
Neural Discovery of Permutation Subgroups

Pavan Karjol

Rohan Kashyap

Prathosh A P

Department of Electrical Communication Engineering,
Indian Institute of Science, Bengaluru, Karnataka

Abstract

We consider the problem of discovering subgroup H of permutation group S_n . Unlike the traditional H -invariant networks wherein H is assumed to be known, we present a method to discover the underlying subgroup, given that it satisfies certain conditions. Our results show that one could discover any subgroup of type $S_k (k \leq n)$ by learning an S_n -invariant function and a linear transformation. We also prove similar results for cyclic and dihedral subgroups. Finally, we provide a general theorem that can be extended to discover other subgroups of S_n . We also demonstrate the applicability of our results through numerical experiments on image-digit sum and symmetric polynomial regression tasks.

works proposed by (LeCun et al., 1995) incorporate translation equivariance and thus preserve translation symmetry. This is highly effective for perception tasks since it enables the model with a notion of locality and symmetry, i.e., the input and label are both invariant to shifts (preserves this property across layers), and has likewise shown substantial gains in image recognition tasks as demonstrated in (Szegedy et al., 2017; He et al., 2016). However, from a group-theoretic perspective, CNN happens to represent a particular case of invariance under the action of a specific group. This leads to studying and understanding its usability when extended to a more general setting, i.e., equivariance or invariance to any generic group action. Thus, learning such representations across neural nets ensures preserving symmetry across the network and efficiently discovering the underlying factors of data variations by utilizing these priors.

1 INTRODUCTION

1.1 Background

Deep Learning has proven to be a successful paradigm for learning the underlying regularities of sensory data such as images, text, and audio (Brown et al., 2020; He et al., 2016; Ramesh et al., 2022). The data in the physical world possess a predefined structure with a low-dimensional manifold approximation within a higher dimensional euclidean space (Cayton, 2005; Schölkopf et al., 1998). However, the task of supervised learning in such a high-dimensional data space demands a large number of data points to counter the curse of dimensionality. Thus, universal function approximations using neural networks in such a setting can be prohibitively expensive to curate large datasets for diverse applications such as medical imaging. This calls for the need for inductive bias to be incorporated into our networks such that they can utilize these priors for learning valuable representations in the feature space. Convolutional Neural Net-

1.2 Group Invariance and Equivariance

Learning symmetries from data has been studied extensively in (Senior et al., 2020; Raviv et al., 2007; Monti et al., 2017; Rossi et al., 2022). Invariant and equivariant classes of functions impose a powerful inductive prior to our models in a statistically efficient manner which aids in learning useful representations on a wide range of data (Bogatskiy et al., 2020; Esteves, 2020). Group equivariant or invariant networks (Cohen et al., 2018; Esteves et al., 2018) exploit the inherent symmetrical structure in the data, i.e., equivariance or invariance to a certain set of group operations (geometric priors) and can thus result in a significant reduction in the sample complexity and lead to better generalization. This has ubiquitous applications in various tasks such as predicting protein interactions (Gainza et al., 2020) and estimating population statistics (Zaheer et al., 2017).

One of the important classes of group invariance networks corresponds to the permutation group (S_n), i.e., the group of all permutations of a set of cardinality n . Zaheer et al. (2017) have focused extensively on the applicability of permutation equivariance and invariance functions on arbitrary objects such as sets. Whereas, (Kicki et al., 2020) proposes a G -invariant network to approximate functions that are in-

variant under the action of any given permutation subgroup of S_n . Moreover, it is crucial to consider subgroups of S_n , since any finite group is isomorphic to a subgroup of S_n (*Cayley's theorem*) for some n . For example, the Quaternionic group Q_8 is isomorphic to a subgroup of S_8 . In addition, other interesting applications of functions correspond to subgroups of S_n . For instance, the area of an n -polygon is a \mathbb{Z}_n -invariant function of the polygon's vertices (Kicki et al., 2020).

1.3 Contributions

In most of the works mentioned earlier, the group (or subgroup) is assumed to be known a priori. This restricted form of modeling choice leads to reduced flexibility (also restrictions). It makes incorporating symmetries into our networks highly infeasible for real-world applications where the underlying structure is unknown. Motivated by this, we demonstrate a general framework, i.e., G -invariant network and a linear transformation for discovering the underlying subgroup of S_n under certain conditions. Our main contributions can be summarized as follows:

In this work, we propose a general framework to discover the underlying subgroup of S_n under a broad set of conditions.

- We prove that we could learn any conjugate group (with respect to G) via a linear transformation and G -invariant network.
- We extend this approach, i.e., a linear transformation and G -invariant network to different classes of subgroups such as permutation group of k (out of n) elements S_k , cyclic subgroups \mathbb{Z}_k and dihedral subgroups D_{2k} . The G -invariant networks for the above families are S_n, \mathbb{Z}_n and D_{2n} respectively. In the latter two cases, k should divide n .
- We prove a general theorem that can guide us to discover other classes of subgroups.
- We substantiate the above results through experiments on image-digit sum and symmetric polynomial regression tasks.

2 PRIOR WORK

2.1 Group Invariant and Equivariant Networks

Significant progress has been made in incorporating invariances to deep neural nets in the last decade (Cohen et al., 2019; Cohen and Welling, 2016b; Ravanbakhsh et al., 2017; Ravanbakhsh, 2020; Wang et al., 2020). We observe that most of the invariant neural networks proposed in the literature assume the knowledge of the underlying symmetry group. Various generalizations, i.e., group equivariant

or invariant neural networks, are presented in (Cohen et al., 2019; Kondor et al., 2018).

Cohen and Welling (2016a) introduce *Group Equivariant Convolutional Neural Networks (G-CNNs)* as a natural extension of the Convolutional Neural Network to construct a representation with the structure of a linear G -space. Further, Cohen et al. (2019) presents a general theory for studying G-CNNs on homogeneous spaces and illustrates a one-to-one correspondence between linear equivariant maps of feature spaces and convolutions kernels. Cohen and Welling (2016b) provides a theoretical framework to study steerable representations in convolutional neural networks and establish mathematical connections between representation learning and representation theory. Ravanbakhsh (2020) presents the universality of invariant and equivariant MLPs with a single hidden layer. Additionally, they show the unconditional universality result for Abelian groups. Kondor and Trivedi (2018) utilize both representation theory and noncommutative harmonic analysis to establish the convolution formulae in a more general setting, i.e., invariance under the action of any compact group.

