
Equivariance Through Parameter-Sharing

Siamak Ravanbakhsh¹ Jeff Schneider¹ Barnabás Póczos¹

Abstract

We propose to study equivariance in deep neural networks through parameter symmetries. In particular, given a group \mathcal{G} that acts discretely on the input and output of a standard neural network layer $\phi_{\mathbf{W}} : \mathbb{R}^M \rightarrow \mathbb{R}^N$, we show that $\phi_{\mathbf{W}}$ is equivariant with respect to \mathcal{G} -action iff \mathcal{G} explains the symmetries of the network parameters \mathbf{W} . Inspired by this observation, we then propose two parameter-sharing schemes to induce the desirable symmetry on \mathbf{W} . Our procedure for tying the parameters achieves \mathcal{G} -equivariance and, under some conditions on the action of \mathcal{G} , it guarantees sensitivity to all other permutation groups outside \mathcal{G} .

Given enough training data, a multi-layer perceptron would eventually learn the domain invariances in a classification task. Nevertheless, success of convolutional and recurrent networks suggests that encoding the domain symmetries through shared parameters can significantly boost the generalization of deep neural networks. The same observation can be made in deep learning for semi-supervised and unsupervised learning in structured domains. This raises an important question that is addressed in this paper: **What kind of priors on input/output structure can be encoded through parameter-sharing?**

This work is an attempt at answering this question, when our priors are in the form discrete domain symmetries. To formalize this type of prior, a family of transformations of input and output to a neural layer are expressed as group “action” on the input and output. The resulting neural network is *invariant* to this action, if transformations of the input within that particular family, does not change the output (e.g., rotation-invariance). However, if the output is transformed, in a predictable way, as we transform the input, the neural layer is *equivariant* to the action of the group.

¹School of Computer Science, Carnegie Mellon University, 5000 Forbes Ave., Pittsburgh, PA, USA 15217. Correspondence to: Siamak Ravanbakhsh <mravanba@cs.cmu.edu>.

Our goal is to show that parameter-sharing can be used to achieve equivariance to any discrete group action.

Application of group theory in machine learning has been the topic of various works in the past (e.g., Kondor, 2008; Bartók et al., 2010). In particular, many probabilistic inference techniques have been extended to graphical models with known symmetry groups (Raedt et al., 2016; Kersting et al., 2009; Bui et al., 2012; Niepert, 2012). Deep and hierarchical models have used a variety of techniques to study or obtain representations that isolate transformations from the “content” (e.g., Hinton et al., 2011; Jayaraman & Grauman, 2015; Lenc & Vedaldi, 2015; Agrawal et al., 2015). The simplest method of achieving equivariance is through data-augmentation (Krizhevsky et al., 2012; Dieleman et al., 2015). Going beyond augmentation, several methods directly apply the group-action, in one way or another, by transforming the data or *its encodings* using group members (Jaderberg et al., 2015; Anselmi et al., 2013; Dieleman et al., 2016). An alternative path to invariance via harmonic analysis. In particular cascade of wavelet transforms is investigated in (Bruna & Mallat, 2013; Oyallon & Mallat, 2015; Sifre & Mallat, 2013). More recently (Cohen & Welling, 2016b) study steerable filters (e.g., Freeman et al., 1991; Hel-Or & Teo, 1998) as a general mean for achieving equivariance in deep networks. Invariance and equivariance through parameter-sharing is also discussed in several prior works (Cohen & Welling, 2016a; Gens & Domingos, 2014).

The desirability of using parameter-sharing for this purpose is mainly due to its simplicity and computational efficiency. However, it also suggests possible directions for *discovering* domain symmetries through regularization schemes. Following the previous work on the study of symmetry in deep networks, we rely on group theory and group-actions to formulate invariances and equivariances of a function. Due to discrete nature of parameter-sharing, our treatment here is limited to permutation groups. Action of a permutation group \mathcal{G} can model discrete transformations of a set of variables, such as translation and 90° rotation of pixels around any center in an image. If the output of a function transforms with a \mathcal{G} -action as we transform its input with a different \mathcal{G} -action, the function is equivariant with respect to action of \mathcal{G} . For example, in a convolution layer, as we translate the input, the feature-maps are also translated. If

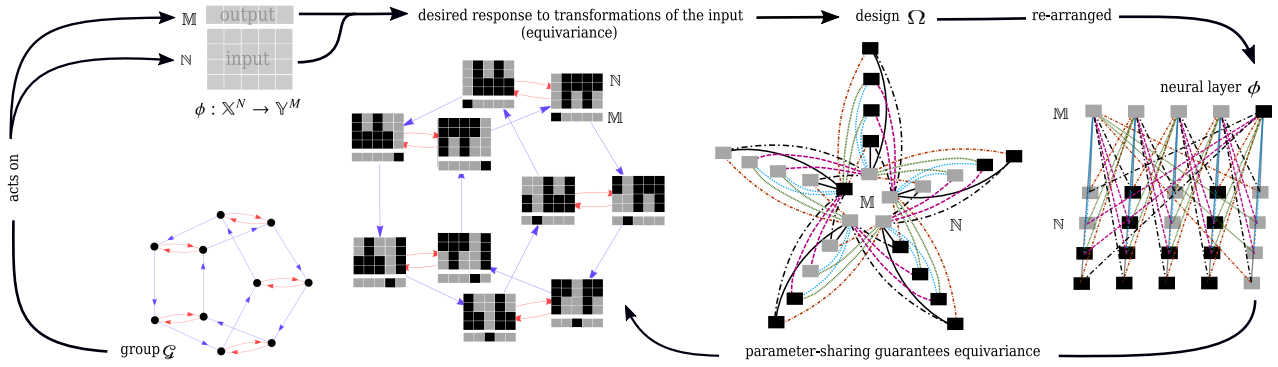


Figure 1. Summary: given a group action on input and output of a neural network layer, define a parameter-sharing for this layer that is equivariant to these actions.

(left) $\mathcal{G} = \mathcal{D}_5$ is a Dihedral group, acting on a 4×5 input image and an output vector of size 5. \mathbb{N} and \mathbb{M} denote the index set of input, and output variables respectively. Here \mathcal{G} is represented using its Cayley diagram.

(middle-left) \mathcal{G} -action for $g \in \mathcal{G}$ is shown for an example input. \mathcal{G} -action on the input is a combination of circular shifts (blue arrows) and vertical flips (red arrows) of the 2D image. \mathcal{G} acts on the output indices \mathbb{M} only through circular shift. A permutation group $\mathcal{G}_{\mathbb{N}, \mathbb{M}}$ encodes the simultaneous “action” of \mathcal{G} on input and output indices.

(middle-right) The structure Ω designed using our procedure, such that its symmetries $\mathcal{Aut}(\Omega)$ subsumes the permutation group $\mathcal{G}_{\mathbb{N}, \mathbb{M}}$.

(right) the same structure Ω unfolded to a bipartite form to better show the resulting parameter-sharing in the neural layer. The layer is equivariant to \mathcal{G} -action: shifting the input will shift the output of the resulting neural network function, while flipping the input does not change the output.

the output does not transform at all, the function is invariant to the action of \mathcal{G} . Therefore, invariance is a special equivariance. In this example, different translations correspond to the action of different members of \mathcal{G} .

