.2. Finally, Section 5 applies this idea to impose a hierarchical structure on 3D point clouds. We show that the equivariant map for this hierarchical structure achieves state-of-the-art performance on the largest benchmark datasets for 3D semantic segmentation.

2 Related Works

Group theory has a long history in signal processing [22], where in particular Fourier transformation and group convolution for Abelian groups have found tremendous success over the past decades. However, among non-commutative groups, wreath product constructions have been the subject of few works. Rockmore [39] give efficient procedures for fast Fourier transforms for wreath products. In a series of related works Foote et al. [15], Mirchandani et al. [34] investigate the wreath product for multi-resolution signal processing. The focus of their work is on the wreath product of cyclic groups for compression and filtering of image data.

Group theory has also found many applications in machine learning [25], and in particular deep learning. Design of invariant MLPs for general groups goes back to Shawe-Taylor [41]. More recently, several works investigate the design of equivariant networks for general finite [38] and infinite groups [7, 10, 26]. In particular, use of the wreath product for design of networks equivariant to a hierarchy of symmetries is briefly discussed in [38]. Equivariant networks have found many applications in learning on various structures, from image and text [31], to sets [36, 50], exchangeable matrices [21], graphs [1, 28, 32], and relational data [18], to signals on spheres [9, 27]. A large body of works investigate equivariance to Euclidean isometries; *e.g.*, [11, 44, 46].

When it comes to equivariant models for compositional structures, contributions have been sparse, with most theoretical work focusing on semidirect product (or more generally using induced representations from a subgroup) to model tensor-fields on homogeneous spaces [8, 10]. Direct product of symmetric groups have been used to model interactions across sets of entities [21] and in generalizations to relational data [18].

Relation to Maron et al. [33] Maron et al. [33] extensively study equivariant networks for direct product of groups; in contrast we focus on wreath product. Since both of these symmetry transformations operate on the Cartesian product of \mathcal{G} -sets (see Fig. 1), the corresponding permutation groups are comparable. Indeed, direct product action is a sub-group of the imprimitive action of the

wreath product. The implication is that if sub-structures in the hierarchy transform independently (e.g., in sequence of sequences), then using the equivariant maps proposed in this paper gives a strictly stronger inductive bias. However, if sub-structures move together (e.g., in an image), then using a model equivariant to direct product of groups is preferable. One could also compare the learning bias of these models by noting that when using wreath product, the number of independent linear operators grows with the “sum” of independent operators for individual building blocks; in contrast, this number grows with the “product” of independent operators on blocks when using direct product. In the following we specifically contrast the use of direct product with wreath product for compositional structures.

3 Preliminaries

3.1 Group Action

A group $\mathcal{G} = \{g\}$ is a set equipped with a binary operation, such that the set is closed under this operation, $gh \in \mathcal{G}$, the operation is associative, $g(hk) = (gh)k$, there exists identity element $e \in \mathcal{G}$, and a unique inverse for each $g \in \mathcal{G}$ satisfying $gg^{-1} = e$. The action of \mathcal{G} on a finite set \mathbb{N} is a function $\alpha : \mathcal{G} \times \mathbb{N} \rightarrow \mathbb{N}$ that transforms the elements of \mathbb{N} , for each choice of $g \in \mathcal{G}$; for short we write $g \cdot n$ instead of $\alpha(g, n)$. Group actions preserve the group structure, meaning that the transformation associated with the identity element is identity $e \cdot n = n$, and composition of two actions is equal to the action of the composition of group elements $(gh) \cdot n = g \cdot (h \cdot n)$. Such a set, with a \mathcal{G} -action defined on it is called a \mathcal{G} -set. The group action on \mathbb{N} naturally extends to $\mathbf{x} \in \mathbb{R}^{|\mathbb{N}|}$, where it defines a permutation of indices $g \cdot (x_1, \dots, x_N) \doteq (x_{g^{-1} \cdot 1}, \dots, x_{g^{-1} \cdot N})$. We often use a permutation matrix $\mathbf{G}^{(g)} \in \{0, 1\}^{N \times N}$ to represent this action – that is $\mathbf{G}^{(g)}\mathbf{x} = g \cdot \mathbf{x}$.

3.2 Equivariant Multilayer Perceptrons

A function $\phi : \mathbb{R}^N \rightarrow \mathbb{R}^M$ is equivariant to a given actions of group \mathcal{G} iff $\phi(\mathbf{G}^{(g)}\mathbf{x}) = \tilde{\mathbf{G}}^{(g)}\phi(\mathbf{x})$ for any $\mathbf{x} \in \mathbb{R}^N$ and $g \in \mathcal{G}$. That is, a symmetry transformation of the input results in the corresponding symmetry transformation of the output. Note that the action on the input and output may in general be different. In particular, when $\tilde{\mathbf{G}}^{(g)} = \mathbf{I}_M$ for all g – that is, the action on the output is trivial – equivariance reduces to invariance. Here, \mathbf{I}_M is the $M \times M$ identity matrix. For simplicity and motivated by practical design choices, we assume the same action on the input and output.

For a feedforward layer $\phi : \mathbf{x} \mapsto \sigma(\mathbf{W}\mathbf{x})$, where σ is a point-wise non-linearity and $\mathbf{W} \in \mathbb{R}^{N \times N}$, the equivariance condition above simplifies to commutativity condition $\mathbf{G}^{(g)}\mathbf{W} = \mathbf{W}\mathbf{G}^{(g)} \forall g \in \mathcal{G}$. This imposes a symmetry on \mathbf{W} in the form of parameter-sharing [38, 47]. While we can use computational means to solve this equation for any finite group, an efficient implementation requires a closed form solution. Several recent works derive the closed form solutions for interesting groups and structures. Note that in general the feedforward layer may have multiple input and output channels, with identical \mathcal{G} -action on each channel. This only require replicating the parameter-sharing pattern in \mathbf{W} , for each combination of input and output channel. An Equivariant MLP is a stack of equivariant feed-forward layers, where the composition of equivariant layers is also equivariant. Therefore, our task in building MLPs equivariant to finite group actions is reduced to finding equivariant linear maps in the form of parameter-sharing matrices satisfying $\mathbf{G}^{(g)}\mathbf{W} = \mathbf{W}\mathbf{G}^{(g)} \forall g \in \mathcal{G}$; see Fig. 4(a.1,a.2) for \mathbf{W} that are equivariant to \mathcal{C}_4 , the group of circular translations of 4 objects (a.1), and symmetric group \mathcal{S}_4 , the group of all permutations of 4 objects (a.2).

4 Equivariant Map for Product Groups

In this section we formalize the *imprimitive action* of the wreath product, which is used in describing the symmetries of hierarchical structures. We then introduce the closed form of linear maps equivariant to the wreath product of two groups. With hierarchies of more than two levels, one only needs to iterate this construction. To put this approach in perspective and to make the distinction clear, first we present the simpler case of direct product; see [33] for an extensive discussion of this setting.

4.1 Equivariant Linear Maps for Direct Product of Groups

The easiest way to combine two groups is through their direct product $\mathcal{G} = \mathcal{H} \times \mathcal{K}$. Here, the underlying set is the Cartesian product of the input sets and group operation is $(h, k)(h', k') \doteq (hh', kk')$. If \mathbb{P} is an \mathcal{H} -set and \mathbb{Q} a \mathcal{K} -set then the group $\mathcal{G} = \mathcal{H} \times \mathcal{K}$ naturally acts on $\mathbb{N} = \mathbb{P} \times \mathbb{Q}$ using $(h, k) \cdot (p, q) \doteq (h \cdot p, k \cdot q)$; see Fig. 2. This type of product is useful in modeling the Cartesian product of structures. The following claim characterises the equivariant map for direct product of two groups using the equivariant map for building blocks.

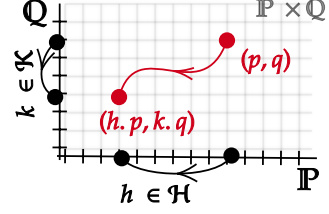


Figure 2: Direct product action.

Claim 1. Let $\mathbf{G}^{(h)}$ represent \mathcal{H} -action, and let $\mathbf{W}_{\mathcal{H}} \in \mathbb{R}^{P \times P}$ be an equivariant linear map for this action. Similarly, let $\mathbf{W}_{\mathcal{K}} : \mathbb{R}^{Q \times Q}$ be equivariant to \mathcal{K} -action given by $\mathbf{G}^{(k)}$ for $k \in \mathcal{K}$. Then, the product group $\mathcal{G} = \mathcal{H} \times \mathcal{K}$ naturally acts on $\mathbb{R}^N = \mathbb{R}^{P \times Q}$ using $\mathbf{G}^{(g)} = \mathbf{G}^{(h)} \otimes \mathbf{G}^{(k)}$, and the Kronecker product $\mathbf{W}_{\mathcal{G}} = \mathbf{W}_{\mathcal{H}} \otimes \mathbf{W}_{\mathcal{K}}$, is a \mathcal{G} -equivariant linear map.

Note that the claim is not restricted to permutation action.² The proof follows from the *mixed-product property* of the Kronecker product³, and the equivariance of $\mathbf{W}_{\mathcal{H}}$ and $\mathbf{W}_{\mathcal{K}}$:

$$\begin{aligned} \mathbf{G}^{(g)} \mathbf{W}_{\mathcal{G}} &= (\mathbf{G}^{(h)} \otimes \mathbf{G}^{(k)}) (\mathbf{W}_{\mathcal{H}} \otimes \mathbf{W}_{\mathcal{K}}) = (\mathbf{G}^{(h)} \mathbf{W}_{\mathcal{H}}) \otimes (\mathbf{G}^{(k)} \mathbf{W}_{\mathcal{K}}) \\ &= (\mathbf{W}_{\mathcal{H}} \mathbf{G}^{(h)}) \otimes (\mathbf{W}_{\mathcal{K}} \mathbf{G}^{(k)}) = (\mathbf{W}_{\mathcal{H}} \otimes \mathbf{W}_{\mathcal{K}}) (\mathbf{G}^{(h)} \otimes \mathbf{G}^{(k)}) = \mathbf{W}_{\mathcal{G}} \mathbf{G}^{(g)} \quad \forall g \in \mathcal{G}. \end{aligned}$$

An implication of the tensor product form of $\mathbf{W}_{\mathcal{H} \times \mathcal{K}}$ is that the number of independent linear operators of the product map (free parameters in the parameter-sharing) is the product of the independent operators of the building blocks. Note that whenever dealing with product of inputs belonging to \mathbb{R}^P and \mathbb{R}^Q , the result which belongs to $\mathbb{R}^{P \times Q}$ and therefore it is vectorized.

