

ON UNIVERSAL EQUIVARIANT SET NETWORKS

Anonymous authors

Paper under double-blind review

ABSTRACT

Using deep neural networks that are either invariant or equivariant to permutations in order to learn functions on unordered sets has become prevalent. The most popular, basic models are DeepSets (Zaheer et al., 2017) and PointNet (Qi et al., 2017). While known to be universal for approximating invariant functions, DeepSets and PointNet are not known to be universal when approximating *equivariant* set functions. On the other hand, several recent equivariant set architectures have been proven equivariant universal (Sannai et al., 2019; Keriven & Peyré, 2019), however these models either use layers that are not permutation equivariant (in the standard sense) and/or use higher order tensor variables which are less practical. There is, therefore, a gap in understanding the universality of popular equivariant set models versus theoretical ones.

In this paper we close this gap by proving that: (i) PointNet is not equivariant universal; and (ii) adding a single linear transmission layer makes PointNet universal. We call this architecture PointNetST and argue it is the simplest permutation equivariant universal model known to date. Another consequence is that DeepSets is universal, and also PointNetSeg, a popular point cloud segmentation network (used e.g., in (Qi et al., 2017)) is universal.

The key theoretical tool used to prove the above results is an explicit characterization of all permutation equivariant polynomial layers. Lastly, we provide numerical experiments validating the theoretical results and comparing different permutation equivariant models.

1 INTRODUCTION

Many interesting tasks in machine learning can be described by functions F that take as input a set, $\mathbf{X} = (\mathbf{x}_1, \dots, \mathbf{x}_n)$, and output some per-element features or values, $F(\mathbf{X}) = (F(\mathbf{X})_1, \dots, F(\mathbf{X})_n)$. *Permutation equivariance* is the property required of F so it is well-defined. Namely, it assures that reshuffling the elements in \mathbf{X} and applying F results in the same output, reshuffled in the same manner. For example, if $\tilde{\mathbf{X}} = (\mathbf{x}_2, \mathbf{x}_1, \mathbf{x}_3, \dots, \mathbf{x}_n)$ then $F(\tilde{\mathbf{X}}) = (F(\mathbf{X})_2, F(\mathbf{X})_1, F(\mathbf{X})_3, \dots, F(\mathbf{X})_n)$.

Building neural networks that are permutation equivariant *by construction* proved extremely useful in practice where the most popular models, DeepSets Zaheer et al. (2017) and PointNet Qi et al. (2017) enjoy small number of parameters, low memory footprint and computational efficiency along with high empirical expressiveness. Although both DeepSets and PointNet are known to be invariant universal (i.e., can approximate arbitrary invariant continuous functions) they are not known to be equivariant universal (i.e., can approximate arbitrary equivariant continuous functions).

On the other hand, several researchers have suggested theoretical permutation equivariant models and proved they are equivariant universal. Sannai et al. (2019) builds a universal equivariant network by taking n copies of $(n - 1)$ -invariant networks and combines them with a layer that is not permutation invariant in the standard (above mentioned) sense. Keriven & Peyré (2019) solves a more general problem of building networks that are equivariant universal over arbitrary high order input tensors \mathbb{R}^{n^d} (including graphs); their construction, however, uses higher order tensors as hidden variables which is of less practical value. Yarotsky (2018) proves that neural networks constructed using a finite set of invariant and equivariant polynomial layers are also equivariant universal, however his network is not explicit (i.e., the polynomials are not characterized for the equivariant case) and also of less practical interest due to the high degree polynomial layers.

In this paper we close the gap between the practical and theoretical permutation equivariant constructions and prove:

Theorem 1.

- (i) *PointNet is not equivariant universal.*
- (ii) *Adding a single linear transmission layer (i.e., $\mathbf{X} \mapsto \mathbf{1}\mathbf{1}^T \mathbf{X}$) to PointNet makes it equivariant universal.*
- (iii) *Using ReLU activation the minimal width required for universal permutation equivariant network satisfies $\omega \leq k_{out} + k_{in} + \binom{n+k_{in}}{k_{in}}$.*

This theorem suggests that, arguably, PointNet with an addition of a single linear layer is the simplest universal equivariant network, able to learn arbitrary continuous equivariant functions of sets. An immediate corollary of this theorem is

Corollary 1. *DeepSets and PointNetSeg are universal.*

PointNetSeg is a network used often for point cloud segmentation (e.g., in Qi et al. (2017)). One of the benefit of our result is that it provides a simple characterization of universal equivariant architectures that can be used in the network design process to guarantee universality.

The theoretical tool used for the proof of Theorem 1 is an explicit characterization of the permutation equivariant polynomials over sets of vectors in \mathbb{R}^k using power-sum multi-symmetric polynomials. We prove:

Theorem 2. *Let $\mathbf{P} : \mathbb{R}^{n \times k} \rightarrow \mathbb{R}^{n \times l}$ be a permutation equivariant polynomial map. Then,*

$$\mathbf{P}(\mathbf{X}) = \prod_{j \in [n]} \mathbf{b}_\alpha \mathbf{q}_\alpha^T, \quad (1)$$

where $\mathbf{b}_\alpha = (\mathbf{x}_1^\alpha, \dots, \mathbf{x}_n^\alpha)^T$, $\mathbf{q}_\alpha = (q_{\alpha,1}, \dots, q_{\alpha,l})^T$, where $q_{\alpha,j} = q_{\alpha,j}(s_1, \dots, s_t)$, $t = \binom{n+k}{k}$, are invariant polynomials; $s_j(\mathbf{X}) = \sum_{i=1}^n \mathbf{x}_i^{\alpha_j}$ are the power-sum multi-symmetric polynomials. On the other hand every polynomial map \mathbf{P} satisfying Equation 1 is equivariant.

This theorem, which extends a result in Golubitsky & Stewart (2002) to sets of vectors using multivariate polynomials, lends itself to expressing arbitrary equivariant polynomials as a composition of entry-wise continuous functions and a single linear transmission, which in turn facilitates the proof of Theorem 1.