The *novelty* of this work is its focus on the “model symmetry” as a gateway to equivariance. This gives us new theoretical guarantees for a “strict” notion of equivariance in neural networks. The core idea is simple: consider a colored bipartite graph Ω representing a neural network layer. Edges of the same color represent tied parameters. This *neural network layer as a function is equivariant to the actions of a given group \mathcal{G} (and nothing more) iff the action of \mathcal{G} is the symmetry group of Ω – i.e., there is a simple bijection between parameter symmetries and equivariences of the corresponding neural network.*

The problem then boils down to **designing** colored bipartite graphs with given symmetries, which constitutes a major part of this paper. Fig. 1 demonstrates this idea.¹

For the necessary background on group theory see the Appendix. In the following, Section 1 formalizes equivariance wrt discrete group action. Section 2 relates the model symmetries a neural layer to its equivariance. Section 3 then builds on this observation to introduce two procedures for parameter-sharing that achieves a desirable equivariance.

¹Throughout this paper, since we deal with finite sets, we use circular shift and circular convolution instead of shift and convolution. The two can be made identical with zero-padding of the input.

Here, we also see how group and graph convolution as well as deep-sets become special instances in our parameter-sharing procedure, which provides new insight and improved design in the case of group convolution. Where input and output of the layer have a one-to-one mapping, we see that the design problem reduces a well-known problem in combinatorics.

1. Group Action and Equivariance

Let $\mathbf{x} = [x_1, \dots, x_N] \in \mathbb{X}^N$ denote a set of variables and $\mathcal{G} = \{g\}$ be a finite group. The discrete action of \mathcal{G} on \mathbf{x} is in the form of permutation of indices in $\mathbb{N} = \{1, \dots, N\}$. This group is a subgroup of the *symmetric group* $\mathcal{S}_{\mathbb{N}}$; the group of all $N!$ permutations of N objects. We use $\vec{\mathbb{N}} = [1, \dots, N]$ to denote the ordered counterpart to \mathbb{N} and the \mathcal{G} -action on this vector $g\vec{\mathbb{N}} \doteq [g1, \dots, gN]$ is a simple permutation. Using $\mathbf{x}_{\vec{\mathbb{N}}}$ to denote \mathbf{x} , the discrete action of $g \in \mathcal{G}$ on $\mathbf{x} \in \mathbb{X}^N$ is given by $g\mathbf{x}_{\vec{\mathbb{N}}} \doteq \mathbf{x}_{g\vec{\mathbb{N}}}$.

\mathcal{G} -action on \mathbb{N} is a permutation group that is not necessarily isomorphic to \mathcal{G} itself. $\mathcal{G}_{\mathbb{N}} \leq \mathcal{G}$ captures the structure of \mathcal{G} when it acts on \mathbb{N} . We use $g_{\mathbb{N}}$ to denote the image of $g \in \mathcal{G}$ in $\mathcal{G}_{\mathbb{N}}$. \mathcal{G} -action is **faithful** iff two groups are isomorphic $\mathcal{G} \cong \mathcal{G}_{\mathbb{N}}$ – that is \mathcal{G} -action preserves its structure. In this case, each $g \in \mathcal{G}$ maps to a *distinct* permutation $g\vec{\mathbb{N}} \neq g'\vec{\mathbb{N}} \forall g, g' \in \mathcal{G}$. Given any \mathcal{G} -action on \mathbb{N} we can efficiently obtain $\mathcal{G}_{\mathbb{N}}$; see Appendix.

Example 1.1 (Cyclic Group) Consider the cyclic group $\mathcal{G} = \mathbb{Z}_6$ and define its action on $\mathbf{x} \in \mathbb{R}^3$ by defining it on the index set $\mathbb{N} = \{1, 2, 3\}$ as $gn \doteq g + n \pmod{3} \forall g \in \mathbb{Z}_6$. This action is not faithful. For example, the action of $g = 1$ and $g = 4$ result in the same permutations of variables in \mathbf{x} ; *i.e.*, single-step of circular shift to the right. With the above action, the resulting permutation group $\mathcal{G}_{\mathbb{N}}$ is isomorphic to $\mathbb{Z}_3 < \mathbb{Z}_6$.

Now consider the same group $\mathcal{G} = \mathbb{Z}_6$ with a different action on \mathbb{N} : $gn \doteq g - n \pmod{3} \forall g \in \mathbb{Z}_6$, where we replaced $(+)$ with $(-)$. Let $\tilde{\mathcal{G}}_{\mathbb{N}}$ be the resulting permutation group. Here again $\tilde{\mathcal{G}}_{\mathbb{N}} \cong \mathbb{Z}_3$. Although isomorphic, $\tilde{\mathcal{G}}_{\mathbb{N}} \neq \mathcal{G}_{\mathbb{N}}$, as they are different permutation groups of \mathbb{N} .

Consider the function $\phi : \mathbb{X}^N \rightarrow \mathbb{Y}^M$ and let $\mathcal{G}_{\mathbb{N}}$ and $\mathcal{G}_{\mathbb{M}}$ be the action of \mathcal{G} on input/output index sets \mathbb{N} and \mathbb{M} .

Definition 1.1 The joint permutation group $\mathcal{G}_{\mathbb{N}, \mathbb{M}}$ is a sub-direct product (or pairing) of $\mathcal{G}_{\mathbb{N}}$ and $\mathcal{G}_{\mathbb{M}}$

$$\mathcal{G}_{\mathbb{N}, \mathbb{M}} = \mathcal{G}_{\mathbb{N}} \odot \mathcal{G}_{\mathbb{M}} \doteq \{(g_{\mathbb{N}}, g_{\mathbb{M}}) \mid g \in \mathcal{G}\}.$$

We are now ready to define equivariance and invariance. $\phi(\cdot)$ is $\mathcal{G}_{\mathbb{N}, \mathbb{M}}$ -equivariant iff

$$g_{\mathbb{N}}\phi(\mathbf{x}) = \phi(g_{\mathbb{M}}\mathbf{x}) \quad \forall \mathbf{x} \in \mathbb{X}^N, (g_{\mathbb{N}}, g_{\mathbb{M}}) \in \mathcal{G}_{\mathbb{N}, \mathbb{M}} \quad (1)$$

Moreover, if $\mathcal{G}_{\mathbb{M}} = \{e\}$ is trivial, we have

$$g_{\mathbb{N}}\phi(\mathbf{x}) = \phi(\mathbf{x}) \quad \forall \mathbf{x} \in \mathbb{X}^N, g_{\mathbb{N}} \in \mathcal{G}_{\mathbb{N}}$$

and $\phi(\cdot)$ is $\mathcal{G}_{\mathbb{N}}$ -invariant.