Example 1 (Convolution in Higher Dimensions). *D-dimensional convolution is a Kronecker (tensor) product of one-dimensional convolutions. The number of parameters grows with the product of kernel width across all dimensions; Fig. 4(a.1) shows the parameter-sharing for circular 1D convolution \mathbf{W}_{C_4} , and (b.1) shows the parameter-sharing for the direct product $\mathbf{W}_{C_3 \times C_4}$.*

Example 2 (Exchangeable Tensors). *Hartford et al. [21] introduce a layer for modeling interactions across multiple sets of entities, e.g., a user-movie rating matrix. Their model can be derived as the Kronecker product of 2-parameter equivariant layer for sets [50]. The number of parameters is therefore 2^D for a rank D exchangeable tensor. Fig. 4(b.1) shows the 2-parameter model \mathbf{W}_{S_4} for sets, and (c.1) shows the parameter-sharing for the direct product $\mathbf{W}_{S_3 \times S_4}$.*

4.2 Wreath Product Action and Equivariance to a Hierarchy of Symmetries

Let us start with an informal definition. Suppose as before \mathbb{P} and \mathbb{Q} are respectively an \mathcal{H} -set and a \mathcal{K} -set. We can attach one copy of \mathbb{Q} to each element of \mathbb{P} . Each of these inner sets or fibers have their own copy of \mathcal{K} acting on them. Action of \mathcal{H} on \mathbb{P} simply permutes these fibers. Therefore the combination of all \mathcal{K} actions on all inner sets combined with \mathcal{H} -action on the outer set defines the action of the wreath product on $\mathbb{P} \times \mathbb{Q}$. Fig. 3 demonstrates how one point (p, q) moves under this action. Next few paragraphs formalize this.

Semidirect Product. Formally, wreath product is defined using *semidirect* product which is a generalization of direct product. In part due to its use in building networks equivariant to Euclidean isometries, application of semidirect product in building equivariant networks is explored in several recent works; see [10] and citations therein. In semidirect product, the underlying set (of group members) is again the product set. However the group operation is more involved. The (external) semi-direct product $\mathcal{G} = \mathcal{K} \rtimes_{\gamma} \mathcal{H}$, requires a *homomorphism* $\gamma : \mathcal{H} \rightarrow \text{Aut}(\mathcal{K})$ that for each choice of

²The tensor product of irreducible representation of two distinct finite groups is an irreducible representation for the product group. Therefore this construction of equivariant maps can be used with a decomposition into irreducible representations for the general linear case.

³ $(\mathbf{A} \otimes \mathbf{B})(\mathbf{C} \otimes \mathbf{D}) = (\mathbf{AC}) \otimes (\mathbf{BD})$

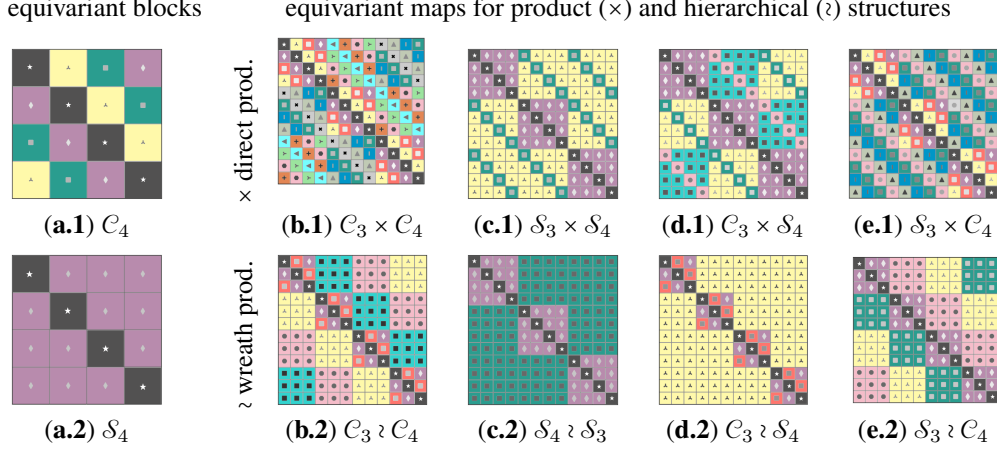


Figure 4: Parameter-sharing pattern in the equivariant linear maps for compositional structures. Equivariant maps for sequence and set are given in the first column (a.1, a.2). The first row shows the maps equivariant to various *products* of sets and sequences, while the second row is the corresponding map for *nested structure*. **Product Structures:** (b.1) product of sequences (as in image); (c.1) product of sets (exchangeable matrices); (d.1) product of a set and a sequence; (e.1) product of a sequence and a set. In the notation, the outer structure appears first in $\mathcal{H} \times \mathcal{K}$. **Hierarchical Structures:** (b.2) sequence of sequences (e.g., multi-resolution sequence model); (c.2) nested set of sets; (d.2) set of sequences; (e.2) sequences of sets (similar to the model used in the application section.) In wreath product notation, the outer structure appears *second* in $\mathcal{K} \wr \mathcal{H}$. Note how the equivariant map for the hierarchy is a scaled version of the map for the outer structure with copies of the inner structure maps appearing as diagonal blocks.

$h \in \mathcal{H}$, re-labels the elements of \mathcal{K} while preserving its group structure. Using γ , the binary operation for the product group $\mathcal{G} = \mathcal{K} \rtimes_{\gamma} \mathcal{H}$ is defined as $(h, k)(h', k') = (hh', k\gamma_h(k'))$. A canonical example is the semidirect product of translations (\mathcal{K}) and rotations (\mathcal{H}), which identifies the group of all rigid motions in the Euclidean space. Here, each rotation defines an automorphism of translations (e.g., moving north becomes moving east after 90° clockwise rotation).

Now we are ready to define the wreath product of two groups. As before, let \mathcal{H} and \mathcal{K} denote two finite groups, and let \mathbb{P} be an \mathcal{H} -set. Define \mathcal{B} as the direct product of $P = |\mathbb{P}|$ copies of \mathcal{K} , and index these copies by $p \in \mathbb{P}$: $\mathcal{B} = \mathcal{K}_1 \times \dots \times \mathcal{K}_p \times \dots \times \mathcal{K}_P$. Each member of this group is a tuple $b = (k_1, \dots, k_p, \dots, k_P)$. Since \mathbb{P} is an \mathcal{H} -set, \mathcal{H} also naturally acts on \mathcal{B} by permuting the *fibers* \mathcal{K}_p . The semidirect product $\mathcal{B} \rtimes \mathcal{H}$ defined using this automorphism of \mathcal{B} is called the wreath product, and written as $\mathcal{K} \wr \mathcal{H}$. Each member of the product group can be identified by the pair (h, b) , where b as a member of the *base* group itself is a P -tuple. This shows that the order of the wreath product group is $|\mathcal{K}|^P |\mathcal{H}|$ which can be much larger than the direct product group $\mathcal{K} \times \mathcal{H}$, whose order is $|\mathcal{K}| |\mathcal{H}|$.

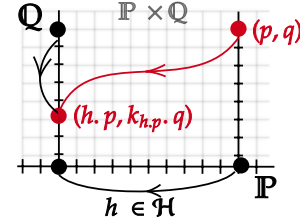


Figure 3: Imprimitive action of wreath product.

4.2.1 Imprimitive Action of Wreath Product

If in addition to \mathbb{P} being an \mathcal{H} -set, \mathcal{K} action on \mathbb{Q} is also defined, the wreath product group acts on $\mathbb{P} \times \mathbb{Q}$ (making it comparable to the direct product). Specifically, $(h, k_1, \dots, k_P) \in \mathcal{K} \wr \mathcal{H}$ acts on $(p, q) \in \mathbb{P} \times \mathbb{Q}$ as follows:

$$(h, k_1, \dots, k_P) \cdot (p, q) \doteq (h \cdot p, k_{h.p} \cdot q).$$

Intuitively, \mathcal{H} permutes the copies of \mathcal{K} acting on each \mathbb{Q} , and itself acts on \mathbb{P} . We can think of \mathbb{P} as the outer structure and \mathbb{Q} as the inner structure; see Fig. 3.

Example 3 (Sequence of Sequences). Consider our early example where both $\mathcal{H} = \mathcal{C}_P$ and $\mathcal{K} = \mathcal{C}_Q$ are cyclic groups with regular action on $\mathbb{P} \cong \mathcal{H}$, $\mathbb{Q} \cong \mathcal{K}$. Each member of the wreath product $\mathcal{C}_Q \wr \mathcal{C}_P$, acts by translating each inner sequence using some $c \in \mathcal{C}_Q$, while $c' \in \mathcal{C}_P$ translates the outer sequence by c' .

