
Monomial Matrix Group Equivariant Neural Functional Networks

Viet-Hoang Tran*

Department of Mathematics
National University of Singapore
hoang.tranviet@u.nus.edu

Thieu N. Vo*

Department of Mathematics
National University of Singapore
thieuvo@nus.edu.sg

Tho Tran-Huu

Department of Mathematics
National University of Singapore
thotranhuu@u.nus.edu.vn

An T. Nguyen

FPT Software AI Center
annt68@fpt.com

Tan M. Nguyen

Department of Mathematics
National University of Singapore
tanmn@nus.edu.sg

Abstract

Neural functional networks (NFNs) have recently gained significant attention due to their diverse applications, ranging from predicting network generalization and network editing to classifying implicit neural representation. Previous NFN designs often depend on permutation symmetries in neural networks' weights, which traditionally arise from the unordered arrangement of neurons in hidden layers. However, these designs do not take into account the weight scaling symmetries of ReLU networks, and the weight sign flipping symmetries of sin or Tanh networks. In this paper, we extend the study of the group action on the network weights from the group of permutation matrices to the group of monomial matrices by incorporating scaling/sign-flipping symmetries. Particularly, we encode these scaling/sign-flipping symmetries by designing our corresponding equivariant and invariant layers. We name our new family of NFNs the Monomial Matrix Group Equivariant Neural Functional Networks (Monomial-NFN). Because of the expansion of the symmetries, Monomial-NFN has much fewer independent trainable parameters compared to the baseline NFNs in the literature, thus enhancing the model's efficiency. Moreover, for fully connected and convolutional neural networks, we theoretically prove that all groups that leave these networks invariant while acting on their weight spaces are some subgroups of the monomial matrix group. We provide empirical evidences to demonstrate the advantages of our model over existing baselines, achieving competitive performance and efficiency. The code is publicly available at <https://github.com/MathematicalAI-NUS/Monomial-NFN>.

1 Introduction

Deep neural networks (DNNs) have become highly versatile modeling tools, finding applications across a broad spectrum of fields such as Natural Language Processing [15, 29, 51, 66], Computer

*Equal contribution. Please correspond to thieuvo@nus.edu.sg.

Vision [27, 37, 61], and the Natural Sciences [31, 49]. There has been growing interest in developing specialized neural networks to process the weights, gradients, or sparsity masks of DNNs as data. These specialized neural networks are called neural functional networks (NFNs) [71]. NFNs have found diverse applications, ranging from predicting network generalization and network editing to classifying implicit neural representations. For instance, NFNs have been employed to create learnable optimizers for neural network training [5, 11, 42, 52], extract information from implicit neural representations of data [43, 52, 60], perform corrective editing of network weights [13, 44, 56], evaluate policies [26], and conduct Bayesian inference using networks as evidence [59].

Developing NFNs is inherently challenging due to their high-dimensional nature. Some early methods to address this challenge assume a restricted training process that effectively reduced the weight space [9, 19, 41]. More recent efforts have focused on building permutation equivariant NFNs that can process neural network weights without such restrictions [35, 45, 71, 72]. These works construct NFNs that are equivariant to permutations of weights, corresponded to the rearrangement of neurons in hidden layers. Such permutations, known as neuron permutation symmetries, preserve the network’s behavior. However, these approaches often overlook other significant symmetries in weight spaces [12, 24]. Notable examples are weight scaling transformations for ReLU networks [7, 12, 46] and sign flipping transformations for sin and tanh networks [14, 22, 38]. Consequently, two weight spaces of a ReLU networks, that differ by a scaling transformation, two weight spaces of a sin, or tanh networks that differ by a sign flipping transformation, can produce different results when processed by existing permutation equivariant NFNs, despite representing the same functions. This highlights a fundamental limitation of the current permutation equivariant NFNs.