$g_{\mathbb{N}}$ and $g_{\mathbb{M}}$ can also be represented using **permutation matrices** $\mathbf{G}_{\mathbb{N}} \in \{0, 1\}^{N \times N}$, and $\mathbf{G}_{\mathbb{M}} \in \{0, 1\}^{M \times M}$. Equivariance relation of (1) then becomes

$$\mathbf{G}_{\mathbb{M}}\phi(\mathbf{x}) = \phi(\mathbf{G}_{\mathbb{N}}\mathbf{x}) \quad \forall \mathbf{x} \in \mathbb{X}^N, (\mathbf{G}_{\mathbb{N}}, \mathbf{G}_{\mathbb{M}}) \in \mathcal{G}_{\mathbb{N}, \mathbb{M}} \quad (2)$$

The following observation shows that the subgroup relationship affects equivariance and invariance.

Observation 1.1 If the function $\phi : \mathbb{X}^N \rightarrow \mathbb{Y}^M$ is $\mathcal{G}_{\mathbb{N}, \mathbb{M}}$ -equivariant, then it is also $\mathcal{H}_{\mathbb{N}, \mathbb{M}}$ -equivariant for any permutation group $\mathcal{H}_{\mathbb{N}, \mathbb{M}} < \mathcal{G}_{\mathbb{N}, \mathbb{M}}$.

Example 1.2 (Reverse Convolution) Consider the cyclic group $\mathcal{G} = \mathbb{Z}_6$ and for $g \in \mathcal{G}$, define the action on $\mathbb{N} = \{1, 2, 3\}$ to be $gn \doteq g + n \pmod{3}$. Also let its action on $\mathbb{M} = \{1, \dots, 6\}$ be $gm \doteq g - n \pmod{6}$. In other words, \mathcal{G} -action on \mathbb{N} performs circular shift to the *right* and its action on \mathbb{M} shifts variables to the *left*. Examples of the permutation matrix representation for

two members of $\mathcal{G}_{\mathbb{N}}$ and $\mathcal{G}_{\mathbb{M}}$ are

$$2_{\mathbb{N}} = \begin{pmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 1 & 0 & 0 \end{pmatrix} \quad 2_{\mathbb{M}} = \begin{pmatrix} 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \end{pmatrix}$$

corresponding to right and left shift on vectors of different lengths. Now consider the function $\phi : \mathbb{R}^N \rightarrow \mathbb{R}^M$

$$\phi_{\mathbf{W}}(\mathbf{x}) = \mathbf{W}\mathbf{x} \quad \mathbf{W}^T = \begin{pmatrix} 0 & a & b & 0 & a & b \\ a & b & 0 & a & b & 0 \\ b & 0 & a & b & 0 & a \end{pmatrix} \quad \forall a, b \in \mathbb{R}$$

Using permutation matrices one could check the equivariance condition (2) for this function. We can show that ϕ is equivariant to $\mathcal{G}_{\mathbb{N}, \mathbb{M}}$. Consider $2 \in \mathbb{Z}_6$ and its images $2_{\mathbb{N}} \in \mathcal{G}_{\mathbb{N}}$ and $2_{\mathbb{M}} \in \mathcal{G}_{\mathbb{M}}$. L.h.s. of (2) is

$$2_{\mathbb{M}}\phi_{\mathbf{W}}(\mathbf{x}) = \begin{pmatrix} 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \end{pmatrix} \begin{pmatrix} 0 & a & b \\ a & b & 0 \\ b & 0 & a \\ 0 & a & b \\ a & b & 0 \\ b & 0 & a \end{pmatrix} \mathbf{x} = \begin{pmatrix} b & 0 & a \\ 0 & a & b \\ a & b & 0 \\ b & 0 & a \\ 0 & a & b \\ a & b & 0 \end{pmatrix} \mathbf{x}$$

which is equal to its r.h.s.

$$\phi_{\mathbf{W}}(2_{\mathbb{N}}\mathbf{x}) = \begin{pmatrix} 0 & a & b \\ a & b & 0 \\ b & 0 & a \\ a & b & 0 \\ b & 0 & a \end{pmatrix} \begin{pmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 1 & 0 & 0 \end{pmatrix} \mathbf{x} = \begin{pmatrix} b & 0 & a \\ 0 & a & b \\ a & b & 0 \\ 0 & a & b \\ a & b & 0 \end{pmatrix} \mathbf{x}$$

for any \mathbf{x} . One could verify this equality for all $g \in \mathbb{Z}_6$.

Now consider the group $\mathcal{H}_{\mathbb{N}, \mathbb{M}} < \mathcal{G}_{\mathbb{N}, \mathbb{M}}$, where $\mathcal{H}_{\mathbb{N}} = \mathcal{G}_{\mathbb{N}}$ and members of $\mathcal{H}_{\mathbb{M}} = \{0, 2, 4\}$, perform left circular shift of length 0, 2 and 4. It is easy to see that $\mathcal{H}_{\mathbb{N}, \mathbb{M}} \cong \mathbb{Z}_3$. Moreover since $\mathcal{H}_{\mathbb{N}, \mathbb{M}} < \mathcal{G}_{\mathbb{N}, \mathbb{M}}$, $\phi(\cdot)$ above is $\mathcal{H}_{\mathbb{N}, \mathbb{M}}$ -equivariant as well. However, one prefers to characterize the equivariance properties of ϕ using $\mathcal{G}_{\mathbb{N}, \mathbb{M}}$ rather than $\mathcal{H}_{\mathbb{N}, \mathbb{M}}$.

The observation above suggests that $\mathcal{G}_{\mathbb{N}, \mathbb{M}}$ -equivariance is not restrictive enough. As an extreme case, a constant function $\phi(\mathbf{x}) = \mathbf{1}$ is equivariant to any permutation group $\mathcal{G}_{\mathbb{N}, \mathbb{M}} \leq \mathcal{S}_{\mathbb{N}} \times \mathcal{S}_{\mathbb{M}}$. In this case equivariance of ϕ with respect to a particular $\mathcal{G}_{\mathbb{N}, \mathbb{M}}$ is not very informative to us. To remedy this, we define a more strict notion of equivariance.

Definition 1.2 we say a function $\phi : \mathbb{X}^N \rightarrow \mathbb{Y}^M$ is **uniquely \mathcal{G} -equivariant** iff it is \mathcal{G} -equivariant and it is “not” \mathcal{H} -equivariant for any $\mathcal{H} > \mathcal{G}$.

2. Symmetry Groups of a Network

Given a group \mathcal{G} , and its discrete action through $\mathcal{G}_{\mathbb{N}, \mathbb{M}}$, we are interested in defining parameter-sharing schemes for a parametric class of functions that guarantees their unique $\mathcal{G}_{\mathbb{N}, \mathbb{M}}$ -equivariance. We start by looking at a single neural layer and relate its unique $\mathcal{G}_{\mathbb{N}, \mathbb{M}}$ -equivariance to the symmetries of a colored multi-edged bipartite graph that de-

finer parameter-sharing. We then show that the idea extends to multiple-layers.

Definition 2.1 A colored multi-edged bipartite graph $\Omega = (\mathbb{N}, \mathbb{M}, \alpha)$ is a triple, where \mathbb{N} and \mathbb{M} are its two sets of nodes, and $\alpha : \mathbb{N} \times \mathbb{M} \rightarrow 2^{\{1, \dots, C\}}$ is the edge function that assigns multiple edge-colors from the set $\{1, \dots, C\}$ to each edge. Non-existing edges receive no color.