We conclude the paper by numerical experiments validating the theoretical results and testing several permutation equivariant networks for the tasks of set classification and regression.

2 PRELIMINARIES

Equivariant maps. Vectors $\mathbf{x} \in \mathbb{R}^k$ are by default column vectors; $\mathbf{0}, \mathbf{1}$ are the all zero and all one vectors/tensors; \mathbf{e}_i is the i -th standard basis vector; \mathbf{I} is the identity matrix; all dimensions are inferred from context or mentioned explicitly. We represent a set of n vectors in \mathbb{R}^k as a matrix $\mathbf{X} \in \mathbb{R}^{n \times k}$ and denote $\mathbf{X} = (\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n)^T$, where $\mathbf{x}_i \in \mathbb{R}^k$, $i \in [n]$, are the columns of \mathbf{X} . We denote by S_n the permutation group of $[n]$; its action on \mathbf{X} is defined by $\sigma \cdot \mathbf{X} = (\mathbf{x}_{\sigma^{-1}(1)}, \mathbf{x}_{\sigma^{-1}(2)}, \dots, \mathbf{x}_{\sigma^{-1}(n)})^T$, $\sigma \in S_n$. That is, σ is reshuffling the rows of \mathbf{X} . The natural class of maps assigning a value or feature vector to every element in an input set is permutation equivariant maps:

Definition 1. *A map $\mathbf{F} : \mathbb{R}^{n \times k} \rightarrow \mathbb{R}^{n \times l}$ satisfying $\mathbf{F}(\sigma \cdot \mathbf{X}) = \sigma \cdot \mathbf{F}(\mathbf{X})$ for all $\sigma \in S_n$ and $\mathbf{X} \in \mathbb{R}^{n \times k}$ is called permutation equivariant.*

Power-sum multi-symmetric polynomials. For a vector $\mathbf{x} \in \mathbb{R}^k$ and a multi-index vector $\alpha = (\alpha_1, \dots, \alpha_k) \in \mathbb{N}^k$ we define $\mathbf{x}^\alpha = x_1^{\alpha_1} \dots x_k^{\alpha_k}$, and $|\alpha| = \sum_{i=1}^k \alpha_i$. Given a vector $\mathbf{x} = (x_1, \dots, x_n) \in \mathbb{R}^n$ the power-sum symmetric polynomials $s_j(\mathbf{x}) = \sum_{i=1}^n x_i^j$, with $j \in [n]$, uniquely characterize \mathbf{x} up to permuting its entries. In other words, for $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$ we have $\mathbf{y} = \sigma \cdot \mathbf{x}$ for some $\sigma \in S_n$ if and only if $s_j(\mathbf{y}) = s_j(\mathbf{x})$ for all $j \in [n]$. An equivalent property is that every S_n

invariant polynomial p can be expressed as a polynomial in the power-sum symmetric polynomials, i.e., $p(\mathbf{x}) = q(s_1(\mathbf{x}), \dots, s_n(\mathbf{x}))$, see Rydh (2007) Corollary 8.4, Briand (2004) Theorem 3.

A generalization of the power-sum symmetric polynomials to *matrices* exists and is called power-sum multi-symmetric polynomials, defined with a bit of notation abuse: $s_\alpha(\mathbf{X}) = \sum_{i=1}^n x_i^\alpha$, where $\alpha \in \mathbb{N}^k$ is a multi-index satisfying $|\alpha| \leq n$. Note that the number of power-sum multi-symmetric polynomials acting on $\mathbf{X} \in \mathbb{R}^{n \times k}$ is $t = \binom{n+k}{k}$. For notation simplicity let $\alpha_1, \dots, \alpha_t$ be a list of all $\alpha \in \mathbb{N}^k$ with $|\alpha| \leq n$. Then we index the collection of power-sum multi-symmetric polynomials as s_1, \dots, s_t .

Similarly to the vector case the numbers $s_j(\mathbf{X})$, $j \in [t]$ characterize \mathbf{X} up to permutation of its rows. That is $\mathbf{Y} = \sigma \cdot \mathbf{X}$ for some $\sigma \in S_n$ iff $s_j(\mathbf{Y}) = s_j(\mathbf{X})$ for all $j \in [t]$. Furthermore, every S_n invariant polynomial $p: \mathbb{R}^{n \times k} \rightarrow \mathbb{R}$ can be expressed as a polynomial in the power-sum multi-symmetric polynomials (see (Rydh, 2007) corollary 8.4), i.e.,

$$p(\mathbf{X}) = q(s_1(\mathbf{X}), \dots, s_t(\mathbf{X})), \quad (2)$$

These polynomials were recently used to encode multi-sets in Maron et al. (2019).

3 EQUIVARIANT MULTI-SYMMETRIC POLYNOMIAL LAYERS

In this section we develop the main theoretical tool of this paper, namely, a characterization of all permutation equivariant polynomial layers. As far as we know, these layers were not fully characterized before.

Theorem 2 provides an explicit representation of arbitrary permutation equivariant polynomial maps. Let $\mathbf{P}: \mathbb{R}^{n \times k} \rightarrow \mathbb{R}^{n \times l}$ using the basis of power-sum multi-symmetric polynomials, $s_i(\mathbf{X})$. The particular use of power-sum polynomials $s_i(\mathbf{X})$ has the advantage it can be encoded efficiently using a neural network: as we will show $s_i(\mathbf{X})$ can be approximated using a PointNet with a single linear transmission layer. This allows approximating an arbitrary equivariant polynomial map using PointNet with a single linear transmission layer. In contrast, Yarotsky (2018) also provides a polynomial characterization of equivariant maps (see Lemma 2.1 and Proposition 2.4 in (Yarotsky, 2018)) however there is no formula given to the invariant/equivariant generating set of polynomials and their efficient implementation/approximation is therefore questionable.

A version of this theorem for vectors instead of matrices (i.e., the case of $k = 1$) appears in Golubitsky & Stewart (2002); we extend their proof to matrices, which is the relevant scenario for ML applications as it allow working with sets of vectors.