2.2 Permutation Invariant and Equivariant Networks

Zaheer et al. (2017) demonstrates the applicability of equivariant and invariant networks on various set-like objects. Further, they show that any permutation invariant function can be expressed in a standard form, i.e., $\rho(\sum_i \phi(x_i))$, which corresponds to an elegant deep neural network architecture. Janossy pooling (Murphy et al., 2018) extends the same to build permutation invariant functions using a generic class of functions. The works, as mentioned earlier, focus mainly on the permutation group S_n .

Recent works by Kicki et al. (2020) and Maron et al. (2019) provide a general architecture invariant to any given subgroup of S_n . Kicki et al. (2020) design a G -invariant neural network for approximating functions (can specifically approximate any G -invariant function) $f : X \rightarrow R$ using G -equivariant network and sum-product formulation, where X is a compact subset of $R^{n \times m}$, for some $n, m > 0$) for any given permutation subgroup G of S_n . They extend this work to study the invariance properties of hierarchical groups $G < H \leq S_n$. However, in most cases, the underlying subgroup is generally unknown.

2.3 Automatic Symmetry Discovery

Dehmamy et al. (2021) introduces the *Lie algebra convolutional network (L-Conv)*, an infinitesimal version of G-Conv, for automatic symmetric discovery. Their framework for continuous symmetries relies on *Lie algebras* rather than *Lie groups* and can thus encode an infinite group without discretizing (Cohen and Welling, 2016a) or summing over irreps. They show that the L-Conv network can serve as a building block for constructing any group equivari-

ant feedforward architecture. They also unveil interesting connections between equivariant loss and Lagrangians in field theory and robustness and Euler-Lagrange equations. However, these apply only to *Lie groups* and are not specific to subgroups of the permutation groups. Anselmi et al. (2019) proposes to learn symmetry-adapted representations and also deduce a regularization scheme for learning these representations without assuming the knowledge of the underlying subgroup (of S_n). However, their proposed solution is implemented in an unsupervised way. Benton et al. (2020) and Zhou et al. (2020) also propose different methods for learning symmetries when the group G is unknown.

3 PRELIMINARIES

This section gives a brief overview of various mathematical concepts used in our work. Let G be a group.

1. **Group action** :- The action of G on a set X is defined using the following map (written as $g \cdot x$, $\forall g \in G$ and $x \in X$):

$$\theta : G \times X \rightarrow X, \quad (1)$$

satisfying the following properties :

- $g_1 \cdot (g_2 \cdot x) = (g_1 g_2) \cdot x \quad \forall g_1, g_2 \in G$ and $x \in X$,
- $1 \cdot x = x, \quad \forall x \in X$

where 1 is the identity element of G .

2. **Group invariant function** :- A function $f : X \rightarrow Y$ is said to be group invariant with respect to G , if,

$$f(x) = f(g \cdot x), \quad \forall g \in G \text{ and } x \in X \quad (2)$$

We call f a G -invariant function.

3. **Group equivariant function** :- A function $f : X \rightarrow Y$ is said to be group equivariant with respect to G , if for any $g \in G$, $\exists \tilde{g} \in G$, such that

$$f(g \cdot x) = \tilde{g} \cdot f(x), \forall x \in X \quad (3)$$

We call f a G -equivariant function.

4. **Conjugate subgroups** :- Two subgroups G_1 and G_2 of G are said to be conjugates, if $\exists g \in G$ such that,

$$G_2 = gG_1g^{-1} := \{gkg^{-1} : k \in G_1\} \quad (4)$$

5. **Normal subgroup** :- A subgroup N is said to be normal in G , if $\forall g \in G$

$$gNg^{-1} = N \quad (5)$$

i.e., there are no subgroups that are conjugate to N .

We describe the notations used for various subgroups of S_n in Table (1). Henceforth, unless explicitly mentioned, we follow the notations mentioned in Table (1).

Table 1: Descriptions of notations

Symbol	Description
S_n	Permutation group of n elements
$S_k^{(0)}$	Permutation subgroup of first k elements
S_k	Permutation subgroup of random k elements
Z_n	Cyclic subgroup of n elements
$Z_k^{(0)}$	Cyclic subgroup of first k elements
Z_k	Cyclic subgroup of random k elements
D_{2n}	Dihedral subgroup of n elements
$D_{2k}^{(0)}$	Dihedral subgroup of first k elements
D_k	Dihedral subgroup of random k elements
A_n	Alternating subgroup of n elements
A_k	Alternating subgroup of random k elements

4 PROPOSED WORK

4.1 Problem statement

We consider the problem of learning an H -invariant function $f : X \rightarrow \mathbb{R}$, where $X = [0, 1]^n \subset \mathbb{R}^n$ and H is the unknown subgroup of S_n . In general, learning such a function is intractable. However, we show that it is possible to learn such a function, i.e., discover the underlying subgroup H , where H belongs to a certain class of subgroups (we explicitly state our conditions in Theorem 4.3, 4.4 and 4.5). The general consequence of our analysis is that learning a H -invariant function is thus equivalent to learning a G -invariant function along with a linear transformation, given that G and H satisfy certain conditions. Since any given G can have several such subgroups, we propose to learn the underlying subgroup H by exploiting the existing structures using a family of G -invariant functions (such as the one mentioned in Zaheer et al. (2017) for the permutation group S_n) and a learnable linear transformation. We formalize these ideas in the coming subsections.

To prove our results, we employ the following theorem regarding S_n -invariant functions (Zaheer et al., 2017), which shows that any such function can be expressed in a canonical form.