Let the $P \times P$ permutation matrix $\mathbf{G}^{(h)}$ represent the action of $h \in \mathcal{H}$, and the $Q \times Q$ matrices $\mathbf{G}^{(h_1)}, \dots, \mathbf{G}^{(h_P)}$ represent the action of h_p on \mathbb{Q} . Then the action of $g \in \mathcal{K} \wr \mathcal{H}$ on (the vectorized) $\mathbb{P} \times \mathbb{Q}$ is the following $PQ \times PQ$ permutation matrix:

$$\mathbf{G}^{(g)} = \begin{bmatrix} \mathbf{G}_{1,1}^{(h)} \mathbf{G}^{(h_1)}, & \dots, & \mathbf{G}_{1,P}^{(h)} \mathbf{G}^{(h_1)} \\ \vdots & \ddots & \vdots \\ \mathbf{G}_{P,1}^{(h)} \mathbf{G}^{(h_P)}, & \dots, & \mathbf{G}_{P,P}^{(h)} \mathbf{G}^{(h_P)} \end{bmatrix} = \sum_{p=1}^P \mathbf{1}_{p,h \cdot p} \otimes \mathbf{G}^{(h_p)} \quad (1)$$

where $\mathbf{1}_{p,h \cdot p}$ is a $P \times P$ matrix whose only nonzero element is at row p and column $h \cdot p$ with the value of 1; and $\mathbf{G}_{p,p'}^{(h)}$ is the element at row p and column p' of the permutation matrix $\mathbf{G}^{(h)}$. In the summation formulation, the Kronecker product $\mathbf{1}_{p,h \cdot p} \otimes \mathbf{G}^{(h_p)}$ puts a copy of permutation matrix $\mathbf{G}^{(h_p)}$ as a block of $\mathbf{G}^{(g)}$. Note that the resulting permutation matrix is different from the Kronecker product; in this case we have $P + 1$ (permutation) matrices participating in creating $\mathbf{G}^{(g)}$, compared with two matrices in the vanilla Kronecker product; see Fig. 5.

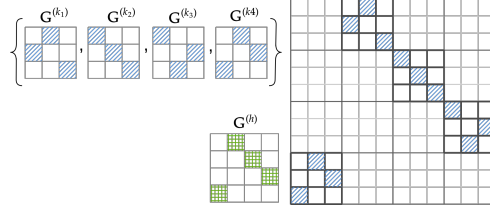


Figure 5: Permutation for the imprimitive wreath product ($\mathbf{G}^{(g)}$) on the right, is built from permutations of the outer structure using $\mathbf{G}^{(h)}$ and independent permutations of inner structures $\mathbf{G}^{(h_1)}, \dots, \mathbf{G}^{(h_4)}$.

4.2.2 Equivariant Map

Consider a hierarchical structure, potentially with more than two levels of hierarchy, such as a set of sequences of images. Moreover, suppose that we have an equivariant map for each individual structure. The question answered by the following theorem is: how to use the equivariant map for each level to construct the equivariant map for the entire hierarchy? For now we only consider two levels of hierarchy; extension to more levels follows naturally, and is discussed later. Note that one instance of the hierarchy (e.g., a sequence of sets) belongs to \mathbb{R}^{PQ} , because each instance of the structure is in \mathbb{R}^Q and the outer structure is of size P – that is vectorization of the input follows the hierarchy.

Theorem 4.1. Let $\mathbf{W}_{\mathcal{K} \wr \mathcal{H}} \in \mathbb{R}^{PQ \times PQ}$ be the matrix of some linear map equivariant to the imprimitive action of $\mathcal{K} \wr \mathcal{H}$ on $\mathbb{P} \times \mathbb{Q}$. Any such map can be written as follows

$$\mathbf{W}_{\mathcal{K} \wr \mathcal{H}} = \mathbf{W}_{\mathcal{H}} \otimes (\mathbf{1}_Q \mathbf{1}_Q^\top) + \mathbf{I}_P \otimes \mathbf{W}_{\mathcal{K}}, \quad \text{where } \mathbf{1}_Q = \overbrace{[1, \dots, 1]}^{Q \text{ times}}, \quad (2)$$

\mathbf{I}_P is the $P \times P$ identity matrix, $\mathbf{W}_{\mathcal{H}} \in \mathbb{R}^{P \times P}$ is \mathcal{H} -equivariant, and $\mathbf{W}_{\mathcal{K}} \in \mathbb{R}^{Q \times Q}$ is equivariant to \mathcal{K} -action.

Proof is in Appendix A. Pictorially, the first term of Eq. (2) scales up the matrix $\mathbf{W}_{\mathcal{H}}$, while the second term adds copies of $\mathbf{W}_{\mathcal{K}}$ on the diagonal blocks of the scaled matrix. Assuming sets of independent equivariant linear operators $\{\mathbf{W}_{\mathcal{K}}\}$ and $\{\mathbf{W}_{\mathcal{H}}\}$, from Eq. (2) it is evident that the number of independent linear operators equivariant to $\mathcal{K} \wr \mathcal{H}$ grows with the sum of those of the building blocks.⁴ These operators may be combined using any parameter to create a parameterized linear map, which in turn can be expressed using parameter-sharing in a vanilla feed-forward layer. This in contrast with direct product in which the number of free parameters has a product form.

Example 4 (Various Hierarchies of Sets and Sequences). Consider two of the most widely used equivariant maps, translation equivariant convolution $\mathbf{W}_{\mathcal{C}_P}$, and \mathcal{S}_Q -equivariant map which has the form $\mathbf{W}_{\mathcal{S}_Q} = w_1 \mathbf{I}_Q + w_2 [50]$. There are four combinations of these structures in a two level hierarchical structure: 1) set of sets $\mathcal{S}_Q \wr \mathcal{S}_P$; 2) sequence of sequences $\mathcal{C}_Q \wr \mathcal{C}_P$; 3) set of sequences $\mathcal{C}_Q \wr \mathcal{S}_P$; 4) sequence of sets $\mathcal{S}_Q \wr \mathcal{C}_P$. Fig. 4(b.2-e.2) show the parameter-sharing matrix for these hierarchical structures, assuming a full kernel in 1D circular convolution.

⁴In the proof we show that the number of parameters is the sum of those of the building blocks minus one.

The reader is invited to contrast the three parameter layer for set of sets, with the four parameter layer for interactions across sets in Fig. 4(c.1,c.2). Similarly, a model for sequence of sequences has far fewer parameters than a model for an image, as seen in Fig. 4(b.1, b.2).

4.2.3 Deeper Hierarchies and Combinations with Direct Product

With more than two levels, the symmetry group involves more than one wreath product, which means that the equivariant map for the hierarchy is given by a recursive application of Theorem 4.1. For example, the equivariant map for $(\mathcal{K} \wr \mathcal{H}) \wr \mathcal{U}$, in which \mathcal{U} acts on some set \mathbb{R} , and $\mathbf{W}_{\mathcal{U}}$ is the corresponding equivariant map, is given by $\mathbf{W}_{(\mathcal{K} \wr \mathcal{H}) \wr \mathcal{U}} = \mathbf{W}_{\mathcal{U}} \otimes (\mathbf{1}_{PQ} \mathbf{1}_{PQ}^{\top}) + \mathbf{I}_R \otimes \mathbf{W}_{\mathcal{K} \wr \mathcal{H}}$, where $\mathbf{W}_{\mathcal{K} \wr \mathcal{H}}$ is in turn given by Eq. (2). Note that wreath product is associative, and so the iterative construction above leads the same equivariant map as $\mathbf{W}_{\mathcal{K} \wr (\mathcal{H} \wr \mathcal{U})}$.

We can also mix and match this construction with that of direct product in Claim 1; for example, to produce the map for exchangeable tensors (product of sets), where each interaction is in the form of an image (hierarchy) – *i.e.*, the group is $(\mathcal{C}_R \times \mathcal{C}_V) \wr (\mathcal{S}_P \times \mathcal{S}_Q)$.

4.3 Efficient Implementation

With equivariant maps for direct product of groups, efficient implementation of $\mathbf{W}_{\mathcal{K} \times \mathcal{H}}$ using efficient implementation for individual blocks $\mathbf{W}_{\mathcal{K}}$, and $\mathbf{W}_{\mathcal{H}}$ is non-trivial – *e.g.*, consider 2D convolution. In contrast, it is possible to use a black-box implementation of the parameter-sharing layers for $\mathbf{W}_{\mathcal{K}}$, and $\mathbf{W}_{\mathcal{H}}$ to construct $\mathbf{W}_{\mathcal{K} \times \mathcal{H}}$. To this end, let $\mathbf{x} \in \mathbb{R}^{P \times Q}$ be the input signal, and $\text{mat}(\mathbf{x}) \in \mathbb{R}^{P \times Q}$ denote its matrix form; for example in a set of sets, each column is an inner set in this matrix form. Throughout, we are assuming a single input-output channel, relying on the idea that having multiple channels simply corresponds to replicating the linear map for each input-output channel combination. Then we can rewrite Eq. (2) as

$$\mathbf{W}_{\mathcal{K} \times \mathcal{H}} \mathbf{x} = \text{vec} \left((\mathbf{W}_{\mathcal{H}} (\text{mat}(\mathbf{x}) \mathbf{1}_Q)) \mathbf{1}_Q^{\top} + \text{mat}(\mathbf{x}) \mathbf{W}_{\mathcal{K}} \right), \quad (3)$$

where the first multiplication $\text{mat}(\mathbf{x}) \mathbf{1}_Q$ pools over columns (inner structures), and after application of the \mathcal{H} -equivariant map $\mathbf{W}_{\mathcal{H}}$ to the pooled value, the result is broadcasted back using right-multiplication by $\mathbf{1}_Q^{\top}$. The second term simply transforms each inner structure using $\mathbf{W}_{\mathcal{K}}$. The overall operation turns out to be simple and intuitive: the inner equivariant map is applied to individual inner structures, and the outer equivariant map is applied to pooled values and broadcasted back.