First, note that it is enough to prove Theorem 1 for $l = 1$ and apply it to every column of \mathbf{P} . Hence, we deal with a vector of polynomials $\mathbf{p}: \mathbb{R}^{n \times k} \rightarrow \mathbb{R}^n$ and need to prove it can be expressed as $\mathbf{p} = \sum_{j \in [n]} \mathbf{b}_j q_j$, for S_n invariant polynomial q_j .

Given a polynomial $p(\mathbf{X})$ and the cyclic permutation $\sigma^{-1} = (123 \dots n)$ the following operation, taking a polynomial to a vector of polynomials, is useful in characterizing equivariant polynomial maps:

$$[p](\mathbf{X}) = \begin{matrix} \circ & & 1 \\ \text{B} & p(\mathbf{X}) & \text{C} \\ \text{B} & p(\sigma \cdot \mathbf{X}) & \text{C} \\ \text{B} & p(\sigma^2 \cdot \mathbf{X}) & \text{C} \\ \text{B} & \vdots & \text{C} \\ \text{B} & p(\sigma^{n-1} \cdot \mathbf{X}) & \text{C} \end{matrix} \quad (3)$$

Theorem 1 will be proved using the following two lemmas:

Lemma 1. *Let $\mathbf{p}: \mathbb{R}^{n \times k} \rightarrow \mathbb{R}^n$ be an equivariant polynomial map. Then, there exists a polynomial $p: \mathbb{R}^{n \times k} \rightarrow \mathbb{R}$, invariant to S_{n-1} (permuting the last $n-1$ rows of \mathbf{X}) so that $\mathbf{p} = [p]$.*

Proof. Equivariance of \mathbf{p} means that for all $\sigma \in S_n$ it holds that

$$\sigma \cdot \mathbf{p}(\mathbf{X}) = \mathbf{p}(\sigma \cdot \mathbf{X}). \quad (4)$$

Choosing an arbitrary permutation $\sigma \in \text{stab}(1) < S_n$, namely a permutation satisfying $\sigma(1) = 1$, and observing the first row in Equation 4 we get $p_1(\mathbf{X}) = p_1(\sigma \cdot \mathbf{X}) = p_1(\mathbf{x}_1, \mathbf{x}_{\sigma^{-1}(2)}, \dots, \mathbf{x}_{\sigma^{-1}(n)})$. Since this is true for all $\sigma \in \text{stab}(1)$ p_1 is S_{n-1} invariant. Next, applying $\sigma = (1i)$ to Equation 4 and observing the first row again we get $p_i(\mathbf{X}) = p_1(\mathbf{x}_i, \dots, \mathbf{x}_1, \dots)$. Using the invariance of p_1 to S_{n-1} we get $\mathbf{p} = \lceil p_1 \rceil$. \square

Lemma 2. Let $p : \mathbb{R}^{n-k} \rightarrow \mathbb{R}$ be a polynomial invariant to S_{n-1} (permuting the last $n-1$ rows of \mathbf{X}) then

$$p(\mathbf{X}) = \prod_{j \in \alpha_j} \mathbf{x}_1^{\alpha_j} q_{\alpha_j}(\mathbf{X}), \quad (5)$$

where q_{α} are S_n invariant.

Proof. Expanding p with respect to \mathbf{x}_1 we get

$$p(\mathbf{X}) = \sum_{j \in \alpha_j} \mathbf{x}_1^{\alpha_j} p_{\alpha_j}(\mathbf{x}_2, \dots, \mathbf{x}_n), \quad (6)$$

for some $m \in \mathbb{N}$. We first claim p_{α} are S_{n-1} invariant. Indeed, note that if $p(\mathbf{X}) = p(\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n)$ is S_{n-1} invariant, i.e., invariant to permutations of $\mathbf{x}_2, \dots, \mathbf{x}_n$, then also its derivatives $\frac{\partial^{|\beta|}}{\partial \mathbf{x}_1^{\beta}} p(\mathbf{X})$ are S_{n-1} permutation invariant, for all $\beta \in \mathbb{N}^k$. Taking the derivative $\partial^{|\beta|} / \partial \mathbf{x}_1^{\beta}$ on both sides of Equation 6 we get that p_{β} is S_{n-1} equivariant.

For brevity denote $p = p_{\alpha}$. Since p is S_{n-1} invariant it can be expressed as a polynomial in the power-sum multi-symmetric polynomials, i.e., $p(\mathbf{x}_2, \dots, \mathbf{x}_n) = r(s_1(\mathbf{x}_2, \dots, \mathbf{x}_n), \dots, s_t(\mathbf{x}_2, \dots, \mathbf{x}_n))$. Note that $s_i(\mathbf{x}_2, \dots, \mathbf{x}_n) = s_i(\mathbf{X}) - \mathbf{x}_1^{\alpha_i}$ and therefore

$$p(\mathbf{x}_2, \dots, \mathbf{x}_n) = r(s_1(\mathbf{X}) - \mathbf{x}_1^{\alpha_1}, \dots, s_t(\mathbf{X}) - \mathbf{x}_1^{\alpha_t}).$$

Since p is a polynomial, expanding its monomials in $s_i(\mathbf{X})$ and $\mathbf{x}_1^{\alpha_i}$ shows p can be expressed as $p = \sum_{j \in \alpha_j} \mathbf{x}_1^{m^j} \tilde{p}_j$, where $m^j \in \mathbb{N}$, and \tilde{p}_j are S_n invariant (as multiplication of invariant S_n polynomials $s_i(\mathbf{X})$). Plugging this in Equation 6 we get Equation 5, possibly with the sum over some $n' > n$. It remains to show n^j can be taken to be at-most n . This is proved in Corollary 5 in Briand (2004) \square