Theorem 4.1 (Deep sets). $f : X = [0, 1]^n \rightarrow \mathbb{R}$ is a permutation invariant (S_n -invariant) continuous function iff it has the representation,

$$f(x) = \rho \left(\sum_{i=1}^n \gamma(x_i) \right), \quad x = [x_1, x_2, \dots, x_n]^T \quad (6)$$

for some continuous outer and inner functions $\rho : \mathbb{R}^{n+1} \rightarrow \mathbb{R}$, $\gamma : [0, 1] \rightarrow \mathbb{R}^{n+1}$.

We get the following result if we consider the permutations of the first k elements.

Corollary 4.1.1. $f : [0, 1]^n \rightarrow \mathbb{R}$ be an S_k^0 -invariant con-

tinuous function iff it has the representation,

$$f(x) = \rho \left(\sum_{i=1}^k \gamma(x_i), x_{k+1}, \dots, x_n \right) \quad (7)$$

Proof. To prove Theorem 4.1, it has been shown that (Zaheer et al., 2017), $\mathcal{X}^{(n)} = \{x_1, x_2, \dots, x_n \in [0, 1]^n : x_1 \leq x_2 \leq x_3 \leq \dots \leq x_n\}$ is homeomorphic to $\sum_{i=1}^n \gamma(X_i)$, where

$$\gamma(t) = [1, t, t^2, \dots, t^n]^T \quad (8)$$

Hence, $\mathcal{X}^{(n:k)} = \{x_1, x_2, \dots, x_n \in [0, 1]^n : x_1 \leq x_2 \leq x_3 \leq \dots \leq x_k\}$ is homeomorphic to $\sum_{i=1}^k \gamma(X_i) \times [0, 1]^{n-k}$.

Let, $E(x) = \left[\sum_{i=1}^k \gamma(x_i), x_{k+1}, \dots, x_n \right]^T$. Then, it is an homeomorphism from $\mathcal{X}^{(n:k)}$ to $Im(E)$ (Image of E). If we set $\rho = fE^{-1}$, we get $\rho(E(x)) = f(x)$. \square

We use the same definition of γ (Zaheer et al., 2017) provided in the eq. (8) in the subsequent results as well. Now, we state our first result using the conjugacy relation between subgroups.

Lemma 4.2. Any S_k -invariant function ψ , can be realized through composition of an $S_k^{(0)}$ -invariant function ϕ and a linear transformation M , i.e., $\psi = \phi \cdot M$. In addition, ψ can be realised through the following form,

$$\psi(x) = \rho \left(\sum_{i=1}^k \gamma(m_i^T x), m_{k+1}^T x, \dots, m_n^T x \right), \quad (9)$$

where m_i is the i^{th} row of M .

Proof. Note that any S_k is conjugate to $S_k^{(0)}$. Thus, $\exists g \in S_n$ such that

$$S_k^{(0)} = gS_k g^{-1} \quad (10)$$

Let $\psi : X \rightarrow R$ be an S_k -invariant function, i.e.,

$$\begin{aligned} \psi(x) &= \psi(h \cdot x), \quad \forall h \in S_k, x \in X \\ \psi((g^{-1}g) \cdot x) &= \psi((g^{-1}ug) \cdot x), \quad \forall u \in S_k^{(0)} \\ (\psi g^{-1})(g \cdot x) &= (\psi g^{-1})(u \cdot (g \cdot x)) \\ (\psi g^{-1})(Mx) &= (\psi g^{-1})(u \cdot (Mx)) \end{aligned} \quad (11)$$

From eq. (11), we see that $\phi = \psi \cdot g^{-1}$ and $M = g$ are the desired $S_k^{(0)}$ -invariant function and the linear transformation respectively and $\phi = \psi \cdot M$. We get the second part of the result by applying Corollary 4.1.1 to ϕ . \square

We could also relax the conjugacy condition, i.e., discover subgroups of type S_k when k itself is unknown. This is formalized in the following result.

Theorem 4.3 (Subgroups of type S_k). Any S_k -invariant function ($k \leq n$) ψ , can be realised using an S_n -invariant function and a linear transformation, in specific, it can be realised through the following form,

$$\psi(x) = (\phi \cdot \hat{M})(x) = \rho \left(\left[\begin{array}{c} (I - M)x \\ \sum_{i=1}^n \gamma(m_i^T x) \end{array} \right] \right) \quad (12)$$

where $\hat{M} = \begin{bmatrix} I - M \\ M \end{bmatrix}$ and

$$\phi(y) = \left[y_1, \dots, y_n, \sum_{i=1}^n \gamma(y_{n+i}) \right]^T$$

Proof. Since S_k is conjugate to $S_k^{(0)}$, it is enough to prove the result for $S_k^{(0)}$ -invariant function. Hence, the goal is to show that $(I - M)X \times \sum_{i=1}^n \gamma(m_i^T X)$ is homeomorphic to $\sum_{m=1}^k \gamma(X_m) \times [0, 1]^{n-k}$ (from Corollary 4.1.1 and Lemma 4.2) for some linear transformation M . Suppose,

$$M = \begin{bmatrix} I_{k \times k} & 0 \\ 0 & 0 \end{bmatrix}, \quad (13)$$

then,

$$\left[\begin{array}{c} (I - M)x \\ \sum_{i=1}^n \gamma(m_i^T x) \end{array} \right] = \left[\begin{array}{c} \mathbf{0}^{(k)} \\ x_{k+1} \\ \vdots \\ x_n \\ B + \sum_{i=1}^k \gamma(x_i) \end{array} \right], \quad (14)$$

where $B = (n - k)\gamma(0)$ and $\mathbf{0}^{(k)}$ is k -dimensional zero vector. Thus, from RHS of the eq. (14), the above claim follows. (Note that, the function $Mx \mapsto \sum_{i=1}^n \gamma(m_i^T x)$ is S_n -invariant and ϕ is S_{2n}^n -invariant function). \square

We now extend our method to cyclic and dihedral subgroups of S_n and state the following result.