Example 5 (Equivariant Map for Multiresolution Image). *Consider a coarse pixelization of an image into small patches. Eq. (3) gives the following recipe for a layer equivariant to independent translations within each patch as well as global translation of the coarse image: 1) apply convolution to each patch independently; 2) pool over each patch, apply convolution to the coarse image, and broadcast back to individual pixels in each patch. 3) add the contribution of these two operations. Notice how pooling over regions, a widely used operation in image processing, appears naturally in this approach. One could also easily extend this to more levels of hierarchy for larger images.*

5 Application: Point-Cloud Segmentation

In this section we consider a simple application of the theory above to large-scale point-cloud segmentation. The layer is the 3D version of equivariant linear map for sequence of sets, which is visualized in Fig. 4(d.2) – that is we combine translation and permutation symmetry by voxelizing the point-cloud; see Fig. 6. Using a hierarchical structure is beneficial compared to both the set model and 3D convolution. In particular, the set model lacks any prior of the Euclidean nature of the data, while the 3D convolution, in order to preserve resolution requires a fine-grained voxelization where each point appears in one voxel. The reader may ask: “how is the wreath product used to characterize the setting when the number of points is changing per voxel?” While our derivation using wreath

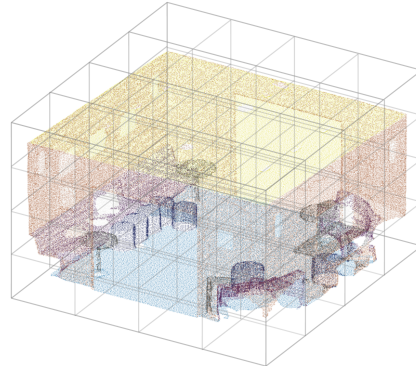


Figure 6: We impose a hierarchy of translation and permutation symmetry by voxelizing the point-cloud.

Table 1: Performance of various models on point cloud segmentation benchmarks.

	SEMANTIC-8		S3DIS		VKITTI		
	OA	mIoU	OA	mIoU	OA	mIoU	mAcc
DEEPSETS[50]	89.3	60.5	67.3	42.7	74.2	42.9	36.8
POINTNET++[36]	85.7	63.1	81.0	54.5	79.7	34.4	47.0
SPG[30]	92.9	76.2	85.5	62.1	84.3	67.3	52.0
CONVPOINT[4]	93.4	76.5	88.8	68.2	-	-	-
KP-FCNN[43]	-	-	-	70.6	-	-	-
RSNET[23]	-	-	-	56.5	-	-	-
PCCN[45]	-	-	-	58.3	-	-	-
SNAPNET[5]	91.0	67.4	-	-	-	-	-
ENGELMANN ET AL. 2018[13]	-	-	-	-	79.7	57.6	35.6
ENGELMANN ET AL. 2017[12]	-	-	-	-	80.6	54.1	36.2
3P-RNN[49]	-	-	-	-	87.8	54.1	41.6
WREATH PRODUCT NET. (ours)	93.9	75.4	90.6	71.2	88.4	68.9	58.6
WREATH PRODUCT NET. + ATTN (ours)	95.2	77.4	95.8	80.1	90.7	69.5	59.2

product assumes a fixed number of points, since the resulting model for the set is independent of the number of data points, we can operate on various points per voxel. This is the same reason why we can apply the same convolution filter to images of different sizes.

The past few years have seen a growing body of work on learning with point cloud data; see [19] for a survey. Many methods use hierarchical aggregation and pooling; this includes the use of furthest point clustering for pooling in POINNET++[37], use of concentric spheres for pooling in SHELLNET, or KD-tree guided pooling in [24]. Several works extend convolution operation to maintain both translation and permutation equivariance in one way or another [3, 4, 45]; see also [6, 17]. Here, objective is not to introduce a radically new procedure, but to show the effectiveness of the approach discussed in previous sections in deep model design from first principles. Indeed, we are able to achieve state-of-the-art in several benchmarks for large point-clouds.

5.1 Equivariance to a Hierarchy of Translations and Permutations

Let $\mathbf{X} \in \mathbb{R}^{N \times C}$ be the input point cloud, where N is the number of points and C is the number of input channels (for concreteness, in this section are including the channel dimension in our formulae.) In addition to 3D coordinates, these channels may include RGB values, normal vectors or any other auxiliary information. Consider a voxelization of the point cloud with a resolution D voxels per dimension, and consider the hierarchy of translation symmetry across voxels and permutation symmetry within each voxel. We may also replace 3D coordinates with relative coordinates within each voxel. Let $\Pi \in \{0, 1\}^{D^3 \times N}$ with one non-zero per column identify the voxel membership. Then the combination of equivariant set layer $\mathbf{X}\mathbf{W}_1 + \mathbf{1}\mathbf{1}^\top \mathbf{X}\mathbf{W}_2$ [50] with 3D convolution using the pool-broadcast interpretation given in Eq. (3), results in the following wreath-product equivariant linear layer

$$\phi(\mathbf{X}) = \mathbf{X}\mathbf{W}_1 + \Pi^\top(\mathbf{W}_3 * (\Pi\mathbf{X})), \quad (4)$$

where $\mathbf{W}_1 \in \mathbb{R}^{C \times C'}$, $*$ denotes the convolution operation, $\mathbf{W}_3 \in \mathbb{R}^{K^3 \times C \times C'}$ is the convolution kernel, with kernel width K , and C' output channels. Here, multiplication with Π and Π^\top performs the pooling and broadcasting from/to points within voxels, respectively. Note that we have dropped the “set” operation $\Pi^\top((\Pi\mathbf{X})\mathbf{W}_2)$ that pools over each voxel, multiplies by the weight and broadcasts back to the points. This is because it can be absorbed in the convolution operation, and therefore it is redundant. In the equation above pooling operation can replace summation (implicit in matrix multiplication) with any other commutative operation; see [35, 48, 51], we use mean-pooling for the experiments. While the layer of Eq. (4) already achieves state-of-the-art in the experiments, we also consider adding an (equivariant) attention mechanism for further improvement; see Appendix B for details.

5.2 Empirical Results

We evaluate our model on two of the largest real world point cloud segmentation benchmarks, SEMANTIC3D [20] and the Stanford Large-Scale 3D Indoor Spaces (S3DIS) [2], as well as a dataset of virtual point cloud scenes, the VKITTI benchmark [16]. As shown in Table 1, in all cases we report new state-of-the-art. Table 2 compares the processing time and the size of different models. The architecture of WREATH PRODUCT NET. is a stack of equivariant layer Eq. (4) plus ReLU nonlinearity and residual connections. WREATH PRODUCT NET.+ATTN also adds the attention mechanism.⁵ Details on architecture and training, as well as further analysis of our results appear in Appendix C.

Outdoor Scene Segmentation - SEMANTIC3D [20] is the largest LiDAR benchmark dataset, consisting of 15 training point clouds and 15 tests point clouds with withheld labels, amassing altogether over 4 billion labeled points from a variety of urban and rural scenes. In particular, rather than working with the smaller REDUCED-8 variation, we run our experiments on the full dataset (SEMANTIC-8)⁶. Table 1 reports mIoU unweighted mean intersection over union metric (mIoU), as well as the overall accuracy (OA) for various methods. In Table 3 we report these measures for different voxelization resolution. At highest resolutions the performance degrades due to overfitting.

Indoor Scene Segmentation - Stanford Large-Scale 3D Indoor Spaces (S3DIS) [2] consists of various 3D RGB point cloud scans from an assortment of room types on six different floor in three buildings on the Stanford campus, totaling almost 600 million points. Table 1 show that we achieve the best overall accuracy as well as mean intersection over union. This is in spite of the fact that our competition use extensive data-augmentation also for this dataset. In addition to random jittering and subsampling employed by both KPCONV and CONVPOINT, KPCONV also uses random dropout of RGB data.

Virtual Scene Segmentation - Virtual KITTI (VKITTI) [16] contains 35 monocular photo-realistic synthetic videos with fully annotated pixel-level labels for each frame and 13 semantic classes in total. Following [12], we project the 2D depth information within these synthetic frames into 3D space, thereby obtaining semantically annotated 3D point clouds. Note that VKITTI is significantly smaller than either SEMANTIC3D or S3DIS, containing only 15 millions points in total.

Conclusion

This paper presents a procedure to design neural networks equivariant to hierarchical symmetries and nested structures. We describe how the imprimitive action of wreath product can formulate the symmetries of the hierarchy, contrast its use case with direct product, and characterize linear maps equivariant to the wreath product of groups. This analysis showed that we can always build an equivariant layer for the hierarchy using equivariant maps for the building blocks, and additional pool-and-broadcast operations. We consider one of the many use cases of this approach to design a deep model for large-scale semantic segmentation of point cloud data, where we are able to achieve state-of-the-art using a simple architecture.

Table 2: Comparison of pre-processing and training time and the number of parameters for SEMANTIC-8.

Method	Pre-Proc. (hrs)	Train (hrs)	# Params. $\times 10^6$
POINTNET	8.82	3.54	3.50
POINTNET++	8.84	7.46	12.40
SNAPNET	13.42	53.44	30.76
SPG	17.43	1.50	0.25
CONVPOINT	13.42	48.74	2.76
OURS	4.39	53.76	5.27
OURS + ATTN	4.39	91.68	47.01

Table 3: Effect of voxelization resolution for SEMANTIC-8.

# Voxels Per Dim	OA	mIoU
2	81.7	62.7
3	85.9	67.3
4	90.6	70.5
5	92.3	72.2
6	93.7	72.8
7	94.4	74.1
8	95.1	75.6
9	94.6	77.1
10	95.2	77.4
12	94.8	73.0
16	90.6	71.5

⁵We have released a PYTORCH implementation of our models at <https://github.com/rw435/wreathProdNet>

⁶The leader-board can be viewed at http://www.semantic3d.net/view_results.php?chl=1

Broader Impact

As deep learning finds its way in various real-world applications, the practitioners are finding more constraints in representing their data in formats and structures amenable to existing deep architectures. The list of basic structures such as images, sets, and graphs that we can approach using deep models has been growing over the past few years. The theoretical contribution of this paper substantially expands this list by enabling deep learning on a hierarchy of structures. This could potentially unlock new applications in data-poor and structure-rich settings. The task we consider in our experiments is deep learning with large point-cloud data, which is finding growing applications, from autonomous vehicles to geographical surveys. While this is not a new task, our empirical results demonstrate the effectiveness of the proposed methodology in dealing with hierarchy in data structure.

Acknowledgments and Disclosure of Funding

This research was in part supported by CIFAR AI Chairs program and NSERC Discovery Grant. Computational resources were provided by Mila and Compute Canada.