Proof. (Theorem 2) Given an equivariant \mathbf{p} as above, use Lemma 1 to write $\mathbf{p} = \lceil p \rceil$ where $p(\mathbf{X})$ is invariant to permuting the last $n-1$ rows of \mathbf{X} . Use Lemma 2 to write $p(\mathbf{X}) = \prod_{j \in \alpha_j} \mathbf{x}_1^{\alpha_j} q_{\alpha_j}(\mathbf{X})$, where q_{α} are S_n invariant. We get,

$$\mathbf{p} = \lceil p \rceil = \prod_{j \in \alpha_j} \mathbf{b}_{\alpha_j} q_{\alpha_j}.$$

The converse direction is immediate after noting that \mathbf{b}_{α} are equivariant and q_{α} are invariant. \square

4 UNIVERSALITY OF SET EQUIVARIANT NEURAL NETWORKS

We consider equivariant deep neural networks $f : \mathbb{R}^{n_{in}} \rightarrow \mathbb{R}^{n_{out}}$,

$$\mathbf{F}(\mathbf{X}) = \mathbf{L}_m \circ \nu \circ \dots \circ \nu \circ \mathbf{L}_1(\mathbf{X}), \quad (7)$$

where $\mathbf{L}_i : \mathbb{R}^{n_{k_i}} \rightarrow \mathbb{R}^{n_{k_{i+1}}}$ are affine equivariant transformations, and ν is an entry-wise non-linearity (e.g., ReLU). We define the width of the network to be $\omega = \max_i k_i$; note that this definition is different from the one used for standard MLP where the width would be $n\omega$, see e.g., Hanin & Sellke (2017). Zaheer et al. (2017) proved that affine equivariant \mathbf{L}_i are of the form

$$\mathbf{L}_i(\mathbf{X}) = \mathbf{X}\mathbf{A} + \frac{1}{n} \mathbf{1}\mathbf{1}^T \mathbf{X}\mathbf{B} + \mathbf{1}\mathbf{c}^T, \quad (8)$$

where $\mathbf{A}, \mathbf{B} \in \mathbb{R}^{k_i \times k_{i+1}}$, and $\mathbf{c} \in \mathbb{R}^{k_{i+1}}$ are the layer's trainable parameters; we call the linear transformation $\mathbf{X} \mapsto \frac{1}{n} \mathbf{1}\mathbf{1}^T \mathbf{X}\mathbf{B}$ a linear transmission layer. Equation 7 with the choice of layers as in Equation 8 is the DeepSets architecture (Zaheer et al., 2017). Taking $\mathbf{B} = 0$ in all layers is

the PointNet architecture (Qi et al., 2017). Another type of architecture of interest is PointNetSeg appearing in (Qi et al., 2017). Variations of this architecture are used for Point cloud segmentation, see e.g. Li et al. (2018) and Engelmann et al. (2017). PointNetSeg uses an invariant version of PointNet (i.e., PointNet composed with an invariant max layer) and concatenates its output as a constant feature to an intermediate layer which is then inputed to another PointNet. Lastly, we will also consider a model we call PointNetST that is PointNet with a Single (linear) Transmission layer; in more details, PointNetST is an equivariant model of the form Equation 7 with layer as in Equation 8 where only a single layer L_i has a non-zero B (see Equation 8). We will prove PointNetST is permutation equivariant universal and therefore arguably the simplest permutation equivariant universal model known to date.

Universality of equivariant deep networks is defined next.

Definition 2. *Permutation equivariant universality¹ of a model $F : \mathbb{R}^n \times \mathbb{R}^{k_{in}} \rightarrow \mathbb{R}^n \times \mathbb{R}^{k_{out}}$ means that for every permutation equivariant continuous function $H : \mathbb{R}^n \times \mathbb{R}^{k_{in}} \rightarrow \mathbb{R}^n \times \mathbb{R}^{k_{out}}$ defined over the cube $K = [0, 1]^n \times \mathbb{R}^{k_{in}} \subset \mathbb{R}^n \times \mathbb{R}^{k_{in}}$, and $\epsilon > 0$ there exists a choice of m (i.e., network depth), k_i (i.e., network width) and the trainable parameters of F so that $\|H(\mathbf{X}) - F(\mathbf{X})\|_7 < \epsilon$ for all $\mathbf{X} \in K$.*

We prove our main theorem next.

Proof. (Theorem 1) Fact (i), namely that PointNet is not equivariant universal is a consequence of the following simple lemma:

Lemma 3. *Let $h = (h_1, \dots, h_n)^T : \mathbb{R}^n \rightarrow \mathbb{R}^n$ be the equivariant linear function defined by $h(\mathbf{x}) = \mathbf{1}\mathbf{1}^T \mathbf{x}$. There is no $f : \mathbb{R} \rightarrow \mathbb{R}$ so that $|h_i(\mathbf{x}) - f(x_i)| < \frac{1}{2}$ for all $i \in [n]$ and $\mathbf{x} \in [0, 1]^n$.*

Proof. Assume such f exists. Let $\mathbf{e}_1 = (1, 0, \dots, 0)^T \in \mathbb{R}^n$. Then,

$$1 = |h_2(\mathbf{e}_1) - h_2(\mathbf{0})| \leq |h_2(\mathbf{e}_1) - f(0)| + |f(0) - h_2(\mathbf{0})| < 1$$

reaching a contradiction. \square

To prove (ii) we first reduce the problem from the class of all continuous equivariant functions to the class of equivariant polynomials. This is justified by the following lemma.