Theorem 4.4 (Cyclic and Dihedral subgroups). If $k|n$, any \mathbb{Z}_k -invariant (or D_{2k} -invariant) function ψ , can be realised using a \mathbb{Z}_n -invariant (or D_{2n} -invariant) function ϕ and a linear transformation, in specific, it can be realised through the following form,

$$\psi(x) = (\phi \cdot \hat{M})(x) \quad (15)$$

where $\hat{M} = \begin{bmatrix} M \\ I - L \end{bmatrix}$ for some $M, L \in \mathbb{R}^{n \times n}$.

Proof. In this proof, without loss of generality, we prove the result for $\mathbb{Z}_k^{(0)}$ -invariant function. Suppose,

$$M = \begin{bmatrix} I_{k \times k} & 0 \\ I_{k \times k} & 0 \\ \vdots & \vdots \\ I_{k \times k} & 0 \end{bmatrix}, \quad L = \begin{bmatrix} I_{k \times k} & 0 \\ 0 & 0 \end{bmatrix}, \quad (16)$$

Since $k|n$, we can stack the $I_{k \times k}$ matrices as shown in eq. (16). Then, $M : X \rightarrow X$ is defined as,

$$x = [x_1, x_2 \dots x_n]^T \mapsto Mx = \begin{bmatrix} x_1, x_2, \dots, x_k, \\ x_1, x_2, \dots, x_k, \\ \vdots \\ x_1, x_2, \dots, x_k \end{bmatrix}^T \quad (17)$$

Under the action of \mathbb{Z}_k ($h \cdot x$, for some $h \in \mathbb{Z}_k$), we get that,

$$x \xrightarrow{h} x' = [x_u, x_{u+1}, \dots, x_k, x_1, \dots, x_{u-1}]^T \quad (18)$$

which corresponds to $(g \cdot (Mx))$, for some $g \in \mathbb{Z}_n$,

$$Mx \xrightarrow{g} Mx' = \begin{bmatrix} x_u, x_{u+1}, \dots, x_k, x_1, \dots, x_{u-1} \\ x_u, x_{u+1}, \dots, x_k, x_1, \dots, x_{u-1} \\ \vdots \\ x_u, x_{u+1}, \dots, x_k, x_1, \dots, x_{u-1} \end{bmatrix}^T \quad (19)$$

Similarly, the converse is also true, i.e., \mathbb{Z}_n -action on Mx corresponds to \mathbb{Z}_k -action on x . Hence, the \mathbb{Z}_k -invariant function of x corresponds to \mathbb{Z}_n -invariance of Mx . Note that, the \mathbb{Z}_n -invariance of the function ϕ is with respect to the first n elements (out of $2n$) of its input vector. Similar proof holds for dihedral groups (D_{2k} and D_{2n}). \square

The above set of techniques can also be extended to other classes of subgroups. In this regard, we state the following general result.

Theorem 4.5. *Any H -invariant function ψ can be learnt through composing a G -invariant function ϕ with a linear transformation M , i.e., $\psi = \phi \cdot M$ if the following conditions hold,*

1. For any $h \in H$, $\exists g \in G$ such that $M(h \cdot x) = g \cdot (Mx)$, $\forall x \in X$
2. For any $g \in G$ such that $g \cdot (Mx) \in R(M)$, $\exists h \in H$ such that $M(h \cdot x) = g \cdot (Mx)$, $\forall x \in X$, where $R(M)$ is the range of M .

Proof. The claim directly follows from the following observations.

Condition (1) states that, any action $h \cdot x$ (action of H on X) corresponds to an action $g \cdot (Mx)$ (action of G on $R(M)$).

Similarly, condition (2) states that, any action $g \cdot (Mx)$ corresponds to an action $h \cdot x$. \square

5 DISCUSSION

The underlying theme from the results stated in the previous section is that we could discover any subgroup belonging to a particular class of subgroups by learning a G -invariant function and a linear transformation. Depending on the class, the chosen G varies. We further elaborate on these observations in the following subsections.

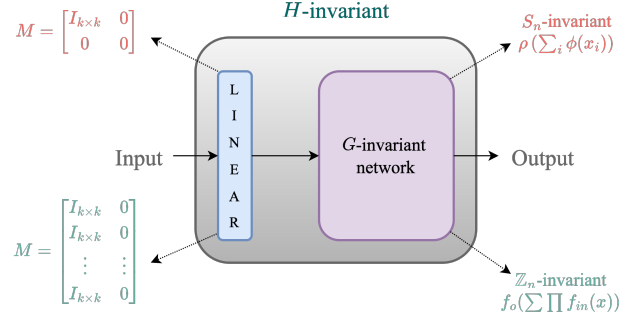


Figure 1: Generic framework for learning H -invariant function. The dotted arrows point towards specific examples of linear and G -invariant functions. The corresponding H -invariant functions are S_k -invariant and \mathbb{Z}_k -invariant.

5.1 Conjugate Groups

In Lemma 4.2, the class of subgroups corresponds to those of type S_k (fixed k) and the corresponding G can be S_k^0 . We observe that, for a fixed k , even if we don't know the exact underlying subgroup S_k (a total of $\binom{n}{k}$ possibilities), we could learn this unknown subgroup. In addition, we also incorporate the canonical form of permutation invariant functions in the resulting architecture. Moreover, this result can be generalized to any class of conjugate subgroups, and the corresponding G is one of these conjugate groups. The significance of this result lies in the fact that a variety of subgroups are related through conjugation. For instance, all \mathbb{Z}_k form one conjugacy class for a given k , and so does A_k 's.

This result is not entirely helpful if the underlying subgroup is normal since it is not conjugate to any other subgroup. However, this is not much of a hindrance since the only non-trivial proper normal subgroup of S_n is A_n , $\forall n \geq 5$.

5.2 S_k, \mathbb{Z}_k and D_{2k} Subgroups

Theorem 4.3 focuses on subgroups of type S_k (varying k and $k \in \{1, 2, \dots, n\}$), and the corresponding G is S_n itself. We incorporate the canonical form of permutation invariant functions here as well. We observe that the number of such subgroups is $2^n - 1$ for a given n . Hence, we could learn any of these subgroups with the standard ar-

chitecture of an S_n -invariant function and a linear transformation. Note that if k is fixed, either of the architectural forms given by Lemma 4.2 and Theorem 4.3 is applicable. We will discuss the corresponding empirical results in the coming sections. Theorem 4.4 considers subgroups of the cyclic \mathbb{Z}_k and dihedral group D_{2k} . The corresponding G -invariant functions are of \mathbb{Z}_n and D_{2n} , respectively.