Contribution. In this paper, we extend the study of symmetries in weight spaces of Fully Connected Neural Networks (FCNNs) and Convolution Neural Networks (CNNs) by formally establishing a group of symmetries that includes both neuron permutations and scaling/sign-flipping transformations. These symmetries are represented by monomial matrices, which share the nonzero pattern of permutation matrices but allow nonzero entries to be any value rather than just 1. We then introduce a novel family of NFNs that are equivariant to groups of monomial matrices, thus incorporating both permutation and scaling/sign-flipping symmetries into the NFN design. We name this new family Monomial Matrix Group Equivariant Neural Functional Networks (Monomial-NFNs). Due to the expanded set of symmetries, Monomial-NFN requires significantly fewer independent trainable parameters compared to baseline NFNs, enhancing the model’s efficiency. By incorporating equivariance to neuron permutations and weight scaling/sign-flipping, our NFNs demonstrate competitive generalization performance compared to existing models. Our contribution is three-fold:

1. We formally describe a group of monomial matrices satisfying the condition that the transformation of weight spaces of FCNNs and CNNs using these group elements does not change the function defined by the networks. For ReLU networks, this group covers permutation and scaling symmetries of the weight spaces, while for sin or tanh networks, this group covers permutation and sign-flipping symmetries. The group is proved to be maximal in certain cases.
2. We design Monomial-NFNs, the first family of NFNs that incorporate scaling and sign-flipping symmetries of weight spaces as far as we are aware. The main building blocks of Monomial-NFNs are the equivariant and invariant linear layers for processing weight spaces.
3. We show that the number of parameters in our equivariant linear layer is much lower than in recent permutation equivariant NFNs. In particular, our method is linear in the number of layers and dimensions of weights and biases, compared to quadratic as in [71]. This demonstrates that Monomial-NFNs have the ability to process weight spaces of large-scale networks.

We evaluate Monomial-NFNs on three tasks: predicting CNN generalization from weights using Small CNN Zoo [64], weight space style editing, and classifying INRs using INRs data [71]. Experimental results show that our model achieves competitive performance and efficiency compared to existing baselines.

Organization. We structure this paper as follows: After summarizing some related work in Section 2, we recall the notions of monomial matrix group and describe their maximal subgroups preserved by some nonlinear activations in Section 3. In Section 4, we formalize the general weight space of FCNNs and CNNs, then discuss the symmetries of these weight spaces using the monomial matrices. In Section 5, we construct monomial matrix group equivariant and invariant layers, which are building blocks for our Monomial-NFNs. In Section 6, we present our experimental results to justify the advantages of Monomial-NFNs over the existing permutation equivariant NFN baselines. The paper ends with concluding remarks. More experimental details are provided in the Appendix.

2 Related Work

Symmetries of Weight Spaces. The challenge of identifying the symmetries in the weight spaces of neural networks, or equivalently, determining the functional equivalence of neural networks, is a well-explored area in academic research [3, 10, 16, 23, 47]. This problem was initially posed by Hecht-Nielsen in [28]. Results for various types of networks have been established as in [1, 2, 12, 14, 22, 38, 63].

Neural Functional Networks. Recent research has focused on learning representations for trained classifiers to predict their generalization performance and other insights into neural networks [8, 20, 64, 54, 53, 55]. In particular, low-dimensional encodings for Implicit Neural Representations (INRs) have been developed for downstream tasks [18, 41]. Other studies have encoded and decoded neural network parameters mainly for reconstruction and generation purposes [6, 21, 34, 48].

Equivariant Neural Functional Networks. Permutations and scaling, for ReLU networks, as well as sign-flipping, for sine or tanh networks, symmetries, are fundamental symmetries of weight networks. Permutation-equivariant NFNs are successfully built in [4, 35, 40, 45, 70, 71, 72]. In particular, the authors in [35, 40] carefully construct computational graphs representing the input neural networks' parameters and process the graphs using graph neural networks. In [4], neural network parameters are efficiently encoded by carefully choosing appropriate set-to-set and set-to-vector functions. The authors in [70] view network parameters as a special case of a collection of tensors and then construct maximally expressive equivariant linear layers for processing any collection of tensors given a description of their permutation symmetries. These methods are applicable to several types of networks, including those with branches or transformers. However, these models were not necessarily equivariant to scaling nor sign-flipping transformations, which are important symmetries of the input neural networks.