Lemma 4. *Equivariant polynomials $P : \mathbb{R}^n \times \mathbb{R}^{k_{in}} \rightarrow \mathbb{R}^n \times \mathbb{R}^{k_{out}}$ are dense in the space of continuous equivariant functions $F : \mathbb{R}^n \times \mathbb{R}^{k_{in}} \rightarrow \mathbb{R}^n \times \mathbb{R}^{k_{out}}$ over the cube K .*

Proof. Take an arbitrary $\epsilon > 0$. Consider the function $f_{ij} : \mathbb{R}^n \times \mathbb{R}^{k_{in}} \rightarrow \mathbb{R}$, which denotes the (i, j) -th output entry of F . By the Stone-Weierstrass Theorem there exists a polynomial $p_{ij} : \mathbb{R}^n \times \mathbb{R}^{k_{in}} \rightarrow \mathbb{R}$ such that $\|f_{ij}(\mathbf{X}) - p_{ij}(\mathbf{X})\|_7 \leq \epsilon$ for all $\mathbf{X} \in K$. Consider the polynomial map $P : \mathbb{R}^n \times \mathbb{R}^{k_{in}} \rightarrow \mathbb{R}^n \times \mathbb{R}^{k_{out}}$ defined by $(P)_{ij} = p_{ij}$. P is in general not equivariant. To finish the proof we will symmetrize P :

$$\begin{aligned} F(\mathbf{X}) - \frac{1}{n!} \times_{\sigma \in S_n} \sigma \cdot P(\sigma^{-1} \cdot \mathbf{X}) &= \frac{1}{n!} \times_{\sigma \in S_n} \sigma \cdot F(\sigma^{-1} \cdot \mathbf{X}) - \frac{1}{n!} \times_{\sigma \in S_n} \sigma \cdot P(\sigma^{-1} \cdot \mathbf{X}) \\ &= \frac{1}{n!} \times_{\sigma \in S_n} \sigma \cdot F(\sigma^{-1} \cdot \mathbf{X}) - P(\sigma^{-1} \cdot \mathbf{X}) \\ &\leq \frac{1}{n!} \times_{\sigma \in S_n} \epsilon \\ &= \epsilon, \end{aligned}$$

where in the first equality we used the fact that F is equivariant. This concludes the proof since $\frac{1}{n!} \times_{\sigma \in S_n} \sigma \cdot P(\sigma^{-1} \cdot \mathbf{X})$ is an equivariant polynomial map. \square

Now, according to Theorem 2 an arbitrary equivariant polynomial $P : \mathbb{R}^n \times \mathbb{R}^{k_{in}} \rightarrow \mathbb{R}^n \times \mathbb{R}^{k_{out}}$ can be written as $P = \sum_{\alpha, j} b_{\alpha}(\mathbf{X}) q_{\alpha}(\mathbf{X})^T$, where $b_{\alpha}(\mathbf{X}) = [x_1^{\alpha}] \in \mathbb{R}^n$ and $q_{\alpha} =$

¹Or just *equivariant universal* in short.

$(q_{\alpha,1}, \dots, q_{\alpha,k_{out}}) \in \mathbb{R}^{k_{out}}$ are invariant polynomials. Remember that every S_n invariant polynomial can be expressed as a polynomial in the $t = \binom{n+k_{in}}{k_{in}}$ power-sum multi-symmetric polynomials $s_j(\mathbf{X}) = \frac{1}{n} \prod_{i=1}^n \mathbf{x}_i^{\alpha_j}$, $j \in [t]$ (we use the $1/n$ normalized version for a bit more simplicity later on). We can therefore write \mathbf{P} as composition of three maps:

$$\mathbf{P} = \mathbf{Q} \circ \mathbf{L} \circ \mathbf{B}, \quad (9)$$

where $\mathbf{B} : \mathbb{R}^{n \times k_{in}} \rightarrow \mathbb{R}^{n \times t}$ is defined by

$$\mathbf{B}(\mathbf{X}) = (\mathbf{b}(\mathbf{x}_1), \dots, \mathbf{b}(\mathbf{x}_n))^T,$$

$\mathbf{b}(\mathbf{x}) = (\mathbf{x}^{\alpha_1}, \dots, \mathbf{x}^{\alpha_t})$; \mathbf{L} is defined as in Equation 8 with $\mathbf{B} = [\mathbf{0}, \mathbf{I}]$ and $\mathbf{A} = [e_1, \dots, e_{k_{in}}, \mathbf{0}]$, where $\mathbf{I} \in \mathbb{R}^{t \times t}$ the identity matrix and $e_i \in \mathbb{R}^t$ represents the standard basis (as-usual). We assume $\alpha_j = e_j \in \mathbb{R}^{k_{in}}$, for $j \in [k_{in}]$. Note that the output of \mathbf{L} is of the form

$$\mathbf{L}(\mathbf{B}(\mathbf{X})) = (\mathbf{X}, \mathbf{1}_{s_1}(\mathbf{X}), \mathbf{1}_{s_2}(\mathbf{X}), \dots, \mathbf{1}_{s_t}(\mathbf{X})).$$

Finally, $\mathbf{Q} : \mathbb{R}^{n \times (k_{in}+t)} \rightarrow \mathbb{R}^{n \times k_{out}}$ is defined by

$$\mathbf{Q}(\mathbf{X}, \mathbf{1}_{s_1}, \dots, \mathbf{1}_{s_t}) = (\mathbf{q}(\mathbf{x}_1, s_1, \dots, s_t), \dots, \mathbf{q}(\mathbf{x}_n, s_1, \dots, s_t))^T,$$

and $\mathbf{q}(\mathbf{x}, s_1, \dots, s_t) = \prod_{j \in [t]} \mathbf{x}^{\alpha_j} q_{\alpha_j}(s_1, \dots, s_t)^T$.

The decomposition in Equation 9 of \mathbf{P} suggests that replacing \mathbf{Q}, \mathbf{B} with Multi-Layer Perceptrons (MLPs) would lead to a universal permutation equivariant network consisting of PointNet with a single linear transmission layer, namely PointNetST.

The \mathbf{F} approximating \mathbf{P} will be defined as

$$\mathbf{F} = \circ \mathbf{L} \circ , \quad (10)$$

where $\circ : \mathbb{R}^{n \times k_{in}} \rightarrow \mathbb{R}^{n \times t}$ and $\circ : \mathbb{R}^{n \times (t+k_{in})} \rightarrow \mathbb{R}^{n \times k_{out}}$ are both of PointNet architecture, namely there exist MLPs $\phi : \mathbb{R}^{k_{in}} \rightarrow \mathbb{R}^t$ and $\psi : \mathbb{R}^{t+k_{in}} \rightarrow \mathbb{R}^{k_{out}}$ so that $\mathbf{B}(\mathbf{X}) = (\phi(\mathbf{x}_1), \dots, \phi(\mathbf{x}_n))^T$ and $\mathbf{L}(\mathbf{B}(\mathbf{X})) = (\psi(\mathbf{x}_1), \dots, \psi(\mathbf{x}_n))^T$. See Figure 1 for an illustration of \mathbf{F} .