5.3 Generalization

Theorem 4.5 presents a general set of conditions to be satisfied to learn any H -invariant function using a G invariant function and a linear transformation. As such, the previous results are specific cases of this Theorem. However, they provide explicit structures of the linear transformation M . These can help design appropriate training techniques to learn the optimum M , while the general result of Theorem 4.5 can guide us towards discovering results for new classes of subgroups.

5.4 Limitations

The proposed framework presumes the knowledge of the underlying class of subgroups a priori (but not the exact subgroup) and an appropriate value of n for S_n , \mathbb{Z}_n or D_{2n} invariant functions. The drawbacks mentioned here are interesting research directions to pursue in the future.

6 EXPERIMENTS

We evaluate the accuracy of our proposed method on image-digit sum and symmetric polynomial regression tasks. The problem of image-digit sum can be modified and cast as learning an S_k -invariant function, while the polynomial regression task intrinsically corresponds to learning a G -invariant function. These are summarized in the following subsections.

6.1 Image-Digit Sum

This task aims to find the sum of k digits using the MNISTm (Loosli et al. (2007)) handwritten digits dataset. It consists of 8 million gray scale 28×28 images of digits $\{0, 1, \dots, 9\}$. We employ a training set of $150k$ samples and a test set of $30k$ samples. We consider the following approaches for evaluation.

1. **Deep Sets- S_k** :- S_k -invariant neural network proposed by Zaheer et al. (2017).
2. **LSTM**:- LSTM network as mentioned in Zaheer et al. (2017).
3. **Proposed method**:- A linear layer followed by an S_n -invariant network.

For the LSTM network and the proposed method, the input is a random sample of n ($n = 10$) images, and the target is the sum of k (k less than n) digit labels. We run separate experiments for each of $k \in \{1, 3, 5, 7, 9\}$. Since all n images are given as input, the two approaches are agnostic of the underlying subgroup. However, we feed only these k of these images as input for the first approach, while the target output remains the same. As such, this task is equivalent to learning an S_k -invariant function.

6.2 Symmetric Polynomial Regression

We evaluate the performance of our method on symmetric polynomial regression tasks as discussed in Kicki et al. (2020), primarily for subgroups of \mathbb{Z}_{10} and \mathbb{Z}_{16} . For all our experiments, we utilize a \mathbb{Z}_n -invariant neural network with a Sum-Product layer as discussed in Kicki et al. (2020) and a linear layer. First, we run our experiments for subgroups of \mathbb{Z}_{10} , i.e., \mathbb{Z}_5 and the group itself (trivial subgroup). We then assess the performance for subgroups of \mathbb{Z}_{16} , namely \mathbb{Z}_2 , \mathbb{Z}_4 , \mathbb{Z}_8 , \mathbb{Z}_{16} using a similar architectural design. We consider the following approaches for evaluation.

1. **G-invariant**:- \mathbb{Z}_k -invariant neural network proposed by Kicki et al. (2020). In this context, $G = \mathbb{Z}_k$.
2. **Simple-FC**:- A stack of fully-connected feedforward layers.
3. **Conv-1D**:- A simple convolutional neural network and feedforward layers.
4. **Proposed method**:- A linear layer followed by a \mathbb{Z}_n -invariant network.

The architectural details of the models considered in our experiments are discussed in the appendix section.

7 RESULTS

7.1 Image-Digit Sum

The test mean absolute errors (MAEs) for the image-digit sum task are shown in Table 2. We observe that the proposed method outperforms the LSTM baseline and is competitive with respect to the Deep Sets method (k input images) when the underlying subgroup S_k is known. In addition, our method converges faster when compared to the LSTM network, which is apparent from the plots for the training and validation errors in Figure 2.

7.2 Symmetric Polynomial Regression

In the \mathbb{Z}_k -invariant polynomial regression task, we train our models for 2500 epochs for each of the subgroups of \mathbb{Z}_5 and \mathbb{Z}_{10} .

Table 2: MAE [$\times 10^{-2}$] for Image Digit-Sum task

Method	S_1	S_3	S_5	S_7	S_9
Deep Sets- S_k	5.61 ± 0.35	7.66 ± 0.26	8.02 ± 0.2	7.68 ± 0.43	6.97 ± 0.39
Proposed	5.73 ± 0.39	7.78 ± 0.49	8.19 ± 0.36	7.84 ± 0.41	7.26 ± 0.58
LSTM	6.23 ± 0.53	9.65 ± 0.57	11.98 ± 0.46	13.35 ± 1.02	12.92 ± 1.42




Figure 2: Training and Validation loss (MAE) for Image-Digit Sum using MNIST dataset.

In Table 3, 4 and 5 we compare the given baselines with our proposed method for the task of discovering unknown subgroups. Our method outperforms the Simple-FC and Conv-1D baseline networks for each of the given subgroups. As expected, it does not match the baseline architecture, the \mathbb{Z}_k -invariant network (the subgroup is known apriori for this baseline) by a significant margin for each of the diverse set of subgroups we have considered in this task. However, in a few cases, we observe large standard deviations and attribute such values to outliers. A detailed version of our results and the mathematical definition of the polynomials is presented in the appendix section.

From Figure 3, it is evident that the \mathbb{Z}_5 -invariant function outperforms both our method and the baselines by a significant margin. The Simple FC and Conv-1D networks have very similar performances and show no prominent effect, even with an increase in data size.