We are interested in the symmetries of this structure. The set of permutations $(\pi_{\mathbb{N}}, \pi_{\mathbb{M}}) \in \mathcal{S}_{\mathbb{N}} \times \mathcal{S}_{\mathbb{M}}$ of nodes (within each part of the bipartite graph) that preserve all edge-colors define the **Automorphism Group** $\mathcal{Aut}(\Omega) \leq \mathcal{S}_{\mathbb{N}} \times \mathcal{S}_{\mathbb{M}}$ – that is $\forall (n, m) \in \mathbb{N} \times \mathbb{M}$

$$(\pi_{\mathbb{N}}, \pi_{\mathbb{M}}) \in \mathcal{Aut}(\Omega) \Leftrightarrow \alpha(n, m) = \alpha((\pi_{\mathbb{N}}n, \pi_{\mathbb{M}}m)) \quad (3)$$

Alternatively, to facilitate the notation, we define the same structure (colored multi-edged bipartite graph) as a **set of binary relations** between \mathbb{N} and \mathbb{M} – that is $\Omega = (\mathbb{N}, \mathbb{M}, \{\Delta_c\}_{1 \leq c \leq C})$ where each relation is associated with one color $\Delta_c = \{(n, m) \mid c \in \alpha(n, m) \forall (n, m) \in \mathbb{N} \times \mathbb{M}\}$. This definition of structure, gives an alternative expression for $\mathcal{Aut}(\Omega)$

$$\begin{aligned} (\pi_{\mathbb{N}}, \pi_{\mathbb{M}}) \in \mathcal{Aut}(\Omega) &\Leftrightarrow \\ \left((n, m) \in \Delta_c \Leftrightarrow (\pi_{\mathbb{N}}n, \pi_{\mathbb{M}}m) \in \Delta_c \right) &\quad \forall c, n, m \end{aligned} \quad (4)$$

The significance of this structure is in that, it defines a parameter-sharing scheme in a neural layer, where the same edge-colors correspond to the same parameters. Consider the function $\phi \doteq [\phi_1, \dots, \phi_M] : \mathbb{R}^N \rightarrow \mathbb{R}^M$

$$\phi_m(\mathbf{x}; \mathbf{w}, \Omega) \doteq \sigma \left(\sum_n \sum_{c \in \alpha(n, m)} w_c x_n \right) \quad \forall m \quad (5)$$

where $\sigma : \mathbb{R} \rightarrow \mathbb{R}$ is a strictly *monotonic* nonlinearity and $\mathbf{w} = [w_1, \dots, w_c, \dots, w_C]$ is the parameter-vector for this layer.


The following key theorem relates the equivariances of $\phi(\cdot; \mathbf{w}, \Omega)$ to the symmetries of Ω .

Theorem 2.1 For any $\mathbf{w} \in \mathbb{R}^C$ s.t., $w_c \neq w_{c'} \forall c, c'$, the function $\phi(\cdot; \mathbf{w}, \Omega)$ is uniquely $\mathcal{Aut}(\Omega)$ -equivariant.

Corollary 2.2 For any $\mathcal{H}_{\mathbb{N}, \mathbb{M}} \leq \mathcal{Aut}(\Omega)$, the function $\phi(\cdot; \mathbf{w}, \Omega)$ is $\mathcal{H}_{\mathbb{N}, \mathbb{M}}$ -equivariant.

The implication is that to achieve unique equivariance for a given group-action, we need to define the parameter-sharing using the structure Ω with symmetry group $\mathcal{G}_{\mathbb{N}, \mathbb{M}}$.

Example 2.1 (Reverse Convolution) Revisiting Example 1.2 we can show that the condition of Theorem 2.1 holds. In this case $\sigma(\mathbf{x}) = \mathbf{x}$ and the parameter-sharing of the matrix \mathbf{W} is visualized below, where we used two different line styles for $a, b \in \mathbb{R}$.



In this figure, the circular shift of variables at the output and input level to the left and right respectively, does not change the edge-colors. For example in both cases node ①'s connection to nodes ③, ⑥ using dashed-lines is preserved.

Six repetitions of this action produces different permutations corresponding to six members of $\mathcal{G}_{\mathbb{N}, \mathbb{M}}$. Therefore $\mathcal{G}_{\mathbb{N}, \mathbb{M}} \leq \mathcal{Aut}(\Omega)$ and according to Corollary 2.2, $\phi(\cdot)$ is $\mathcal{G}_{\mathbb{N}, \mathbb{M}}$ equivariant. Moreover, using Theorem 3.3 of the next section, we can prove that these six permutations are the “only” edge-color preserving ones for this structure, resulting in *unique equivariance*.


Matrix Form. To write (5) in a matrix form, if there are multiple edges between two nodes n, m , we need to merge them. In general, by assigning on distinct color to any set in the range of $\alpha : \mathbb{N} \times \mathbb{M} \rightarrow 2^{\{1, \dots, C\}}$ we can w.l.o.g. reduce multiple edges to a single edge. In other words we can rewrite ϕ using $\mathbf{W} \in \mathbb{R}^{M \times N}$

$$\phi(\mathbf{x}; \mathbf{w}; \Omega) = \sigma(\mathbf{W}\mathbf{x}) \quad W_{m, n} = \sum_{c \in \alpha(n, m)} w_c \quad (6)$$

Using this notation, and due to strict monotonicity of the nonlinearity $\sigma(\cdot)$, Theorem 2.1 simply states that for all $(g_{\mathbb{N}}, g_{\mathbb{M}}) \in \mathcal{Aut}(\Omega)$, $\mathbf{x} \in \mathbb{R}^N$ and \mathbf{W} given by (6)

$$\mathbf{G}_{\mathbb{M}} \mathbf{W} \mathbf{x} = \mathbf{W} \mathbf{G}_{\mathbb{N}} \mathbf{x}. \quad (7)$$

Example 2.2 (Permutation-Equivariant Layer) Consider all permutations of indices \mathbb{N} and $\mathbb{M} = \mathbb{N}$.



We want to define a neural layer such that all permutations of the input $g_{\mathbb{N}} \in \mathcal{G}_{\mathbb{N}} = \mathcal{S}_{\mathbb{N}}$ result in the same

permutation of the output $g_M = g_N$. Consider the following colored bipartite graph, for a special case where $N = M = 4$. It is easy to show that color-preserving permutations of this structure are $\mathcal{Aut}(\Omega) = \mathcal{S}_N \circ \mathcal{S}_M = \{(g, g) \mid g \in \mathcal{S}_N\} \cong \mathcal{S}_N$: On one hand, for $(\pi_N, \pi_M) \in \mathcal{S}_N \times \mathcal{S}_M$, having $\pi_N = \pi_M$ clearly preserves the colors. On the other hand, if $\pi_N \neq \pi_M$, there exists $u \in \mathbb{N}$ (also in \mathbb{M}) such that $\pi_N u \neq \pi_M u$. Therefore (π_N, π_M) does not preserve the relation $\Delta = \{(n, n) \mid n \in \mathbb{N}\}$ corresponding to dashed edges, and therefore $(\pi_N, \pi_M) \notin \mathcal{Aut}(\Omega)$. This proves $\mathcal{Aut}(\Omega) = \mathcal{S}_N \circ \mathcal{S}_M$. The function (5) for this Ω is

$$\phi(\mathbf{x}; \mathbf{w} = [w_1, w_2], \Omega) = \sigma(w_1 \mathbf{I} \mathbf{x} + w_2 \mathbf{1} \mathbf{1}^T \mathbf{x}).$$

Ravanbakhsh et al. (2016); Zaheer et al. (2017) derive the same permutation equivariant layer, by proving the commutativity in (7), while here it follows from Corollary 2.2.