To build the MLPs ϕ, ψ we will first construct ψ to approximate \mathbf{q} , that is, we use the universality of MLPS (see (Hornik, 1991; Sonoda & Murata, 2017; Hanin & Sellke, 2017)) to construct ψ so that $\|\psi(\mathbf{x}, s_1, \dots, s_t) - \mathbf{q}(\mathbf{x}, s_1, \dots, s_t)\|_7 < \frac{\epsilon}{2}$ for all $(\mathbf{x}, s_1, \dots, s_t) \in [0, 1]^{k_{in}+t}$. Furthermore, as ψ over $[0, 1]^{k_{in}+t}$ is uniformly continuous, let δ be such that if $\mathbf{z}, \mathbf{z}^0 \in [0, 1]^{k_{in}+t}$, $\|\mathbf{z} - \mathbf{z}^0\|_7 < \delta$ then $\|\psi(\mathbf{z}) - \psi(\mathbf{z}^0)\|_7 < \frac{\epsilon}{2}$. Now, we use universality again to construct ϕ approximating \mathbf{b} , that is we take ϕ so that $\|\phi(\mathbf{x}) - \mathbf{b}(\mathbf{x})\|_7 < \delta$ for all $\mathbf{x} \in [0, 1]^{k_{in}}$.

$$\|\mathbf{F}(\mathbf{X}) - \mathbf{P}(\mathbf{X})\|_7 \leq \| \mathbf{L}(\mathbf{B}(\mathbf{X})) - \mathbf{L}(\mathbf{B}(\mathbf{X})) \|_7 + \| \mathbf{L}(\mathbf{B}(\mathbf{X})) - \mathbf{Q}(\mathbf{L}(\mathbf{B}(\mathbf{X}))) \|_7 \\ = \text{err}_1 + \text{err}_2$$

First, $\|\mathbf{L}(\mathbf{B}(\mathbf{X})) - \mathbf{L}(\mathbf{B}(\mathbf{X}))\|_7 < \delta$ for all $\mathbf{X} \in K$ and therefore $\text{err}_1 < \frac{\epsilon}{2}$. Second, note that if $\mathbf{X} \in K$ then $\mathbf{B}(\mathbf{X}) \in [0, 1]^{n \times t}$ and $\mathbf{L}(\mathbf{B}(\mathbf{X})) \in [0, 1]^{n \times (k_{in}+t)}$. Therefore by construction of ψ we have $\text{err}_2 < \frac{\epsilon}{2}$.

To prove (iii) we use the result in Hanin & Sellke (2017) (see Theorem 1) bounding the width of an MLP approximating a function $f : [0, 1]^{d_{in}} \rightarrow \mathbb{R}^{d_{out}}$ by $d_{in} + d_{out}$. Therefore, the width of the MLP ϕ is bounded by $k_{in} + t$, where the width of the MLP ψ is bounded by $t + k_{in} + k_{out}$, proving the bound. \square

Before we prove Corollary 1 let us recall the PointNetSeg model from Qi et al. (2017) in detail. We write the model \mathbf{N} as follows: First, we apply $\mathbf{L}(\mathbf{B}(\mathbf{X})) = (\psi(\mathbf{x}_1), \dots, \psi(\mathbf{x}_n))^T$. Let $(\eta(\mathbf{x}_1), \dots, \eta(\mathbf{x}_n))^T \in \mathbb{R}^{n \times k_1}$ be the output of the first layer of the mlp and let $\max(\mathbf{L}(\mathbf{B}(\mathbf{X}))) \in \mathbb{R}^{k_1}$ be the vector where the i 'th coordinate is the maximal value over the i 'th column in $\mathbf{L}(\mathbf{B}(\mathbf{X}))$. Then $\mathbf{N}(\mathbf{X}) = ((\eta(\mathbf{x}_1), \max(\mathbf{L}(\mathbf{B}(\mathbf{X}))), \dots, (\eta(\mathbf{x}_n), \max(\mathbf{L}(\mathbf{B}(\mathbf{X}))))^T$, with $\mathbf{L}(\mathbf{B}(\mathbf{X})) = (\phi(\mathbf{x}_1), \dots, \phi(\mathbf{x}_n))^T$.

We can now prove Corollary 1.




Figure 1: The construction of the universal network (PointNetST).

Proof. (Corollary 1)

The fact that the DeepSets model is equivariant universal is immediate. Indeed, The PointNetST model can be obtained from the DeepSets model by setting $\mathbf{B} = 0$ in all but one layer, with \mathbf{B} as in Equation 8.

For the PointNetSeg model note that by Theorem 1 in Qi et al. (2017) every invariant function $f : \mathbb{R}^{n \times k_{in}} \rightarrow \mathbb{R}^t$ can be approximated by a network of the form

$$\bar{\phi} \circ \max \circ \bar{\psi}$$

Where $\bar{\psi}(x_1, \dots, x_n) = (\bar{\psi}(x_1), \dots, \bar{\psi}(x_n))^T \in \mathbb{R}^{n \times K}$, with $\bar{\psi} : \mathbb{R}^{k_{in}} \rightarrow \mathbb{R}^K$, $\bar{\phi} : \mathbb{R}^K \rightarrow \mathbb{R}^t$ multi layer perceptrons. In particular, for every $\varepsilon > 0$ there exists $\bar{\phi}, \bar{\psi}$ such that $\|\bar{\phi} \circ \max \circ \bar{\psi}(\mathbf{X}) - (s_1(\mathbf{X}), \dots, s_t(\mathbf{X}))\|_1 < \varepsilon$ for every $\mathbf{X} \in [0, 1]^{n \times k_{in}}$ where $s_1(\mathbf{X}), \dots, s_t(\mathbf{X})$ are the power sum multi symmetric polynomials. The rest of the proof follows closely the proof of Theorem 1. \square

5 EXPERIMENTS

We conducted experiments in order to validate our theoretical observations. We compared the results of several equivariant models, as well as baseline (full) MLP, on three equivariant learning tasks: a classification task (knapsack) and two regression tasks (squared norm and Fiedler eigen vector). For all tasks we compare results of 6 different models: DeepSets, PointNet, PointNetSeg, PointNetST, and PointNetQT. PointNetQT is PointNet with a single quadratic equivariant transmission layer as defined in the appendix. In all experiments we used a network of the form Equation 7 with $m = 6$ depth and varying width, fixed across all layers.