Table 3: MAE [$\times 10^{-2}$] for $\mathbb{Z}_5 : \mathbb{Z}_{10}$

Method	Train	Validation	Test
\mathbb{Z}_5 -invariant	2.65 ± 0.91	7.32 ± 0.55	7.53 ± 0.576
Proposed	4.48 ± 1.25	24.56 ± 6.93	24.78 ± 6.45
Conv-1D	20.90 ± 4.91	32.96 ± 1.31	32.33 ± 1.18
Simple-FC	23.86 ± 3.87	33.57 ± 2.07	33.14 ± 2.11

7.3 Effect of the data size on the performance

This section aims to assess the effect of the dataset size in learning \mathbb{Z}_k -invariant functions using our proposed method and hope to gain a better understanding in such a setting. To analyze our model performance with respect to data size,

Table 4: MAE [$\times 10^{-2}$] for $\mathbb{Z}_{10} : \mathbb{Z}_{10}$

Method	Train	Validation	Test
\mathbb{Z}_{10} -invariant	6.89 ± 1.31	16.68 ± 0.55	17.16 ± 0.56
Proposed	14.52 ± 1.72	39.69 ± 4.13	40.11 ± 4.17
Conv-1D	35.71 ± 2.71	52.96 ± 0.70	50.63 ± 1.33
Simple-FC	46.13 ± 2.27	54.62 ± 1.34	51.64 ± 0.89

Table 5: MAE [$\times 10^{-2}$] for $\mathbb{Z}_4 : \mathbb{Z}_{16}$

Method	Train	Validation	Test
\mathbb{Z}_4 -invariant	1.21 ± 0.25	3.41 ± 0.4	3.54 ± 0.39
Proposed	3.32 ± 1.65	23.70 ± 4.87	24.69 ± 5.25
Conv-1D	8.39 ± 3.02	31.34 ± 0.77	31.10 ± 0.87
Simple-FC	7.27 ± 5.03	30.82 ± 1.74	30.83 ± 1.61

we use 16, 32, and 64 data points for training (as mentioned in Kicki et al. (2020)), we randomly sample these values from [0,1] and use 480 and 4800 as validation and test sets respectively to assess the generalization ability for each of these methods as mentioned above. We report the mean and standard deviation values across 10 randomly initialized iterations.





Figure 3: The MAE value comparisons using the test dataset for all the models we have considered for the $\mathbb{Z}_5 : \mathbb{Z}_{10}$ task. The X-axis represents the size of the training set (16, 32, 64).

We also examine the Simple-FC and Conv-1D network by increasing its parameter count, i.e., varying the number of neurons in each layer. However, we observe no significant gains in doing so, as mentioned in the appendix section for at least a few subgroups.


7.4 Interpretability

7.4.1 Image-Digit Sum

The resulting M matrix is interpretable, and we consistently observe the expected pattern for the image-digit sum task. Note that any row-permuted version of the matrix structure, as shown in eq. (13) will work since the transformed space is still homeomorphic. The M matrices for S_5 and S_9 (ex-


 Figure 4: M matrices for S_5 and S_9 after training.

tracted after training) are depicted in Figure 4. The columns with dark green squares match the actual indices.


 Figure 5: (a) M matrix for $\mathbb{Z}_4 : \mathbb{Z}_{16}$ (b) Reference matrix.

7.4.2 Polynomial Regression

We observe that the M -matrix extracted after training (Figure 5.a) does not exactly capture the expected pattern, i.e., a stack of identity matrices (Figure 5.b), even though it nearly masks most of the irrelevant columns ($n - k$). The former behavior (lack of exact structure) explains the difference in performance with respect to the \mathbb{Z}_k -invariant network, while the latter (masking behavior) describes the superior model performance compared to other baselines. Also, the masking of irrelevant columns already conveys the underlying subgroup; thus, we use this information to estimate the true indices. We estimate the significant indices using the L_1 -norm of columns of M and the mean as the threshold. The results (for different number of training data points N and different $\mathbb{Z}_k : \mathbb{Z}_n$'s) of the success rate of the estimation are given in Table 6, where we count the estimation as success when the estimated indices exactly match the true indices; otherwise, as a failure. We run each experiment for 10 trials. We get high estimation accuracy in most of the cases except for $N = 16$. The estimated indices can be used to run a \mathbb{Z}_k -invariant network (or pro-

posed method with fixed M) and obtain better performance on regression tasks.

Table 6: Estimation Accuracy (in %)

$Z_k : Z_n$	16	32	64
$Z_4 : Z_{16}$	100	100	100
$Z_5 : Z_{10}$	80	100	100
$Z_8 : Z_{16}$	30	80	100

8 CONCLUSION

In this work, we studied the problem of discovering the underlying subgroup of S_n , i.e., learning a H -invariant function where H is an unknown subgroup of S_n . We proved that we could learn any H -invariant function using a G -invariant function and a linear transformation provided H belongs to a specific class of subgroups. We considered various subgroups, such as conjugate subgroups, permutation subgroups of k elements, and cyclic and dihedral subgroups, and illustrated unique structures of the corresponding linear transformations. We demonstrated the validity of our theoretical analysis through empirical results. We also discussed the limitations of our method, which may lead to exciting research directions in the future.

References

- Anselmi, F., Evangelopoulos, G., Rosasco, L., and Poggio, T. (2019). Symmetry-adapted representation learning. *Pattern Recognition*, 86:201–208.
- Benton, G., Finzi, M., Izmailov, P., and Wilson, A. G. (2020). Learning invariances in neural networks from training data. *Advances in neural information processing systems*, 33:17605–17616.
- Bogatskiy, A., Anderson, B., Offermann, J., Roussi, M., Miller, D., and Kondor, R. (2020). Lorentz group equivariant neural network for particle physics. In *International Conference on Machine Learning*, pages 992–1002. PMLR.
- Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J. D., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., et al. (2020). Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.
- Cayton, L. (2005). Algorithms for manifold learning. *Univ. of California at San Diego Tech. Rep.*, 12(1-17):1.
- Cohen, T. and Welling, M. (2016a). Group equivariant convolutional networks. In *International conference on machine learning*, pages 2990–2999. PMLR.
- Cohen, T. S., Geiger, M., Köhler, J., and Welling, M. (2018). Spherical cnns. *arXiv preprint arXiv:1801.10130*.