References

- [1] Marjan Albooyeh, Daniele Bertolini, and Siamak Ravanbakhsh. Incidence networks for geometric deep learning. *arXiv preprint arXiv:1905.11460*, 2019.
- [2] Iro Armeni, Ozan Sener, Amir R Zamir, Helen Jiang, Ioannis Brilakis, Martin Fischer, and Silvio Savarese. 3d semantic parsing of large-scale indoor spaces. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1534–1543, 2016.
- [3] Matan Atzmon, Haggai Maron, and Yaron Lipman. Point convolutional neural networks by extension operators. *arXiv preprint arXiv:1803.10091*, 2018.
- [4] Alexandre Boulch. Generalizing discrete convolutions for unstructured point clouds. *arXiv preprint arXiv:1904.02375*, 2019.
- [5] Alexandre Boulch, Joris Guerry, Bertrand Le Saux, and Nicolas Audebert. Snapnet: 3d point cloud semantic labeling with 2d deep segmentation networks. *Computers & Graphics*, 71: 189–198, 2018.
- [6] Christopher Choy, JunYoung Gwak, and Silvio Savarese. 4d spatio-temporal convnets: Minkowski convolutional neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3075–3084, 2019.
- [7] Taco S Cohen and Max Welling. Group equivariant convolutional networks. *arXiv preprint arXiv:1602.07576*, 2016.
- [8] Taco S Cohen and Max Welling. Steerable cnns. *arXiv preprint arXiv:1612.08498*, 2016.
- [9] Taco S Cohen, Mario Geiger, Jonas Köhler, and Max Welling. Spherical cnns. *arXiv preprint arXiv:1801.10130*, 2018.
- [10] Taco S Cohen, Mario Geiger, and Maurice Weiler. A general theory of equivariant cnns on homogeneous spaces. In *Advances in Neural Information Processing Systems*, pages 9142–9153, 2019.
- [11] Sander Dieleman, Jeffrey De Fauw, and Koray Kavukcuoglu. Exploiting cyclic symmetry in convolutional neural networks. *arXiv preprint arXiv:1602.02660*, 2016.
- [12] Francis Engelmann, Theodora Kontogianni, Alexander Hermans, and Bastian Leibe. Exploring spatial context for 3d semantic segmentation of point clouds. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 716–724, 2017.
- [13] Francis Engelmann, Theodora Kontogianni, Jonas Schult, and Bastian Leibe. Know what your neighbors do: 3d semantic segmentation of point clouds. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 0–0, 2018.

- [14] Jin Fang, Feilong Yan, Tongtong Zhao, Feihu Zhang, Dingfu Zhou, Ruigang Yang, Yu Ma, and Liang Wang. Simulating lidar point cloud for autonomous driving using real-world scenes and traffic flows. *arXiv preprint arXiv:1811.07112*, 1, 2018.
- [15] Richard Foote, Gagan Mirchandani, Daniel N Rockmore, Dennis Healy, and Tim Olson. A wreath product group approach to signal and image processing. i. multiresolution analysis. *IEEE Transactions on Signal Processing*, 48(1):102–132, 2000.
- [16] Adrien Gaidon, Qiao Wang, Yohann Cabon, and Eleonora Vig. Virtual worlds as proxy for multi-object tracking analysis. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4340–4349, 2016.
- [17] Benjamin Graham and Laurens van der Maaten. Submanifold sparse convolutional networks. *arXiv preprint arXiv:1706.01307*, 2017.
- [18] Devon Graham and Siamak Ravanbakhsh. Deep models for relational databases. *arXiv preprint arXiv:1903.09033*, 2019.
- [19] Yulan Guo, Hanyun Wang, Qingyong Hu, Hao Liu, Li Liu, and Mohammed Bennamoun. Deep learning for 3d point clouds: A survey. *arXiv preprint arXiv:1912.12033*, 2019.
- [20] Timo Hackel, Nikolay Savinov, Lubor Ladicky, Jan D Wegner, Konrad Schindler, and Marc Pollefeys. Semantic3d. net: A new large-scale point cloud classification benchmark. *arXiv preprint arXiv:1704.03847*, 2017.
- [21] Jason Hartford, Devon R Graham, Kevin Leyton-Brown, and Siamak Ravanbakhsh. Deep models of interactions across sets. *arXiv preprint arXiv:1803.02879*, 2018.
- [22] Richard B Holmes. Mathematical foundations of signal processing. 2. the role of group theory. Technical report, MASSACHUSETTS INST OF TECH LEXINGTON LINCOLN LAB, 1987.
- [23] Qiangui Huang, Weiyue Wang, and Ulrich Neumann. Recurrent slice networks for 3d segmentation of point clouds. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2626–2635, 2018.
- [24] Roman Klokov and Victor Lempitsky. Escape from cells: Deep kd-networks for the recognition of 3d point cloud models. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 863–872, 2017.
- [25] Imre Risi Kondor. *Group theoretical methods in machine learning*, volume 2. Columbia University, 2008.
- [26] Risi Kondor and Shubhendu Trivedi. On the generalization of equivariance and convolution in neural networks to the action of compact groups. *arXiv preprint arXiv:1802.03690*, 2018.
- [27] Risi Kondor, Zhen Lin, and Shubhendu Trivedi. Clebsch–gordan nets: a fully fourier space spherical convolutional neural network. In *Advances in Neural Information Processing Systems*, pages 10137–10146, 2018.
- [28] Risi Kondor, Hy Truong Son, Horace Pan, Brandon Anderson, and Shubhendu Trivedi. Covariant compositional networks for learning graphs. *arXiv preprint arXiv:1801.02144*, 2018.
- [29] Loic Landrieu and Mohamed Boussaha. Point cloud oversegmentation with graph-structured deep metric learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7440–7449, 2019.
- [30] Loic Landrieu and Martin Simonovsky. Large-scale point cloud semantic segmentation with superpoint graphs. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4558–4567, 2018.
- [31] Yann LeCun, Yoshua Bengio, et al. Convolutional networks for images, speech, and time series. *The handbook of brain theory and neural networks*, 3361(10):1995, 1995.
- [32] Haggai Maron, Heli Ben-Hamu, Nadav Shamir, and Yaron Lipman. Invariant and equivariant graph networks. *arXiv preprint arXiv:1812.09902*, 2018.

- [33] Haggai Maron, Or Litany, Gal Chechik, and Ethan Fetaya. On learning sets of symmetric elements. *ArXiv*, abs/2002.08599, 2020.
- [34] Gagan Mirchandani, Richard Foote, Daniel N Rockmore, Dennis Healy, and Tim Olson. A wreath product group approach to signal and image processing. ii. convolution, correlation, and applications. *IEEE transactions on signal processing*, 48(3):749–767, 2000.
- [35] Ryan L Murphy, Balasubramaniam Srinivasan, Vinayak Rao, and Bruno Ribeiro. Janossy pooling: Learning deep permutation-invariant functions for variable-size inputs. *arXiv preprint arXiv:1811.01900*, 2018.
- [36] Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 652–660, 2017.
- [37] Charles Ruizhongtai Qi, Li Yi, Hao Su, and Leonidas J Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. In *Advances in neural information processing systems*, pages 5099–5108, 2017.
- [38] Siamak Ravanbakhsh, Jeff Schneider, and Barnabas Poczos. Equivariance through parameter-sharing. In *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *JMLR: WCP*, August 2017.
- [39] Daniel N Rockmore. Fast fourier transforms for wreath products. *Applied and Computational Harmonic Analysis*, 2(3):279–292, 1995.
- [40] Yudai Sadakuni, Ryosuke Kusakari, Kazuya Onda, and Yoji Kuroda. Construction of highly accurate depth estimation dataset using high density 3d lidar and stereo camera. In *2019 IEEE/SICE International Symposium on System Integration (SII)*, pages 554–559. IEEE, 2019.
- [41] John Shawe-Taylor. Building symmetries into feedforward networks. In *Artificial Neural Networks, 1989., First IEE International Conference on (Conf. Publ. No. 313)*, pages 158–162. IET, 1989.
- [42] Lyne Tchampi, Christopher Choy, Iro Armeni, JunYoung Gwak, and Silvio Savarese. Segcloud: Semantic segmentation of 3d point clouds. In *2017 International Conference on 3D Vision (3DV)*, pages 537–547. IEEE, 2017.
- [43] Hugues Thomas, Charles R Qi, Jean-Emmanuel Deschaud, Beatriz Marcotegui, François Goulette, and Leonidas J Guibas. Kpconv: Flexible and deformable convolution for point clouds. *arXiv preprint arXiv:1904.08889*, 2019.
- [44] 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.
- [45] Shenlong Wang, Simon Suo, Wei-Chiu Ma, Andrei Pokrovsky, and Raquel Urtasun. Deep parametric continuous convolutional neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2589–2597, 2018.
- [46] Maurice Weiler and Gabriele Cesa. General $e(2)$ -equivariant steerable cnns. In *Advances in Neural Information Processing Systems*, pages 14334–14345, 2019.
- [47] Jeffrey Wood and John Shawe-Taylor. Representation theory and invariant neural networks. *Discrete applied mathematics*, 69(1-2):33–60, 1996.
- [48] Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. How powerful are graph neural networks? *arXiv preprint arXiv:1810.00826*, 2018.
- [49] Xiaoqing Ye, Jiamao Li, Hexiao Huang, Liang Du, and Xiaolin Zhang. 3d recurrent neural networks with context fusion for point cloud semantic segmentation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 403–417, 2018.

- [50] Manzil Zaheer, Satwik Kottur, Siamak Ravanbakhsh, Barnabas Poczos, Ruslan R Salakhutdinov, and Alexander J Smola. Deep sets. In *Advances in neural information processing systems*, pages 3391–3401, 2017.
- [51] Muhan Zhang, Zhicheng Cui, Marion Neumann, and Yixin Chen. An end-to-end deep learning architecture for graph classification. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.

A Proof of Theorem 4.1

Proof. We show that both of the two terms in Eq. (2) are equivariant to $\mathcal{K} \wr \mathcal{H}$ -action as expressed by the permutation matrix of Eq. (1).