Multiple Layers. For deep networks, the equivariance of the composition $\phi_2 \circ \phi_1$ to \mathcal{G} -action follows from that of individual layer $\phi_1 : \mathbb{X}^N \rightarrow \mathbb{Y}^M$ and $\phi_2 : \mathbb{Y}^M \rightarrow \mathbb{Z}^O$. Assuming ϕ_1 is $\mathcal{G}_{N,M}$ -equivariant and ϕ_2 is $\mathcal{G}_{M,O}$ -equivariant, where \mathcal{G} -action on \mathbb{M} is shared between the two layers, it follows that $\phi_2 \circ \phi_1$ is $\mathcal{G}_{N,O}$ -equivariant, where $\mathcal{G}_{N,O} = \mathcal{G}_N \circ \mathcal{G}_O$. This is because $\forall g \in \mathcal{G}$ and $\mathbf{x} \in \mathbb{X}^N$

$$\phi_2(\phi_1(g_N \mathbf{x})) = \phi_2(g_M \phi_1(\mathbf{x})) = g_O \phi_2(\phi_1(\mathbf{x})). \quad (8)$$

3. Structure Design

Consider the definition of neural layer (5) that employs parameter-sharing according to Ω . Given \mathcal{G} -action on \mathbb{N} and \mathbb{M} , we are interested in designing structures Ω such that $\mathcal{Aut}(\Omega) = \mathcal{G}_{N,M}$. According to the Theorem 2.1, it then follows that ϕ is uniquely $\mathcal{G}_{N,M}$ -equivariant. Here, we give the sufficient conditions and the design recipe to achieve this.

For this we briefly review some group properties that are used in later developments.

transitivity We say that \mathcal{G} -action on \mathbb{N} is transitive iff $\forall n_1, n_2 \in \mathbb{N}$, there exists at least one action $g \in \mathcal{G}$ (or $g_N \in \mathcal{G}_N$) such that $gn_1 = n_2$.

regularity The group action is free or **semi-regular** iff $\forall n_1, n_2 \in \mathbb{N}$, there is at most one $g \in \mathcal{G}$ such at $gn_1 = n_2$, and the action is *regular* iff it is both transitive and free – i.e., for any pair $n_1, n_2 \in \mathbb{N}$, there is uniquely one $g \in \mathcal{G}$ such that $gn_1 = n_2$. Any free action is also faithful. We use a similar terminology for \mathcal{G}_N . That is we call \mathcal{G}_N semi-regular iff $\forall n_1, n_2 \in \mathbb{N}$ at

most one $g_N \in \mathcal{G}_N$ moves n_1 to n_2 and \mathcal{G}_N is regular if this number is exactly one.

orbit The orbit of $n \in \mathbb{N}$ is all the members to which it can be moved, $\mathcal{G}n = \{gn \mid g \in \mathcal{G}\}$. The orbits of $n \in \mathbb{N}$ form an equivalence relation² This equivalence relation *partitions* \mathbb{N} into orbits $\mathbb{N} = \bigcup_{1 \leq p \leq P} \mathcal{G}n_p$, where n_p is an arbitrary **representative** of the partition $\mathcal{G}n_p \subseteq \mathbb{N}$. Note that the \mathcal{G} -action on \mathbb{N} is always transitive on its orbits – that is for any $n, n' \in \mathcal{G}n_p$, there is at least one $g \in \mathcal{G}$ such that $n = gn'$. Therefore, for a semi-regular \mathcal{G} -action, the action of \mathcal{G} on the orbits $\mathcal{G}n_p \forall 1 \leq p \leq P$ is regular.

Example 3.1 (Mirror Symmetry) Consider $\mathcal{G} = \mathbb{Z}_2 = \{e = 0, 1\}$ ($1 + 1 = 0$) acting on \mathbb{N} , where the only non-trivial action is defined as flipping the input: $1_N[1, \dots, N] = [N, N - 1, \dots, 1]$.

\mathcal{G} is faithful in its action on \mathbb{N} , however \mathcal{G}_N is not transitive – e.g., N cannot be moved to $N - 1$. If N is even, then \mathcal{G} -action is semi-regular. This is because otherwise the element in the middle $n = \lceil \frac{N}{2} \rceil$ is moved to itself by two different actions $e, 1 \in \mathcal{G}$. Furthermore, if N is even, \mathcal{G} -action has $\frac{N}{2}$ orbits and \mathcal{G}_2 acts on these orbits regularly. If N is odd, \mathcal{G} -action has $\lceil \frac{N}{2} \rceil$ orbits. However, its action on the orbit of the middle element $\mathcal{G} \lceil \frac{N}{2} \rceil$ is not regular.

In the following, Section 3.1 proposes a procedure for parameter-sharing in a fully connected layer. Although simple, this design is dense and does not guarantee “unique” of equivariance. Section 3.2 proposes an alternative design with sparse connections that in some settings ensures unique equivariance. Section 3.3 investigates the effect of having multiple input and output channels in the neural layer and Section 3.4 studies a special case of $\mathcal{G}_N = \mathcal{G}_M$, where input and output indices have a one-to-one mapping.

3.1. Dense Design

Consider a complete bipartite graph with \mathbb{N} and \mathbb{M} as its two parts and edges $(n, m) \in \mathbb{N} \times \mathbb{M}$. The action of $\mathcal{G}_{N,M}$ partitions the edges into orbits $\{\mathcal{G}_{N,M}(n_p, m_q)\}_{n_p, m_q}$, where (n_p, m_q) is a representative edge from an orbit. Painting each orbit with a different color gives

$$\Omega = (\mathbb{N}, \mathbb{M}, \{\Delta_{p,q} = \mathcal{G}_{N,M}(n_p, m_q)\}). \quad (9)$$

Therefore two edges (n, m) and (n', m') have the same color iff an action in $\mathcal{G}_{N,M}$ moves one edge to the other.

² $n \sim n' \Leftrightarrow \exists g \text{ s.t., } n = gn' \Leftrightarrow n \in \mathcal{G}n' \Leftrightarrow n' \in \mathcal{G}n$.

Proposition 3.1 $\mathcal{G}_{\mathbb{N},\mathbb{M}} \leq \Omega$ for Ω of (9).

Corollary 3.2 $\phi(\cdot; \mathbf{w}, \Omega)$, for structure (9), is equivariant to $\mathcal{G}_{\mathbb{N},\mathbb{M}}$.


Example 3.2 (Nested Subsets and Wreath Product)

The permutation-equivariant layer that we saw in Example 2.2 is useful for defining neural layers for set structure. If our data-structure is in the form of nested subsets, then we require equivariance to permutation of variables within each set as well as permutation of subsets. Here, we show how to use our dense design for this purpose.