- Cohen, T. S., Geiger, M., and Weiler, M. (2019). A general theory of equivariant cnns on homogeneous spaces. *Advances in neural information processing systems*, 32.
- Cohen, T. S. and Welling, M. (2016b). Steerable cnns. *arXiv preprint arXiv:1612.08498*.
- Dehmamy, N., Walters, R., Liu, Y., Wang, D., and Yu, R. (2021). Automatic symmetry discovery with lie algebra convolutional network. *Advances in Neural Information Processing Systems*, 34:2503–2515.
- Derksen, H. and Kemper, G. (2001). Computational invariant theory. *Book manuscript*.
- Esteves, C. (2020). Theoretical aspects of group equivariant neural networks. *arXiv preprint arXiv:2004.05154*.
- Esteves, C., Allen-Blanchette, C., Makadia, A., and Daniilidis, K. (2018). Learning so (3) equivariant representations with spherical cnns. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 52–68.
- Gainza, P., Sverrisson, F., Monti, F., Rodola, E., Boscaini, D., Bronstein, M., and Correia, B. (2020). Deciphering interaction fingerprints from protein molecular surfaces using geometric deep learning. *Nature Methods*, 17(2):184–192.
- He, K., Zhang, X., Ren, S., and Sun, J. (2016). Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778.
- Kicki, P., Ozay, M., and Skrzypczyński, P. (2020). A computationally efficient neural network invariant to the action of symmetry subgroups. *arXiv preprint arXiv:2002.07528*.
- Kondor, R., Lin, Z., and Trivedi, S. (2018). Clebsch-gordan nets: a fully fourier space spherical convolutional neural network. *Advances in Neural Information Processing Systems*, 31.
- Kondor, R. and Trivedi, S. (2018). On the generalization of equivariance and convolution in neural networks to the action of compact groups. In *International Conference on Machine Learning*, pages 2747–2755. PMLR.
- LeCun, Y., Bengio, Y., et al. (1995). Convolutional networks for images, speech, and time series. *The handbook of brain theory and neural networks*, 3361(10):1995.
- Loosli, G., Canu, S., and Bottou, L. (2007). Training invariant support vector machines using selective sampling. *Large scale kernel machines*, 2.
- Maron, H., Fetaya, E., Segol, N., and Lipman, Y. (2019). On the universality of invariant networks. In *International conference on machine learning*, pages 4363–4371. PMLR.
- Monti, F., Boscaini, D., Masci, J., Rodola, E., Svoboda, J., and Bronstein, M. M. (2017). Geometric deep learning on graphs and manifolds using mixture model cnns. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5115–5124.
- Murphy, R. L., Srinivasan, B., Rao, V., and Ribeiro, B. (2018). Janosy pooling: Learning deep permutation-invariant functions for variable-size inputs. *arXiv preprint arXiv:1811.01900*.
- Ramesh, A., Dhariwal, P., Nichol, A., Chu, C., and Chen, M. (2022). Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*.
- Ravanbakhsh, S. (2020). Universal equivariant multilayer perceptrons. In *International Conference on Machine Learning*, pages 7996–8006. PMLR.
- Ravanbakhsh, S., Schneider, J., and Póczos, B. (2017). Equivariance through parameter-sharing. In *International conference on machine learning*, pages 2892–2901. PMLR.
- Raviv, D., Bronstein, A. M., Bronstein, M. M., and Kimmel, R. (2007). Symmetries of non-rigid shapes. In *2007 IEEE 11th International Conference on Computer Vision*, pages 1–7. IEEE.
- Rossi, E., Monti, F., Leng, Y., Bronstein, M., and Dong, X. (2022). Learning to infer structures of network games. In *International Conference on Machine Learning*, pages 18809–18827. PMLR.
- Schölkopf, B., Smola, A., and Müller, K.-R. (1998). Non-linear component analysis as a kernel eigenvalue problem. *Neural computation*, 10(5):1299–1319.
- Senior, A. W., Evans, R., Jumper, J., Kirkpatrick, J., Sifre, L., Green, T., Qin, C., Žídek, A., Nelson, A. W., Bridgland, A., et al. (2020). Improved protein structure prediction using potentials from deep learning. *Nature*, 577(7792):706–710.
- Szegedy, C., Ioffe, S., Vanhoucke, V., and Alemi, A. A. (2017). Inception-v4, inception-resnet and the impact of residual connections on learning. In *Thirty-first AAAI conference on artificial intelligence*.
- Wang, R., Walters, R., and Yu, R. (2020). Incorporating symmetry into deep dynamics models for improved generalization. *arXiv preprint arXiv:2002.03061*.
- Zaheer, M., Kottur, S., Ravanbakhsh, S., Póczos, B., Salakhutdinov, R. R., and Smola, A. J. (2017). Deep sets. *Advances in neural information processing systems*, 30.
- Zhou, A., Knowles, T., and Finn, C. (2020). Meta-learning symmetries by reparameterization. *arXiv preprint arXiv:2007.02933*.

Supplementary Materials

9 Appendix

9.1 Remarks regarding theoretical results

- To obtain the function ρ mentioned in Lemma 4.2, Theorem 4.3, and Theorem 4.4 of the main paper, we use the same technique presented in the proof of Corollary 4.1.1, i.e., $\rho = \psi E^{-1}$, where E is the corresponding homeomorphism.
- Lemma 4.2, Theorem 4.3, and Theorem 4.4 are special cases of Theorem 4.5. This claim directly follows once we specify the corresponding group actions. As such, the proofs of Lemma 4.2 and Theorem 4.4 already describe the required group actions. However, this is not obvious in Theorem 4.3. In that case, we observe the following:

$$x = [x_1, x_2 \dots x_n]^T \mapsto Mx = [0, 0, \dots, 0, x_{k+1}, x_{k+2}, \dots, x_n, x_1, x_2, \dots, x_k, 0, 0, \dots, 0]^T, \quad (20)$$

where $M = \begin{bmatrix} I_{k \times k} & 0 \\ 0 & 0 \end{bmatrix}$ as mentioned in eq. (13) of the main paper.