Part 1. We show that $\mathbf{G}^{(g)}$ commutes with the first term in Eq. (2) for any $g \in \mathcal{K} \wr \mathcal{H}$:

$$(\mathbf{W}_{\mathcal{H}} \otimes (\mathbf{1}_Q \mathbf{1}_Q^\top)) \mathbf{G}^{(g)} = (\mathbf{W}_{\mathcal{H}} \otimes (\mathbf{1}_Q \mathbf{1}_Q^\top)) \left(\sum_{p=1}^P \mathbf{1}_{p, \mathcal{H} \cdot p} \otimes \mathbf{G}^{(h_p)} \right) \quad (5)$$

$$= \sum_{p=1}^P (\mathbf{W}_{\mathcal{H}} \mathbf{1}_{p, \mathcal{H} \cdot p}) \otimes \overbrace{((\mathbf{1}_Q \mathbf{1}_Q^\top) \mathbf{G}^{(h_p)})}^{=(\mathbf{1}_Q \mathbf{1}_Q^\top)} \quad (6)$$

$$= \mathbf{W}_{\mathcal{H}} \left(\sum_{p=1}^P \mathbf{1}_{p, \mathcal{H} \cdot p} \right) \otimes (\mathbf{1}_Q \mathbf{1}_Q^\top) \quad (7)$$

$$= (\mathbf{W}_{\mathcal{H}} \mathbf{G}^{(h)}) \otimes (\mathbf{1}_Q \mathbf{1}_Q^\top) = (\mathbf{G}^{(h)} \mathbf{W}_{\mathcal{H}}) \otimes (\mathbf{1}_Q \mathbf{1}_Q^\top) \quad (8)$$

$$= \sum_{p=1}^P (\mathbf{1}_{p, \mathcal{H} \cdot p} \mathbf{W}_{\mathcal{H}}) \otimes ((\mathbf{1}_Q \mathbf{1}_Q^\top) \mathbf{G}^{(h_p)}) \quad (9)$$

$$= \sum_{p=1}^P (\mathbf{1}_{p, \mathcal{H} \cdot p} \otimes \mathbf{G}^{(h_p)}) (\mathbf{W}_{\mathcal{H}} \otimes (\mathbf{1}_Q \mathbf{1}_Q^\top)) \quad (10)$$

$$= \left(\sum_{p=1}^P \mathbf{1}_{p, \mathcal{H} \cdot p} \otimes \mathbf{G}^{(h_p)} \right) (\mathbf{W}_{\mathcal{H}} \otimes (\mathbf{1}_Q \mathbf{1}_Q^\top)) = \mathbf{G}^{(g)} (\mathbf{W}_{\mathcal{H}} \otimes (\mathbf{1}_Q \mathbf{1}_Q^\top)). \quad (11)$$

In Eq. (5) we substitute from Eq. (1). In Eq. (6) we pulled the summation out, and then used the mixed product property of Kronecker product. We also note that a permutation of the uniformly constant matrix $\mathbf{1}_Q \mathbf{1}_Q^\top$ is constant. In Eq. (7), since the summation can be restricted to the first term we pull out $\mathbf{W}_{\mathcal{H}}$. In Eq. (8) we use the fact that the summation of the previous line is the permutation matrix $\mathbf{G}^{(h)}$ and apply the key assumption that $\mathbf{W}_{\mathcal{H}}$ and $\mathbf{W}_{\mathcal{K}}$ are equivariant to \mathcal{H} and \mathcal{K} -action respectively. The following lines repeat the procedure so far in reverse order to commute $\mathbf{G}^{(g)}$ and $(\mathbf{W}_{\mathcal{H}} \otimes (\mathbf{1}_Q \mathbf{1}_Q^\top))$.

Part 2. For the second term, again we use the mixed-product property and equivariance of the input maps

$$(\mathbf{I}_P \otimes \mathbf{W}_{\mathcal{K}}) \mathbf{G}^{(g)} = (\mathbf{I}_P \otimes \mathbf{W}_{\mathcal{K}}) \left(\sum_{p=1}^P \mathbf{1}_{p, \mathcal{H} \cdot p} \otimes \mathbf{G}^{(h_p)} \right) \quad (12)$$

$$= \sum_{p=1}^P (\mathbf{I}_P \mathbf{1}_{p, \mathcal{H} \cdot p}) \otimes (\mathbf{W}_{\mathcal{K}} \mathbf{G}^{(h_p)}) \quad (13)$$

$$= \sum_{p=1}^P (\mathbf{1}_{p, \mathcal{H} \cdot p} \mathbf{I}_P) \otimes (\mathbf{G}^{(h_p)} \mathbf{W}_{\mathcal{K}}) \quad (14)$$

$$= \sum_{p=1}^P (\mathbf{1}_{p, \mathcal{H} \cdot p} \otimes \mathbf{G}^{(h_p)}) (\mathbf{I}_P \otimes \mathbf{W}_{\mathcal{K}}) = \mathbf{G}^{(g)} (\mathbf{I}_P \otimes \mathbf{W}_{\mathcal{K}}) \quad (15)$$

where in Eq. (14) we used the fact that the identity matrix commutes with any matrix, as well as the equivariance of $\mathbf{G}^{(h_p)}$. Eqs. (13) and (15) simply use the mixed product property. Putting the two parts together shows the equivariance of the first and second term in Eq. (2), completing the proof. \square

A.1 Proof of Maximality

Previously we proved that the linear maps of Eq. (2) is equivariant to imprimitive action of wreath product $\mathcal{K} \wr \mathcal{H}$ given by Eq. (1). Here, we prove that any $\mathcal{K} \wr \mathcal{H}$ equivariant linear map has this form.

Claim 2. *Assuming that any \mathcal{K} -equivariant (\mathcal{H} -equivariant) linear map $\mathbf{W}_{\mathcal{K}}$ ($\mathbf{W}_{\mathcal{H}}$) can be written as a linear combination of K (H) independent linear operators, then any $\mathbf{W}_{\mathcal{K} \wr \mathcal{H}}$ as defined in Eq. (2) is a linear combination of $K + H - 1$ independent linear bases.*

Proof. Consider the two terms (I) $\mathbf{W}_{\mathcal{H}} \otimes (\mathbf{1}_Q \mathbf{1}_Q^\top)$ and (II) $\mathbf{I}_P \otimes \mathbf{W}_{\mathcal{K}}$ in Eq. (2), where (I) has H independent bases and (II) has K independent bases. If they had independent bases, their summation would have $K + H$ linear bases. However, (I) and (II) have exactly one shared basis: $\mathbf{W}_{\mathcal{H}} = \mathbf{I}_P$ is \mathcal{H} -equivariant, and $\mathbf{W}_{\mathcal{K}} = \mathbf{1}_Q \mathbf{1}_Q^\top$ is \mathcal{K} -equivariant. Therefore, $\mathbf{I}_P \otimes (\mathbf{1}_Q \mathbf{1}_Q^\top)$ is a basis for both (I) and (II). \square

Next, we show that any $\mathcal{K} \wr \mathcal{H}$ -equivariant linear map can be written using as a weighted sum of $K + H - 1$ independent linear operators. Together with proof of equivariance of $\mathbf{W}_{\mathcal{K} \wr \mathcal{H}}$ these two facts prove that any $\mathcal{K} \wr \mathcal{H}$ equivariant linear map has the form of Eq. (2).

Claim 3. *Any linear map that is equivariant to the imprimitive action of $\mathcal{K} \wr \mathcal{H}$ can a linear combination of $K + H - 1$ equivariant linear bases.*

Proof. Let $\mathbf{W} \in \mathbb{R}^{N \times N}$ be the matrix representing this equivariant map, satisfying $\mathbf{W} \mathbf{G}^g = \mathbf{G}^g \mathbf{W} \quad \forall g \in \mathcal{K} \wr \mathcal{H}$. Since $\mathbf{G}^{g^\top} \mathbf{W} \mathbf{G}^g = \mathbf{W} \quad \forall g \in \mathcal{K} \wr \mathcal{H}$, all elements of the matrix \mathbf{W} that are in the same orbit are tied together. This constraint means that the number of unique values in \mathbf{W} are equal to the number of orbits in this simultaneous action on rows and columns of \mathbf{W} . Using a property of Kronecker product, we can also write the equation above as $\text{vec}(\mathbf{G}^g \mathbf{W} \mathbf{G}^{g^\top}) = (\mathbf{G}^g \otimes \mathbf{G}^g) \text{vec}(\mathbf{W})$. Therefore, the diagonal group action is given by $\mathbf{G}^g \otimes \mathbf{G}^g$. Using *Burnside's lemma* we get the number of orbits L in this action to be

$$L = \frac{1}{|\mathcal{G}|} \sum_{g \in \mathcal{K} \wr \mathcal{H}} \text{Tr}(\mathbf{G}^{(g)} \otimes \mathbf{G}^{(g)}) = \sum_{g \in \mathcal{K} \wr \mathcal{H}} \text{Tr}(\mathbf{G}^{(g)})^2 \quad (16)$$

Our objective is to show that $L = H + K - 1$. Our strategy is to replace the definition of \mathbf{G}^g from Eq. (1) into Eq. (16) and simplify. This simplification involves numerous steps.