Equivariant classification. For classification, we chose to learn the multidimensional knapsack problem, which is known to be NP-hard. We are given a set of 4-vectors, represented by $\mathbf{X} \in \mathbb{R}^{n \times 4}$, and our goal is to learn the equivariant classification function $f : \mathbb{R}^{n \times 4} \rightarrow \{0, 1\}^n$ defined by the following optimization problem:

$$\begin{aligned} f(\mathbf{X}) = \arg \max_z & \sum_{i=1}^n x_{i1} z_i \\ \text{s.t.} & \sum_{i=1}^n x_{ij} z_i \leq w_j, \quad j = 2, 3, 4 \\ & z_i \in \{0, 1\}, \quad i \in [n] \end{aligned}$$

Intuitively, given a set of vectors $\mathbf{X} \in \mathbb{R}^{n \times 4}$, $(\mathbf{X})_{ij} = x_{ij}$, where each row represents an element in a set, our goal is to find a subset maximizing the value while satisfying budget constraints. The first column of \mathbf{X} defines the value of each element, and the three other columns the costs. In Subsection 5.1 we detail how we generated this dataset.

Equivariant regression. The first equivariant function we considered for regression is the function $f(\mathbf{X}) = \frac{1}{n} \sum_{i=1}^n \sum_{j=1}^k (\mathbf{X}_{i,j} - \frac{1}{2})^2$. Hanin & Sellke (2017) showed this function cannot be approximated by MLPs of small width. We drew $10k$ training examples and $1k$ test examples i.i.d. from a $\mathcal{N}(\frac{1}{2}, 1)$ distribution (per entry of \mathbf{X}).

The second equivariant function we considered is defined on point clouds $\mathbf{X} \in \mathbb{R}^{n \times 3}$. For each point cloud we computed a graph by connecting every point to its 10 nearest neighbors. We then computed the absolute value of the first non trivial eigenvector of the graph Laplacian. We used the ModelNet dataset (Wu et al., 2015) which contains $\sim 9k$ training meshes and $\sim 2k$ test meshes. The point clouds are generated by randomly sampling 512 points from each mesh.

Figure 2 summarizes train and test accuracy of the 6 models after training (training details in Subsection 5.1) as a function of the network width ω . We have tested 15 ω values equidistant in $[5, \frac{nk_{in}}{2}]$.

As can be see in the graphs, in all three datasets the equivariant universal models (PointNetST, PointNetQT, DeepSets, PointNetSeg) achieved comparable accuracy. PointNet, which is not

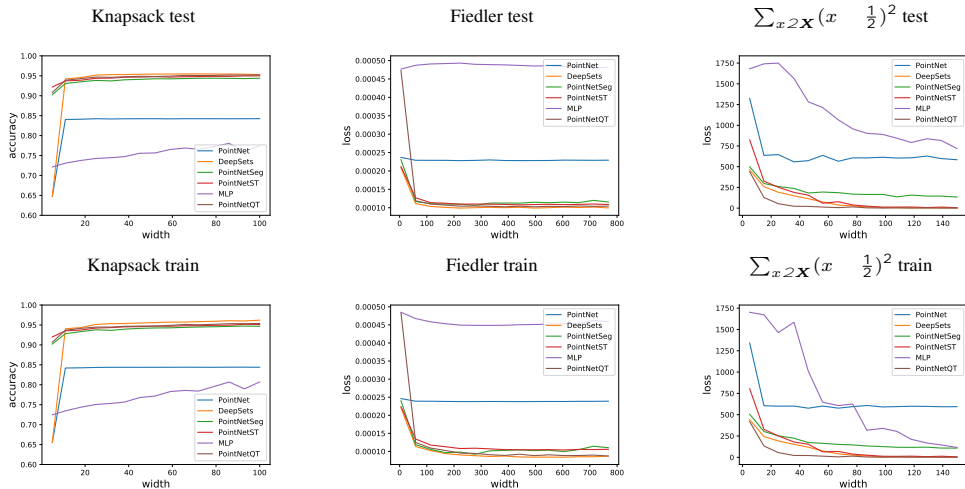


Figure 2: Classification and regression tasks with permutation equivariant models. All the universal permutation equivariant models perform similarly, while the equivariant non-universal PointNet demonstrates reduced performance consistently; MLP baseline (with the same number of parameters as the equivariant models) performs poorly.

equivariant universal, consistently achieved inferior performance compared to the universal models, as expected by the theory developed in this paper. The non-equivariant MLP, although universal, used the same width (i.e., same number of parameters) as the equivariant models and was able to over-fit only on one train set (the quadratic function); its performance on the test sets was inferior by a large margin to the equivariant models.

An interesting point is that although the width used in the experiments is much smaller than the bound $k_{out} + k_{in} + \binom{n+k_{in}}{k_{in}}$ established by Theorem 1, the universal models are still able to learn well the functions we tested on. This raises the question of the tightness of this bound, which we leave to future work.

Lastly, we see that adding a quadratic transmission layer (PointNetQT) has a marginal effect on the model performance, while adding a single linear transmission layer to a PointNet network gives similar results to the DeepSets architecture, while having lower memory footprint.