Under the action of $S_k^{(0)}$ ($h \cdot x$, for some $h \in S_k^{(0)}$), we get,

$$x \xrightarrow{h} x' = [x_{h(1)}, x_{h(2)}, \dots, x_{h(k)}, x_{k+1}, \dots, x_n]^T \quad (21)$$

which corresponds to ($g \cdot (Mx)$, for some $g \in S_{2n}^n$) (S_{2n}^n is the group of permutations of first n elements out of $2n$ elements),

$$Mx \xrightarrow{g} Mx' = [0, 0, \dots, 0, x_{g(k+1)}, x_{g(k+2)}, \dots, x_n, x_{g(1)}, x_{g(2)}, \dots, x_{g(k)}, 0, 0, \dots, 0]^T \quad (22)$$

Similarly, the action of S_{2n}^n on $R(M)$ (range of M) corresponds to the action of $S_k^{(0)}$ on X .

In the following subsections, we describe the architectures of various models and additional results considered in our experiments.

Table 7: MAE [$\times 10^{-2}$] Image Digit-Sum task

Method	S_1	S_3	S_5	S_7	S_9
LSTM	6.23 ± 0.53	9.65 ± 0.57	11.98 ± 0.46	13.35 ± 1.02	12.92 ± 1.42
Conv-1D	36.32 ± 0.12	19.11 ± 0.49	27.92 ± 0.41	35.42 ± 0.19	40.83 ± 0.11
Simple FC	25.26 ± 0.01	18.18 ± 0.15	35.27 ± 0.07	44.51 ± 0.58	51.79 ± 0.89

9.2 Image-Digit sum

1. **Deep Sets- S_k** :- S_k -invariant neural network proposed by Zaheer et al. (2017). It consists of two networks, γ , and ρ . Each element in the input is passed through the γ network, followed by the sum operation. The result is then fed to the second network ρ . The network γ is a feed-forward network consisting of three *dense* layers with *tanh* activation, and the second network is a *dense* layer.
2. **LSTM**:- The LSTM network used for comparison in Zaheer et al. (2017). It consists of two *dense* layers, an *LSTM* layer followed by two *dense* layers. The activation used is *tanh* function.
3. **Proposed method**:- An S_n -invariant network follows a linear layer. The S_n -invariant network has the same architecture as Deep Sets (the first approach) except the input layer.

Table 8: Definitions of the various polynomials used in the main paper.

INVARIANCE	POLYNOMIAL
$\mathbb{Z}_2 : \mathbb{Z}_{16}$	$x_1x_2^2 + x_2x_1^2$
$\mathbb{Z}_4 : \mathbb{Z}_{16}$	$x_1x_2^2 + x_2x_3^2 + x_3x_4^2 + x_4x_1^2$
$\mathbb{Z}_5 : \mathbb{Z}_{10}$	$x_1x_2^2 + x_2x_3^2 + x_3x_4^2 + x_4x_5^2 + x_5x_1^2$
$\mathbb{Z}_8 : \mathbb{Z}_{16}$	$x_1x_2^2 + x_2x_3^2 + \dots + x_7x_8^2 + x_8x_1^2$
$\mathbb{Z}_{10} : \mathbb{Z}_{10}$	$x_1x_2^2 + x_2x_3^2 + \dots + x_9x_{10}^2 + x_{10}x_1^2$
$\mathbb{Z}_{16} : \mathbb{Z}_{16}$	$x_1x_2^2 + x_2x_3^2 + \dots + x_{15}x_{16}^2 + x_{16}x_1^2$

9.3 Comparison between S_k -invariant networks with backbone as $S_k^{(0)}$ and S_n

As specified in Section 5.2 of the main paper, any S_k -invariant network can be realized through either an $S_k^{(0)}$ or an S_n -invariant network and a linear layer when k is fixed. In general, we observed that the S_n -invariant network as the backbone does better than the $S_k^{(0)}$ network. We attribute this to the expressivity power of the linear transformation (based on its specific structure) when an S_n invariant network is used.

9.4 Symmetric Polynomial Regression

- G-invariant**:- \mathbb{Z}_k -invariant network implemented using the design described in Kicki et al. (2020). As discussed in Kicki et al. (2020), it is a composition of a \mathbb{Z}_k -equivariant network and a Sum-Product Layer. It then uses a Multi-Layer Perceptron to process the \mathbb{Z}_k -invariant representation of the input and thus predicts the polynomial output. The network is thus invariant under the action of the given permutation subgroup \mathbb{Z}_k .
- Simple-FC**:- This is an abbreviation of a fully-connected neural network without the Reynolds operator, i.e., group averaging for this baseline implementation (Derksen and Kemper (2001)).
- Conv-1D**:- This is an abbreviation of the 1D Convolutional neural network equipped with fully-connected layers.
- Proposed method**:- To discover the underlying subgroup, we use a \mathbb{Z}_n -invariant neural network with the addition of a linear layer. The architectural design of the \mathbb{Z}_n -invariant function is the same as the G -invariant network.

The hyperparameters of the above models are given in Table. (6-9) of Kicki et al. (2020).

Table 9: MAE [$\times 10^{-2}$] Polynomial Regression

Method	$\mathbb{Z}_2 : \mathbb{Z}_{16}$	$\mathbb{Z}_8 : \mathbb{Z}_{16}$	$\mathbb{Z}_{16} : \mathbb{Z}_{16}$
\mathbb{Z}_k -invariant	1.26 ± 0.25	14.30 ± 1.04	17.16 ± 0.59
Proposed	22.27 ± 4.25	39.31 ± 5.12	40.11 ± 4.17
Conv-1D	21.67 ± 0.69	50.15 ± 1.03	68.38 ± 3.13
Simple FC	21.61 ± 1.54	45.92 ± 2.83	51.64 ± 0.89

9.4.1 Additional results

The test errors for $\mathbb{Z}_2 : \mathbb{Z}_{16}$, $\mathbb{Z}_8 : \mathbb{Z}_{16}$ and $\mathbb{Z}_{16} : \mathbb{Z}_{16}$ invariant functions are provided in Table 9. The details of all the polynomials used in our work are presented in Table 8.