$$L = \frac{1}{|\mathcal{H}||\mathcal{K}|^P} \sum_{\hat{h} \in \mathcal{H}, \hat{h}_1, \dots, \hat{h}_P \in \mathcal{K}} \text{Tr} \left(\sum_{p=1}^P \mathbf{1}_{p, \hat{h}_p} \otimes \mathbf{G}^{(\hat{h}_p)} \right)^2 \quad (17)$$

$$= \frac{1}{|\mathcal{H}||\mathcal{K}|^P} \sum_{\hat{h} \in \mathcal{H}, \hat{h}_1, \dots, \hat{h}_P \in \mathcal{K}} \text{Tr} \left(\sum_{p=1}^P \mathbf{G}_{p,p}^{(\hat{h})} \mathbf{G}^{(\hat{h}_p)} \right)^2 \quad (18)$$

$$= \frac{1}{|\mathcal{H}||\mathcal{K}|^P} \sum_{\hat{h} \in \mathcal{H}, \hat{h}_1, \dots, \hat{h}_P \in \mathcal{K}} \left(\sum_{p=1}^P \mathbf{G}_{p,p}^{(\hat{h})} \text{Tr}(\mathbf{G}^{(\hat{h}_p)}) \right) \left(\sum_{p'=1}^P \mathbf{G}_{p',p'}^{(\hat{h})} \text{Tr}(\mathbf{G}^{(\hat{h}_{p'})}) \right) \quad (19)$$

$$= \frac{1}{|\mathcal{H}||\mathcal{K}|^P} \sum_{\hat{h} \in \mathcal{H}, \hat{h}_1, \dots, \hat{h}_P \in \mathcal{K}} \left(\sum_{p=1}^P \mathbf{G}_{p,p}^{(\hat{h})} \text{Tr}(\mathbf{G}^{(\hat{h}_p)}) \right)^2 + \left(\sum_{p=1}^P \sum_{p' \neq p, p'=1}^P \mathbf{G}_{p,p}^{(\hat{h})} \mathbf{G}_{p',p'}^{(\hat{h})} \text{Tr}(\mathbf{G}^{(\hat{h}_p)}) \text{Tr}(\mathbf{G}^{(\hat{h}_{p'})}) \right) \quad (20)$$

where in Eq. (18) we observed that the contribution from $\mathbf{G}^{(\hat{h}_p)}$ to the trace of $\mathcal{G}^{(g)}$ is non-zero only if $\mathbf{G}_{p,p}^{(\hat{h})} = 1$. Note that all the \mathbf{G} matrices are permutation matrices and in the following we will continue to use this fact without elaborating. In Eq. (19) we used linearity of trace, and in Eq. (20) we divided the product of the previous line into $p = p'$ and $p \neq p'$. As we show below, the first term simplifies to K and the second terms simplifies to $H - 1$, showing that $L = K + H - 1$. We start by

simplifying the first term of Eq. (20) below:

$$\frac{1}{|\mathcal{H}||\mathcal{K}|^P} \sum_{\hat{h} \in \mathcal{H}, \hat{h}_1, \dots, \hat{h}_P \in \mathcal{K}} \left(\sum_{p=1}^P \mathbf{G}_{p,p}^{(\hat{h})} \text{Tr}(\mathbf{G}^{(\hat{h}_p)})^2 \right) \quad (21)$$

$$= \frac{1}{|\mathcal{H}||\mathcal{K}|^P} \sum_{p=1}^P \left(\sum_{\hat{h} \in \mathcal{H}} \mathbf{G}_{p,p}^{(\hat{h})} \right) \left(\sum_{\hat{h}_1, \dots, \hat{h}_P \in \mathcal{K}} \text{Tr}(\mathbf{G}^{(\hat{h}_p)})^2 \right) \quad (22)$$

$$= \frac{1}{|\mathcal{H}||\mathcal{K}|^P} \left(\sum_{\hat{h} \in \mathcal{H}} \sum_{p=1}^P \mathbf{G}_{p,p}^{(\hat{h})} \right) \left(|\mathcal{K}|^{P-1} \sum_{\hat{h} \in \mathcal{K}} \text{Tr}(\mathbf{G}^{(\hat{h})})^2 \right) \quad (23)$$

$$= \frac{1}{|\mathcal{H}||\mathcal{K}|^P} \left(\sum_{\hat{h} \in \mathcal{H}} \text{Tr}(\mathbf{G}^{(\hat{h})}) \right) (|\mathcal{K}|^{P-1} |\mathcal{K}| K) \quad (24)$$

$$= \frac{1}{|\mathcal{H}||\mathcal{K}|^P} |\mathcal{H}||\mathcal{K}|^P K = K \quad (25)$$

where beside simple algebra, in Eq. (24), we used the Burnside lemma on $\mathbf{W}_{\mathcal{K}}$, which gives us the number of unique parameters (orbits) as $K = \frac{1}{|\mathcal{K}|} \sum_{\hat{h} \in \mathcal{K}} \text{Tr}(\mathbf{G}^{(\hat{h})})^2$. Similarly for \mathcal{H} -action we get $\frac{1}{|\mathcal{H}|} \sum_{\hat{h} \in \mathcal{H}} \text{Tr}(\mathbf{G}^{(\hat{h})})^2 = H$. In Eq. (25) we used the fact that \mathcal{H} action on \mathbb{P} is transitive (*i.e.*, it has a single orbit), to get $|\mathcal{H}| = \sum_{\hat{h} \in \mathcal{H}} \text{Tr}(\mathbf{G}^{(\hat{h})})$. Similarly, due to transitivity of \mathcal{K} -action on \mathbb{Q} we have $|\mathcal{K}| = \sum_{\hat{h} \in \mathcal{K}} \text{Tr}(\mathbf{G}^{(\hat{h})})$. Next, we use these same equalities to simplify the second term of Eq. (20):

$$\frac{1}{|\mathcal{H}||\mathcal{K}|^P} \sum_{\hat{h} \in \mathcal{H}, \hat{h}_1, \dots, \hat{h}_P \in \mathcal{K}} \sum_{p=1}^P \sum_{p' \neq p, p'=1}^P \mathbf{G}_{p,p}^{(\hat{h})} \mathbf{G}_{p',p'}^{(\hat{h})} \text{Tr}(\mathbf{G}^{(\hat{h}_p)}) \text{Tr}(\mathbf{G}^{(\hat{h}_{p'})}) \quad (26)$$

$$= \frac{1}{|\mathcal{H}||\mathcal{K}|^P} \sum_{p=1}^P \sum_{p' \neq p, p'=1}^P \left(\sum_{\hat{h} \in \mathcal{H}} \mathbf{G}_{p,p}^{(\hat{h})} \mathbf{G}_{p',p'}^{(\hat{h})} \right) \left(|\mathcal{K}|^{P-2} \left(\sum_{\hat{h} \in \mathcal{K}} \text{Tr}(\mathbf{G}^{(\hat{h}_p)}) \right)^2 \right) \quad (27)$$

$$= \frac{1}{|\mathcal{H}||\mathcal{K}|^P} \left(\sum_{\hat{h} \in \mathcal{H}} \left(\sum_{p=1}^P \sum_{p'=1}^P \mathbf{G}_{p,p}^{(\hat{h})} \mathbf{G}_{p',p'}^{(\hat{h})} \right) - \left(\sum_{p=1}^P \mathbf{G}_{p,p}^{(\hat{h})} \right)^2 \right) (|\mathcal{K}|^{P-2} |\mathcal{K}|^2) \quad (28)$$

$$= \frac{1}{|\mathcal{H}|} \sum_{\hat{h} \in \mathcal{H}} \left(\sum_{p=1}^P \mathbf{G}_{p,p}^{(\hat{h})} \right) \left(\sum_{p'=1}^P \mathbf{G}_{p',p'}^{(\hat{h})} \right) - \left(\sum_{p=1}^P \mathbf{G}_{p,p}^{(\hat{h})} \right) \quad (29)$$

$$= \frac{1}{|\mathcal{H}|} \sum_{\hat{h} \in \mathcal{H}} \text{Tr}(\mathbf{G}^{(\hat{h})})^2 - \text{Tr}(\mathbf{G}^{(\hat{h})}) = H - 1. \quad (30)$$

Eqs. (25) and (30) together with Eq. (16) show that $L = H + K - 1$, that any linear map $\mathbf{W}_{\mathcal{G}}$ equivariant to imprimitive action of $\mathcal{K} \wr \mathcal{H}$, can be written using $H + K - 1$ independent linear bases. Considering this proof together with our proof for the first part (that the equivariant map of Eq. (2) has $H + K - 1$ independent linear bases) completes the proof that any $\mathcal{K} \wr \mathcal{H}$ -equivariant linear map is of the form Eq. (2). \square

B Adaptive Pooling and Attention

While the layer of Eq. (4) performs competitively, to improve its performance we add a novel attention mechanism. To this end we complement the pool-and-broadcast scheme of Eq. (4) with an adaptive pooling mechanism using a learned $\tilde{\Pi} : \mathbb{R}^N \rightarrow \Delta_L$ in Eq. (4), where Δ_L is the L -dimensional probability simplex. Here, one may interpret L as the number of *latent classes*, analogous to D^3 voxels. The pooling matrix $\tilde{\Pi}$, is a function of input through a linear map $\mathbf{W}_3^{(\ell)} \in \mathbb{R}^{C \times L}$ followed by Softmax : $\mathbb{R}^L \rightarrow (0, 1)^L$, so that the probability of each point belonging to different classes sums to one – that is

$$\tilde{\Pi}_n = \text{Softmax}((\mathbf{X}\mathbf{W}_3)_n) \forall n \in [N]. \quad (31)$$

We then use a linear function that models the interaction between latent classes, for each pair of channels. This is done using a rank four tensor $\mathbf{W}_4 \in \mathbb{R}^{L \times L \times C \times C}$. As before the result is broadcasted

back using $\tilde{\Pi}_n^\top$, giving us the adaptive pooling layer, in which the output for channel c is given by

$$\phi(\mathbf{X})_c = \sum_{c'=1}^{C'} \tilde{\Pi}^\top \mathbf{W}_{4,c,c'} \tilde{\Pi} \mathbf{X}_c. \quad (32)$$

Intuitively this layer can decide which subset of nodes to pool, and therefore acts like an equivariant attention mechanism. In experiments, we see further improvement in results when adding this nonlinear layer to the linear layer of Eq. (4).

C Details on Experiments

C.1 Architecture and Data Processing

Our results are produced using minor architectural variations and limited hyperparameter search. Concretely, our best-performing models on SEMANTIC3D and S3DIS involve 56 residual-style blocks, where each block comprises of 2 equivariant linear maps of Eq. (4). Each map involves a single 3D convolution operator with periodic padding and a $3 \times 3 \times 3$ kernel. We use an identity map to facilitate skip connections within these blocks. The number channels is fixed 64, except the final block, which maps to the number of segmentation classes (8 for SEMANTIC3D and 13 for S3DIS). When incorporating adaptive pooling, we utilize 5, 10, 15, 25 and 50 latent classes L for blocks 1-19, 19-29, 30-39, 40-49 and 50-56, respectively and inclusively. Our best-performing model on VKITTI uses 22 such residual-style blocks with identical architectural hyperparameters. For adaptive pooling, we then use 5, 10, 15, and 20 latent classes for blocks 1-9, 10-14, 15-19, and 20-22, respectively and inclusively.