We use a special indexing for the input to better identify the exchangeability of variables. We assume D subsets, each of which has d variables $\mathbf{x} = [x_{1,1}, \dots, x_{1,d}, x_{2,1}, \dots, x_{D,d}]$.

The group of our interest is the **wreath product** $\mathcal{S}_d \wr \mathcal{S}_D$. This type of group product can be used to build hierarchical and nested structures with different type of symmetries at each level. Nesting subsets corresponds to the most basic form of such hierarchical constructions. We use (n, n') to index input variables and (m, m') for output variables.

The following figure shows the resulting parameter-sharing for an example with $D = 2, d = 3$.



How did we arrive at this structure Ω ? Recall Our objective is to define parameter-sharing so that $\phi_{\mathbf{w}} : \mathbb{R}^{dD} \rightarrow \mathbb{R}^{dD}$ is equivariant to the action of $\mathcal{G} = \mathcal{S}_d \wr \mathcal{S}_D$ – i.e., permutations within sets at two levels. This group-action identifies three partitions of edges (seen in the figure): I) $((n, n'), (n, n')) \forall n, n'$ connects each variable to its counterpart (dashed orange); II) $((n, n'), (n, m')) \forall n, n' \neq m'$ connects each variable to other variables within the same subset; III) $((n, n'), (m, m')) \forall n \neq m$ is the set of edges from one subset to another. According to the Corollary 3.2 this parameter-sharing guarantees equivariance.

This fully-connected design is useful when the group $\mathcal{G}_{\mathbb{N},\mathbb{M}}$ is large; for example when dealing with $\mathcal{S}_{\mathbb{N}}$. However, for

smaller groups it could be very inefficient in practice, as sometimes we can achieve equivariance through a sparse structure Ω . As an example, consider the 2D circular convolution layer. It is easy to show that according to this design, the convolution filter will be the same size as the input image. While this achieves the desirable equivariance, it is inefficient and does not generalize as well as a convolution layer with small filters. Moreover, the dense design does not guarantee “unique” equivariance. We next show under some conditions on $\mathcal{G}_{\mathbb{N},\mathbb{M}}$ the sparse design can produce this stronger guarantee.

3.2. Sparse Design

Our sparse construction uses orbits and symmetric generating sets:

- Let us denote the orbits of \mathcal{G} -action on \mathbb{M} and \mathbb{N} by $\{\mathcal{G}n_p \mid 1 \leq p \leq P\}$ and $\{\mathcal{G}m_q \mid 1 \leq q \leq Q\}$ respectively, where P and Q are the total number of orbits and n_p, m_q are (arbitrary) representative members of orbits $\mathcal{G}n_p, \mathcal{G}m_q$ respectively. Note that in contrast to previous section, here we are considering the orbit of variables rather than the edges.
- The set $\mathcal{A} \subseteq \mathcal{G}$ is called the generating set of \mathcal{G} ($\langle \mathcal{A} \rangle = \mathcal{G}$), iff every member of \mathcal{G} can be expressed as a combination of members of \mathcal{A} . If the generating set is closed under inverse $a \in \mathcal{A} \Rightarrow a^{-1} \in \mathcal{A}$ we call it a **symmetric generating set**.

Define the structure Ω as

$$\Omega = (\mathbb{N}, \mathbb{M}, \{\Delta_{p,q,a}\}_{1 \leq p \leq P, 1 \leq q \leq Q, a \in \mathcal{A}})$$

$$\Delta_{p,q,a} = \{(g_{\mathbb{N}} a n_p, g_{\mathbb{M}} m_q) \mid (g_{\mathbb{N}}, g_{\mathbb{M}}) \in \mathcal{G}_{\mathbb{N},\mathbb{M}}\}. \quad (10)$$

In words, we have one color per each combination of orbits (p, q) and members of the generating set $a \in \mathcal{A}$. The following theorem relates the symmetry group of this structure to \mathcal{G} .

Theorem 3.3 $\mathcal{G}_{\mathbb{N},\mathbb{M}} \leq \text{Aut}(\Omega)$ for Ω of (10). Moreover if $\mathcal{G}_{\mathbb{N}}$ and $\mathcal{G}_{\mathbb{M}}$ are both semi-regular, then $\mathcal{G}_{\mathbb{N},\mathbb{M}} = \text{Aut}(\Omega)$.

Note that this result holds for any choice of a symmetric generating set \mathcal{A} in defining Ω . Therefore, in designing sparse layers, one seeks a minimal \mathcal{A} .

Corollary 3.4 The function $\phi(\cdot, \mathbf{w}, \Omega)$, using the structure (10) is $\mathcal{G}_{\mathbb{N},\mathbb{M}}$ -equivariant. If $\mathcal{G}_{\mathbb{N}}$ and $\mathcal{G}_{\mathbb{M}}$ are semi-regular, this function is “uniquely” $\mathcal{G}_{\mathbb{N},\mathbb{M}}$ -equivariant.

Now, assuming \mathcal{G} -action is semi-regular on both \mathbb{N} and \mathbb{M} , using (arbitrarily chosen) representatives $\{n_p\}_{1 \leq p \leq P}$ and $\{m_q\}_{1 \leq q \leq Q}$ for orbits in \mathbb{N} and \mathbb{M} , we can rewrite the expression (5) of the structured neural layer for the structure above. Here, components of $\phi = [\phi_1, \dots, \phi_M]$ are enumerated for $1 \leq q \leq Q, g_M \in \mathcal{G}_M$:

$$\phi_{g_M m_q}(\mathbf{x}; \mathbf{w}) = \sigma \left(\sum_{1 \leq p \leq P} \sum_{a \in \mathcal{A}} w_{q,p,a} x_{g_N a n_p} \right) \quad (11)$$

where $\mathbf{w} \in \mathbb{R}^{P \times Q \times |\mathcal{A}|}$ is the set of unique parameters, and each element $\phi_{g_M m_q}$ depends on subset of parameters $\{w_{q,p,a}\}_{p,a}$ identified by q and a subset of inputs $\{x_{a, g_N n_p}\}_{p,a}$ identified by g_N .

Example 3.3 (Dihedral Group of Fig. 1) In the example of Fig. 1, the number of orbits of \mathcal{G} -action on \mathbb{N} is $P = 2$ and for \mathbb{M} this is $Q = 1$. The symmetric generating set is the generating set that is used in the Cayley diagram, with the addition of inverse shift (inverse of the blue arrow). We then used (10) to build the structure of Fig. 1 (right).

Example 3.4 (Reverse Convolution) The parameter-sharing structure of reverse convolution in Examples 1.2 and 2.1 is produced using our sparse design. In these examples, both \mathcal{G}_N and \mathcal{G}_M are regular. Therefore the proposed parameter-sharing provides *unique* equivariance.

3.3. Multiple Channels

In this section, we extend our results to multiple input and output channels. Up to this point, we considered a neural network layer $\phi : \mathbb{R}^N \rightarrow \mathbb{R}^M$. Here, we want to see how to achieve $\mathcal{G}_{N,M}$ -equivariance for $\phi : \mathbb{R}^{N \times K} \rightarrow \mathbb{R}^{M \times K'}$, where K and K' are the number of input and output channels.