Our method makes the first step toward incorporating both permutation and non-permutation symmetries into NFNs. In particular, the model proposed in our paper is equivariant to permutations and scaling, for ReLU networks, or sign-flipping, for sine and tanh networks. This leads to a significant reduction in the number of parameters, a property that is particularly useful for large neural networks in modern deep learning, while achieving comparable or better results than those in the literature. The authors in [32, 67] have also developed NFNs that incorporates scaling symmetries.

3 Monomial Matrices Preserved by a Nonlinear Activation

Given two sets X, Y , and a group G acts on them, a function $\phi: X \rightarrow Y$ is called G -equivariant if $\phi(g \cdot x) = g \cdot \phi(x)$ for all $x \in X$ and $g \in G$. If G acts trivially on Y , then we say ϕ is G -invariant. In this paper, we consider NFNs which are equivariant with respect to certain symmetries of deep weight spaces. These symmetries will be represented by monomial matrices. In Subsection 3.1, we recall the notion of monomial matrices, as well as their actions on space of matrices. We then formalize the maximal group of matrices preserved by the activations ReLU, sin and Tanh in Subsection 3.2.

3.1 Monomial Matrices and Monomial Matrix Group Actions

All matrices considered in this paper have real entries and n is a positive integer.

Definition 3.1 (See [50, page 46]). A matrix of size $n \times n$ is called a *monomial matrix* (or *generalized permutation matrix*) if it has exactly one non-zero entry in each row and each column, and zeros elsewhere. We will denote by \mathcal{G}_n the set of such all matrices.

Permutation matrices and *invertible diagonal matrices* are special cases of monomial matrices. In particular, a permutation matrix is a monomial matrix in which the non-zero entries are all equal to 1. In case the nonzero entries of a monomial matrix are in the diagonal, it becomes an invertible diagonal matrix. We will denote by \mathcal{P}_n and Δ_n the sets of permutation matrices and invertible diagonal matrices of size $n \times n$, respectively. It is well-known that the groups \mathcal{G}_n , \mathcal{P}_n , and Δ_n are subgroups of the general linear group $\text{GL}(n)$.

Permutation matrix group \mathcal{P}_n is a representation of the permutation group S_n , which is the group of all permutations of the set $\{1, 2, \dots, n\}$ with group operator as the composition. Indeed, for each permutation $\pi \in S_n$, we denote by P_π the square matrix obtained by permuting n columns of the identity matrix I_n by π . We call P_π the *permutation matrix* corresponding to π . The correspondence $\pi \mapsto P_\pi$ defines a group homomorphism $\rho: S_n \rightarrow \text{GL}(n)$ with the image $\mathcal{P}_n = \rho(S_n)$.

Each monomial matrix in \mathcal{G}_n is a product of an invertible diagonal matrix in Δ_n and a permutation matrix in \mathcal{P}_n , i.e.

$$\mathcal{G}_n = \{DP : D \in \Delta_n \text{ and } P \in \mathcal{P}_n\}. \quad (1)$$

In general, we have $PD \neq DP$. However, for $D = \text{diag}(d_1, d_2, \dots, d_n)$ and $P = P_\pi$, we have $PD = (PDP^{-1})P$ which is again a product of the invertible diagonal matrix

$$PDP^{-1} = \text{diag}(d_{\pi^{-1}(1)}, d_{\pi^{-1}(2)}, \dots, d_{\pi^{-1}(n)}) \quad (2)$$

and the permutation matrix P . As an implication of Eq. (2), there is a group homomorphism $\varphi: \mathcal{P}_n \rightarrow \text{Aut}(\Delta_n)$, defined by the conjugation, i.e. $\varphi(P)(D) = PDP^{-1}$ for all $P \in \mathcal{P}_n$ and $D \in \Delta_n$. The map φ defines the group \mathcal{G}_n as the semidirect product $\mathcal{G}_n = \Delta_n \rtimes_\varphi \mathcal{P}_n$ (see [17]). For convenience, we sometimes denote element DP of \mathcal{G}_n as a pair (D, P) .