5.1 IMPLEMENTATION DETAILS

Knapsack data generation. We constructed a dataset of $10k$ training examples and $1k$ test examples consisting of 50×4 matrices. We took $w_1 = 100, w_2 = 80, w_3 = 50$. To generate $\mathbf{X} \in \mathbb{R}^{50 \times 4}$, we draw an integer uniformly at random between 1 and 100 and randomly choose 50 integers between 1 as the first column of \mathbf{X} . We also randomly chose an integer between 1 and 25 and then randomly chose 150 integers in that range as the three last columns of \mathbf{X} . The labels for each input \mathbf{X} were computed by a standard dynamic programming approach, see Martello & Toth (1990).

Optimization. We implemented the experiments in Pytorch Paszke et al. (2017) with the Adam Kingma & Ba (2014) optimizer for learning. For the classification we used the cross entropy loss and trained for 150 epochs with learning rate 0.001, learning rate decay of 0.5 every 100 epochs and batch size 32. For the quadratic function regression we trained for 150 epochs with learning rate of 0.001, learning rate decay 0.1 every 50 epochs and batch size 64; for the regression to the leading eigen vector we trained for 50 epochs with learning rate of 0.001 and batch size 32.

REFERENCES

- Emmanuel Briand. When is the algebra of multisymmetric polynomials generated by the elementary multisymmetric polynomials. *Beiträge zur Algebra und Geometrie*, 45, 01 2004.
- Francis Engelmann, Theodora Kontogianni, Alexander Hermans, and Bastian Leibe. Exploring spatial context for 3d semantic segmentation of point clouds. In *IEEE International Conference on Computer Vision, 3DRMS Workshop, ICCV*, 2017.
- Martin Golubitsky and Ian Stewart. The symmetry perspective. 2002.
- Boris Hanin and Mark Sellke. Approximating continuous functions by relu nets of minimal width. *CoRR*, abs/1710.11278, 2017. URL <http://arxiv.org/abs/1710.11278>.
- Kurt Hornik. Approximation capabilities of multilayer feedforward networks. *Neural networks*, 4(2): 251–257, 1991.
- Nicolas Keriven and Gabriel Peyré. Universal invariant and equivariant graph neural networks. *arXiv preprint arXiv:1905.04943*, 2019.
- Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Jiaxin Li, Ben M Chen, and Gim Hee Lee. So-net: Self-organizing network for point cloud analysis. *arXiv preprint arXiv:1803.04249*, 2018.
- Haggai Maron, Heli Ben-Hamu, Hadar Serviansky, and Yaron Lipman. Provably powerful graph networks. *CoRR*, abs/1905.11136, 2019. URL <http://arxiv.org/abs/1905.11136>.
- Silvano Martello and Paolo Toth. *Knapsack Problems: Algorithms and Computer Implementations*. John Wiley & Sons, Inc., New York, NY, USA, 1990. ISBN 0-471-92420-2.
- Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in PyTorch. In *NIPS Autodiff Workshop*, 2017.
- 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*, pp. 652–660, 2017.
- David Rydh. A minimal set of generators for the ring of multisymmetric functions. In *Annales de l’institut Fourier*, volume 57, pp. 1741–1769, 2007.
- Akiyoshi Sannai, Yuuki Takai, and Matthieu Cordonnier. Universal approximations of permutation invariant/equivariant functions by deep neural networks. *arXiv preprint arXiv:1903.01939*, 2019.
- Sho Sonoda and Noboru Murata. Neural network with unbounded activation functions is universal approximator. *Applied and Computational Harmonic Analysis*, 43(2):233–268, 2017.
- Zhirong Wu, Shuran Song, Aditya Khosla, Fisher Yu, Linguang Zhang, Xiaoou Tang, and Jianxiong Xiao. 3d shapenets: A deep representation for volumetric shapes. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2015.
- Dmitry Yarotsky. Universal approximations of invariant maps by neural networks. *arXiv preprint arXiv:1804.10306*, 2018.
- Manzil Zaheer, Satwik Kottur, Siamak Ravanbakhsh, Barnabas Poczos, Ruslan R Salakhutdinov, and Alexander J Smola. Deep sets. In *Advances in neural information processing systems*, pp. 3391–3401, 2017.

A APPENDIX

One potential application of Theorem 2 is augmenting an equivariant neural network (Equation 7) with equivariant polynomial layers $\mathbf{P} : \mathbb{R}^{n \times k} \rightarrow \mathbb{R}^{n \times l}$ of some maximal degree d . This can be done in the following way: look for all solutions to $\alpha, \beta_1, \beta_2, \dots \in \mathbb{N}^k$ so that $|\alpha| + \sum_{i=1}^n |\beta_i| \leq d$. Any solution to this equation will give a basis element of the form $\mathbf{p}(\mathbf{X}) = [\mathbf{x}_1^\alpha] \odot \prod_{i=1}^n \mathbf{x}_i^{\beta_i}$.

In the paper we tested PointNetQT, an architecture that adds to PointNet a single quadratic equivariant layer. We opted to use only the quadratic transmission operators: For a matrix $\mathbf{X} \in \mathbb{R}^{n \times k}$ we define $L(\mathbf{X}) \in \mathbb{R}^{n \times k}$ as follows:

$$L(\mathbf{X}) = \mathbf{X}\mathbf{W}_1 + \mathbf{1}\mathbf{1}^T\mathbf{X}\mathbf{W}_2 + (\mathbf{1}\mathbf{1}^T\mathbf{X}) \odot (\mathbf{1}\mathbf{1}^T\mathbf{X})\mathbf{W}_3 + (\mathbf{X} \odot \mathbf{X})\mathbf{W}_4 + (\mathbf{1}\mathbf{1}^T\mathbf{X}) \odot \mathbf{X}\mathbf{W}_5,$$

where \odot is a pointwise multiplication and $\mathbf{W}_i \in \mathbb{R}^{n \times k}, i \in [5]$ are the learnable parameters.