First, we extend the action of \mathcal{G} on \mathbb{N} and \mathbb{M} to $\mathbb{N}^K = [\underbrace{\mathbb{N}, \dots, \mathbb{N}}_{K \text{ times}}]$ as well as $\mathbb{M}^{K'}$, to accommodate multiple chan-

nels. For this, simply repeat the \mathcal{G} -action on each component. \mathcal{G} -action on multiple input channels is equivalent to sub-direct product $\underbrace{\mathcal{G}_N \odot \dots \odot \mathcal{G}_N}_{K \text{ times}} \cong \mathcal{G}_N$. The same applies

to \mathcal{G}_M .

This repetition, multiplies the orbits of \mathcal{G}_N , one for each channel, so that instead of having P and Q orbits on the input \mathbb{N} and output \mathbb{M} sets, we have $K \times P$ and $K' \times Q$ orbits on the input \mathbb{N}^K and output $\mathbb{M}^{K'}$. This increases the number of parameters by a factor of $K \times K'$.


The important implication is that, **orbits and multiple channels are treated identically by both dense and sparse designs.**

Example 3.5 (Group Convolution) The idea of group-convolution is studied by Cohen & Welling (2016a); see also (Olah, 2014). The following claim relates the function of this type of layer to our sparse design.

Claim 3.5 *Under the following conditions the neural layer (5) using our sparse design (10) performs group convolution: I) there is a bijection between the output and \mathcal{G} (i.e., $\mathbb{M} = \mathcal{G}$) and; II) \mathcal{G}_N is transitive.*


This also identifies the limitations of group-convolution even in the setting where $\mathbb{M} = \mathcal{G}$: When \mathcal{G}_N is semi-regular and not transitive ($P > 1$), group convolution is not guaranteed to be uniquely equivariant while the sparse parameter-sharing of (10) provides this guarantee.

For demonstration consider the following example in **equivariance to mirror symmetry**.



This figure shows the bipartite structure for $\mathcal{G} = \mathbb{Z}_2 = \{0, 1\}$ and $\mathcal{A} = \{1\}$. \mathcal{G} -action is horizontal flip of the input and the output. On the *right*, $\mathbb{M} = \mathcal{G}$ while on the *left* \mathcal{G}_M -action has two orbits. Orbits are identified by line-style and color of the circles. In a neural layer with this parameter-sharing, when we flip the input variables (around the mirror line) the output is also flipped.

The representatives in each orbit on \mathbb{N} and \mathbb{M} is identified with a star. Note that each combination of orbits p and q has a parameter of its own, identified with different edge-styles. While this construction guarantees “unique” \mathcal{G} -equivariance, if instead we use the same parameters across orbits (as suggested by the original group convolution) we get the parameter-sharing of the figure below *middle*.



In this case, the resulting neural layer has the desired equivariance (*right*). However, it is equivariant to the action of a larger group $\mathcal{G}_{\mathbb{N}, \mathbb{M}} \cong \mathbb{Z}_2 \times \mathbb{Z}_2 > \mathbb{Z}_2$, in which 1 in the second \mathbb{Z}_2 group exchanges variables across the orbits on \mathbb{N} (*left* in figure above).

3.4. $\mathcal{G}_{\mathbb{N}} = \mathcal{G}_{\mathbb{M}}$


In semi-supervised and un-supervised applications, we often need to produce a single output y_n for each input $x_n \forall n \in \mathbb{N}$ – that is $\mathbb{N} = \mathbb{M}$. We can ensure this by having a relation $\Delta_{c^*} = \{(n, n) \mid n \in \mathbb{N}\}$ in Ω that guarantees any $(\pi_{\mathbb{N}}, \pi_{\mathbb{M}}) \in \mathcal{Aut}(\Omega)$ applies the same permutation to \mathbb{N} and $\mathbb{M} = \mathbb{N}$ – i.e., $\pi_{\mathbb{N}} = \pi_{\mathbb{M}}$. The resulting structure $\Omega = (\mathbb{N}, \mathbb{N}, \{\Delta_c\}_{1 \leq c \leq C} \cup \{\Delta_{c^*}\})$ can be also interpreted as a colored multi-edged *directed* graph (digraph). This is because we can collapse the two parts by identifying $n \in \mathbb{N}$ with $n \in \mathbb{M}$.

Therefore, the symmetry-group of the original bipartite structure, is isomorphic to symmetry group of a colored multi-edged *digraph* on \mathbb{N} . Achieving unique \mathcal{G} -equivariance then reduces to answering the following question: *when could we express a permutation group $\mathcal{G} \leq \mathcal{S}_{\mathbb{N}}$ as the symmetry group $\mathcal{Aut}(\Omega)$ of a colored multi-edged digraph with N nodes?*

This problem is well-studied under the class of **concrete representation problems** (Babai, 1994). Permutation groups \mathcal{G} that can be expressed in this way are called **2-closed groups** (Wielandt, 1969). The recipe for achieving $\mathcal{G}_{\mathbb{N}} \leq \mathcal{Aut}(\Omega)$ is similar to our dense construction of Section 3.1³ The 2-closure $\bar{\mathcal{G}}_{\mathbb{N}}$ of a group $\mathcal{G}_{\mathbb{N}}$ is then, the greatest permutation group $\bar{\mathcal{G}}_{\mathbb{N}} \leq \mathcal{S}_{\mathbb{N}}$ with the same orbit on $\mathbb{N} \times \mathbb{N}$ as $\mathcal{G}_{\mathbb{N}}$. It is known that for example semi-regular permutation groups are 2-closed $\bar{\mathcal{G}}_{\mathbb{N}} = \mathcal{G}_{\mathbb{N}}$. This result also follows a corollary of our Theorem 3.3 for sparse design of (10).

Example 3.6 (Equivariance to $\times 90^\circ$ Rotations)

Figure below compares the digraph representation of Ω produced using (*left*) our sparse design, and (*right*) our dense design.



³In a fully connected digraph, the edges that belong to the same orbit by \mathcal{G} -action on $\mathbb{N} \times \mathbb{N}$, receive the same color.

Multiples of $\pm 90^\circ$ rotation is produced as the action of cyclic group \mathbb{Z}_4 on eight input output variables – that is $\mathbb{N} = \mathbb{M} = \{1, \dots, 8\}$. \mathbb{Z}_4 -action is semi-regular with two orbits; these orbits the two inner and outer set of four nodes. The representatives of each orbit in our sparse design is indicated using filled circles. The generating set consists of $\mathcal{A} = \{1, 3\}$, rotation by 90° and its inverse, rotation by 270° . Each edge in each of these figures, has a corresponding edge in the opposite direction, within a different relation. To avoid over-crowding the figure, we have dropped this edge from the drawing above, unless both edges belong to the same relation.