Additionally, while previous works rely on stochastic point cloud subsampling during pre-processing and reprojection during post-processing, we simply split large point clouds into subsets of 1 million points each to ameliorate memory issues. We then compute a $9 \times 9 \times 9$ voxelization over each sample, and train on mini-batches of 4 such samples. To generate predictions, we stitch predictions on these smaller samples together.

C.2 Datasets: Training, Validation and Test

Outdoor Scene Segmentation - SEMANTIC3D SEMANTIC3D [20] We hold out 4 of the available 15 training point clouds to use as a validation set, as in [30].

Table 4 reports the overall accuracy (OA), and mean intersection over union (mIoU) for our method and the competition. We achieve both the best overall accuracy as well as the best mean intersection over union.

Method	OA	mIoU	Man-Made Terrain	Natural Vegetation	High Vegetation	Low Vegetation	Buildings	Hardscape	Scanning Artifacts	Cars
DEEPSETS[50]	89.3	60.5	90.8	76.3	41.9	22.1	94.0	46.5	26.8	85.4
POINTNET[36]	85.7	63.1	81.9	78.1	64.3	51.7	75.9	36.4	43.7	72.6
SNAPNET[5]	91.0	67.4	89.6	79.5	74.8	56.1	90.9	36.5	34.3	77.2
SPG[30]	92.9	76.2	91.5	75.6	78.3	71.7	94.4	56.8	52.9	88.4
CONVPOINT[4]	93.4	76.5	92.1	80.6	76.0	71.9	95.6	47.3	61.1	87.7
WREATH PRODUCT NET. (ours)	93.9	75.4	93.4	84.0	76.4	68.2	96.8	45.6	47.4	91.9
WREATH PRODUCT NET. + ATTN (ours)	94.6	77.1	95.2	87.1	75.3	67.1	96.1	51.3	51.0	93.4

Table 4: Performance of various models on the *full* SEMANTIC3D dataset (semantic-8). Higher is better, bolded is best. mIoU is unweighted mean intersection over union metric. OA is overall accuracy. Per class splits show mIoU.

Indoor Scene Segmentation - Stanford Large-Scale 3D Indoor Spaces (S3DIS) The S3DIS dataset [2] consists of various 3D RGB point cloud scans from an assortment of room types on six different floor in three buildings on the Stanford campus, totaling almost 600 million points. Following previous works by [12, 30, 36, 37, 42], we perform 6-fold cross validation with micro-averaging, computing all metrics once over the merged predictions of all test folds.

Here, we achieve the best overall accuracy as well as mean intersection over union. This is in spite of the fact that our competition use extensive data-augmentation also for this dataset. In particular, both KPCONV and CONVPOINT use random jittering and subsampling.

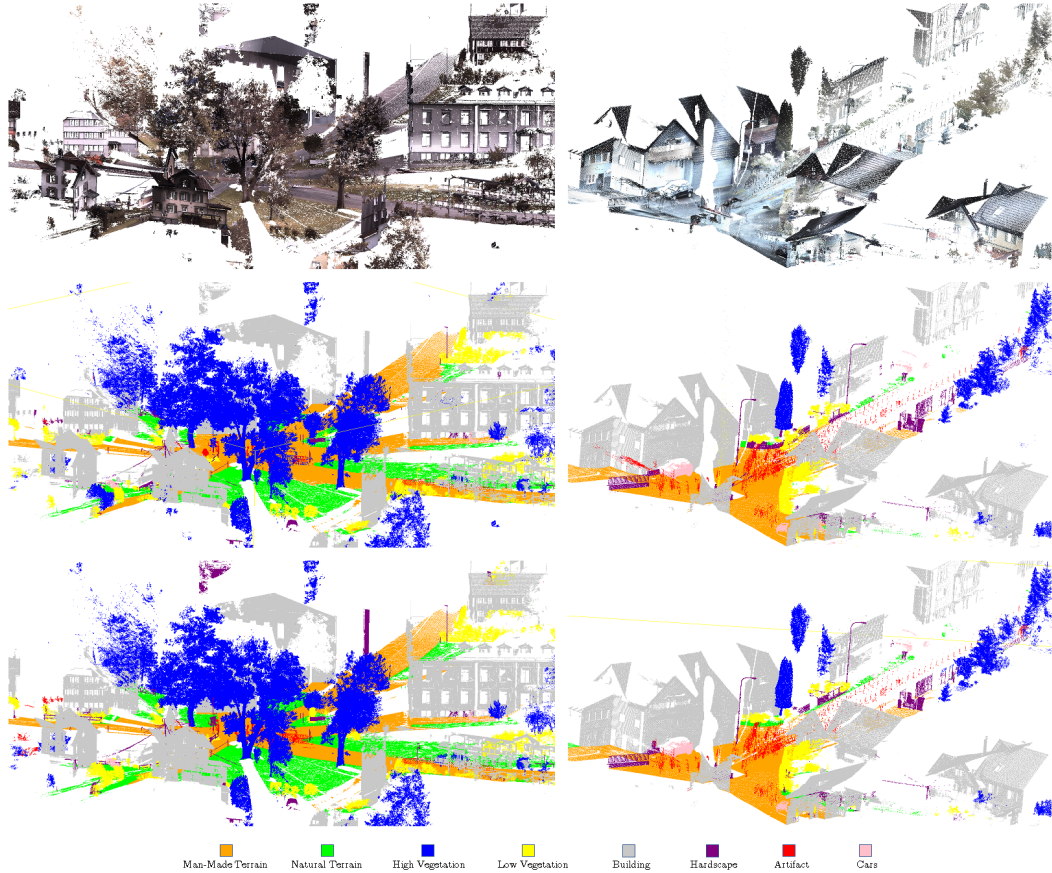


Figure 7: **(top left)** RGB scan of a point cloud with 280 994 028 points from the SEMANTIC3D dataset; **(middle left)** segmentation results from our method; **(bottom left)** ground truth segmentation; **(top right)** RGB scan of a point cloud with 19 767 991 points from the SEMANTIC3D dataset. **(middle right)** segmentation results from our method; **(bottom right)** ground truth segmentation;

Method	OA	mIoU	Ceiling	Floor	Wall	Beam	Column	Window	Door	Chair	Table	Bookcase	Sofa	Board	Clutter
DEEPSETS[50]	67.3	42.7	81.1	72.4	67.2	16.9	25.8	44.2	48.5	51.0	49.8	21.7	24.4	17.2	34.6
POINTNET[36]	78.5	47.6	88.0	88.7	69.3	42.4	23.1	47.5	51.6	42.0	54.1	38.2	9.6	29.4	35.2
RSNET[23]	N/A	56.5	92.5	92.8	78.6	32.8	34.4	51.6	68.1	60.1	59.7	50.2	16.4	44.9	52.0
PCCN[45]	N/A	58.3	92.3	96.2	75.9	0.27	6.0	69.5	63.5	65.6	66.9	68.9	47.3	59.1	46.2
SPG[30]	85.5	62.1	89.9	95.1	76.4	62.8	47.1	55.3	68.4	73.5	69.2	63.2	45.9	8.7	52.9
CONVPOINT[4]	88.8	68.2	95.0	97.3	81.7	47.1	34.6	63.2	73.2	75.3	71.8	64.9	59.2	57.6	65.0
KP-FCNN[43]	N/A	70.6	93.6	92.4	83.1	63.9	54.3	66.1	76.6	57.8	64.0	69.3	74.9	61.3	60.3
WREATH PRODUCT NET. (ours)	90.6	71.2	94.3	96.7	80.6	65.0	76.2	62.1	71.9	64.3	62.7	60.2	68.8	63.5	59.9
WREATH PRODUCT NET. + ATTN (ours)	95.8	80.1	97.2	97.8	88.6	72.3	82.0	73.6	78.6	75.1	72.0	73.1	71.4	82.1	77.2

Table 5: Performance of various models on the S3DIS dataset (micro-averaged over all 6 folds). Higher is better, bolded is best. mIoU is unweighted mean intersection over union metric. OA is overall accuracy. Per class splits show mIoU.

C.3 Virtual Scene Segmentation - Virtual KITTI

The VKITTI dataset [16] contains 35 monocular photo-realistic synthetic videos with fully annotated pixel-level labels for each frame and 13 semantic classes in total. Following [12], we project the 2D depth information within these synthetic frames into 3D space, thereby obtaining semantically annotated 3D point clouds. Similar to the training and evaluation scheme in [12, 29], we separate the original set of sequences into 6 non-overlapping subsequences and use a 6 fold cross-validation protocol (with micro-averaging similar to the methodology on S3DIS).

Note that VKITTI is significantly smaller than either SEMANTIC3D or S3DIS, containing only 15 millions points in total. We hypothesize that these smaller, and sparser point clouds provide little geometric signal outside vegetation and road structure. This partially explains our only incremental improvement

Method	OA	mIoU	mAcc
DEEPSETS[50]	74.2	42.9	36.8
POINTNET[36]	79.7	34.4	47.0
ENGELMANN ET AL. 2018[13]	79.7	57.6	35.6
ENGELMANN ET AL. 2017[12]	80.6	54.1	36.2
3P-RNN[49]	87.8	54.1	41.6
SPG[29]	84.3	67.3	52.0
WREATH PRODUCT NET. (ours)	88.4	68.9	58.6
WREATH PRODUCT NET. + ATTN (ours)	90.7	69.5	59.2

Table 6: Performance of various models on the VKITTI dataset (micro-averaged over all 6 folds). Higher is better, best results are in bold. mIoU is unweighted mean intersection over union metric. OA is overall accuracy. mAcc is mean accuracy.

over the state of the art (see Table 6). However, we expect future simulations of point cloud scenes to become increasingly dense, in line with increasingly powerful LiDAR scanners for real world applications [14, 40], where our simple baselines can potentially produce significant gain both in accuracy and computation.