Example 3.7 (Graph Convolution) Consider the setting where we use the (normalized) adjacency matrix $\mathbf{B} \in \{0, 1\}^{N \times N}$ (or Laplacian) of a graph Λ , to identify parameter-sharing in a neural network layer. For a single input/output channel, this is often in the form of $\mathbf{A}\mathbf{x}$, where $\mathbf{x} \in \mathbb{R}^N$ and $\mathbf{A} = w_1\mathbf{B} + w_2\mathbf{I}$ has different parameters for diagonal and off-diagonal values (e.g., Kipf & Welling, 2016; Bruna et al., 2013; Henaff et al., 2015); for multiple channels see Section 3.3. The following corollary of Theorem 2.1 identifies the equivariance of $\mathbf{A}\mathbf{x}$.

Corollary 3.6 *Given the digraph Λ and its binary adjacency matrix $\mathbf{B} \in \{0, 1\}^{N \times N}$, then $(w_1\mathbf{B} + w_2\mathbf{I})\mathbf{x}$ is uniquely equivariant to the symmetry-group of Λ .*

Since two graphs on N nodes can have identical symmetries, one implication of this corollary is that graph-convolution has identical equivariances for graphs with the same symmetry groups.

4. Conclusion

This work is a step towards designing neural network layers with a given equivariance and invariance properties. Our approach was to relate the equivariance properties of the neural layer to the symmetries of the parameter-matrix.

We then proposed two parameter-sharing scheme that achieves equivariance wrt any discrete group-action. Moreover under some conditions, we guarantee sensitivity wrt other group actions. This is important because even a trivial constant function is invariant to all transformations. It is therefore essential to be able to draw the line between equivariance/invariance and sensitivity in a function. To our knowledge, our work presents the first results of its kind on guarantees regarding both variance and equivariance with respect to group actions.

Acknowledgment

This research is supported in part by DOE grant DESC0011114 and NSF grant IIS1563887.

References

- Agrawal, Pulkit, Carreira, Joao, and Malik, Jitendra. Learning to see by moving. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 37–45, 2015.
- Anselmi, Fabio, Leibo, Joel Z, Rosasco, Lorenzo, Mutch, Jim, Tacchetti, Andrea, and Poggio, Tomaso. Unsupervised learning of invariant representations in hierarchical architectures. *arXiv preprint arXiv:1311.4158*, 2013.
- Babai, Laszlo. Automorphism groups, isomorphism, reconstruction (chapter 27 of the handbook of combinatorics). *University of Chicago*, 1994.
- Bartók, Gábor, Szepesvári, Csaba, and Zilles, Sandra. Models of active learning in group-structured state spaces. *Information and Computation*, 208(4):364–384, 2010.
- Bruna, Joan and Mallat, Stéphane. Invariant scattering convolution networks. *IEEE transactions on pattern analysis and machine intelligence*, 35(8):1872–1886, 2013.
- Bruna, Joan, Zaremba, Wojciech, Szlam, Arthur, and LeCun, Yann. Spectral networks and locally connected networks on graphs. *arXiv preprint arXiv:1312.6203*, 2013.
- Bui, Hung Hai, Huynh, Tuyen N, and Riedel, Sebastian. Automorphism groups of graphical models and lifted variational inference. *arXiv preprint arXiv:1207.4814*, 2012.
- Cohen, Taco S and Welling, Max. Group equivariant convolutional networks. *arXiv preprint arXiv:1602.07576*, 2016a.
- Cohen, Taco S and Welling, Max. Steerable cnns. *arXiv preprint arXiv:1612.08498*, 2016b.
- Dieleman, Sander, Willett, Kyle W, and Dambre, Joni. Rotation-invariant convolutional neural networks for galaxy morphology prediction. *Monthly notices of the royal astronomical society*, 450(2):1441–1459, 2015.
- Dieleman, Sander, De Fauw, Jeffrey, and Kavukcuoglu, Koray. Exploiting cyclic symmetry in convolutional neural networks. *arXiv preprint arXiv:1602.02660*, 2016.
- Freeman, William T, Adelson, Edward H, et al. The design and use of steerable filters. *IEEE Transactions on Pattern analysis and machine intelligence*, 13(9):891–906, 1991.
- Gens, Robert and Domingos, Pedro M. Deep symmetry networks. In *Advances in neural information processing systems*, pp. 2537–2545, 2014.
- Hel-Or, Yacov and Teo, Patrick C. A common framework for steerability, motion estimation, and invariant feature detection. In *Circuits and Systems, 1998. ISCAS’98. Proceedings of the 1998 IEEE International Symposium on*, volume 5, pp. 337–340. IEEE, 1998.
- Henaff, Mikael, Bruna, Joan, and LeCun, Yann. Deep convolutional networks on graph-structured data. *arXiv preprint arXiv:1506.05163*, 2015.
- Hinton, Geoffrey E, Krizhevsky, Alex, and Wang, Sida D. Transforming auto-encoders. In *International Conference on Artificial Neural Networks*, pp. 44–51. Springer, 2011.
- Jaderberg, Max, Simonyan, Karen, Zisserman, Andrew, et al. Spatial transformer networks. In *Advances in Neural Information Processing Systems*, pp. 2017–2025, 2015.
- Jayaraman, Dinesh and Grauman, Kristen. Learning image representations equivariant to ego-motion. In *Proc. ICCV*, 2015.
- Kersting, Kristian, Ahmadi, Babak, and Natarajan, Sri-raam. Counting belief propagation. In *Proceedings of the Twenty-Fifth Conference on Uncertainty in Artificial Intelligence*, pp. 277–284. AUAI Press, 2009.
- Kipf, Thomas N and Welling, Max. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*, 2016.
- Kondor, Risi. Group theoretical methods in machine learning. *Columbia University*, 2008.
- Krizhevsky, Alex, Sutskever, Ilya, and Hinton, Geoffrey E. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pp. 1097–1105, 2012.
- Lenc, Karel and Vedaldi, Andrea. Understanding image representations by measuring their equivariance and equivalence. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 991–999, 2015.
- Niepert, Mathias. Markov chains on orbits of permutation groups. *arXiv preprint arXiv:1206.5396*, 2012.
- Olah, Christopher. Groups and group convolutions. <http://colah.github.io/posts/2014-12-Groups-Convolution/>, 2014.

Oyallon, Edouard and Mallat, Stéphane. Deep roto-translation scattering for object classification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 2865–2873, 2015.

Raedt, Luc De, Kersting, Kristian, Natarajan, Sriraam, and Poole, David. Statistical relational artificial intelligence: Logic, probability, and computation. *Synthesis Lectures on Artificial Intelligence and Machine Learning*, 10(2): 1–189, 2016.

Ravanbakhsh, Siamak, Schneider, Jeff, and Poczos, Barnabas. Deep learning with sets and point clouds. *arXiv preprint arXiv:1611.04500*, 2016.

Sifre, Laurent and Mallat, Stéphane. Rotation, scaling and deformation invariant scattering for texture discrimination. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 1233–1240, 2013.

Wielandt, H. *Permutation groups through invariant relations and invariant functions*. Dept. of Mathematics, Ohio State University, 1969.

Zaheer, Manzil, Kottur, Satwik, Ravanbakhsh, Siamak, Poczos, Barnabas, Salakhutdinov, Ruslan, and Smola, Alexander J. Deep sets. *CoRR*, abs/1703.06114, 2017. URL <http://arxiv.org/abs/1703.06114>.