The groups \mathcal{G}_n , \mathcal{P}_n and Δ_n act on the left and the right of \mathbb{R}^n and $\mathbb{R}^{n \times m}$ in a canonical way (by matrix-vector or matrix-matrix multiplications). More precisely, we have:

Proposition 3.2. *Let $\mathbf{x} \in \mathbb{R}^n$ and $A = (A_{ij}) \in \mathbb{R}^{n \times m}$. Then for $D = \text{diag}(d_1, \dots, d_n) \in \Delta_n$, $\bar{D} = \text{diag}(\bar{d}_1, \dots, \bar{d}_m) \in \Delta_m$, $P_\pi \in \mathcal{P}_n$, and $P_\sigma \in \mathcal{P}_m$, we have:*

$$\begin{aligned} P_\pi \cdot \mathbf{x} &= (x_{\pi^{-1}(1)}, x_{\pi^{-1}(2)}, \dots, x_{\pi^{-1}(n)})^\top, \\ D \cdot \mathbf{x} &= (d_1 \cdot x_1, d_2 \cdot x_2, \dots, d_n \cdot x_n)^\top, \\ (D \cdot P_\pi \cdot A \cdot P_\sigma \cdot \bar{D})_{ij} &= d_i \cdot A_{\pi^{-1}(i)\sigma(j)} \cdot \bar{d}_j, \\ (D \cdot P_\pi \cdot A \cdot (\bar{D} \cdot P_\sigma)^{-1})_{ij} &= d_i \cdot A_{\pi^{-1}(i)\sigma^{-1}(j)} \cdot \bar{d}_j^{-1}. \end{aligned}$$

The above proposition can be verified by a direct computation, and is used in subsequent sections.

3.2 Monomial Matrices Preserved by a Nonlinear Activation

We characterize the maximal matrix groups preserved by the activations $\sigma = \text{ReLU}, \sin$ or \tanh . Here, ReLU is the rectified linear unit activation function which has been used in most of modern neural networks, sin is the sine function which is often used as an activation function in implicit neural representations [57], and tanh is the hyperbolic tangent activation function. Different variants of the results in this subsection can also be found in [24, 68]. We refine them using the terms of monomial matrices and state explicitly here for the completeness of the paper.

Definition 3.3. A matrix $A \in \text{GL}(n)$ is said to be *preserved by an activation* σ if and only if $\sigma(A \cdot \mathbf{x}) = A \cdot \sigma(\mathbf{x})$ for all $\mathbf{x} \in \mathbb{R}^n$.

We adopt the term *matrix group preserved by an activation* from [68]. This term is then referred to as the *intertwiner group of an activation* in [24].

Proposition 3.4. *For every matrix $A \in \text{GL}(n)$, we have:*

- (i) *A is preserved by the activation ReLU if and only if $A \in \mathcal{G}_n^{>0}$. Here, $\mathcal{G}_n^{>0}$ is the subgroup of \mathcal{G}_n containing only monomial matrices whose nonzero entries are positive numbers.*
- (ii) *A is preserved by the activation $\sigma = \sin$ or \tanh if and only if $A \in \mathcal{G}_n^{\pm 1}$. Here, $\mathcal{G}_n^{\pm 1}$ is the subgroup of \mathcal{G}_n containing only monomial matrices whose nonzero entries are ± 1 .*

A detailed proof of Proposition 3.4 can be found in Appendix C.1. As a consequence of the above theorem, $\mathcal{G}_n^{>0}$ (respectively, $\mathcal{G}_n^{\pm 1}$) is the maximal matrix subgroup of the general linear group $\text{GL}(n)$ that is preserved by the activation ReLU (respectively, sin and tanh).

Remark 3.5. *Intuitively, $\mathcal{G}_n^{>0}$ is generated by permuting and positive scaling the coordinates of vectors in \mathbb{R}^n , while $\mathcal{G}_n^{\pm 1}$ is generated by permuting and sign flipping. Formally, these groups can be written as the semidirect products:*

$$\mathcal{G}_n^{>0} = \Delta_n^{>0} \rtimes_\varphi \mathcal{P}_n, \quad \text{and} \quad \mathcal{G}_n^{\pm 1} = \Delta_n^{\pm 1} \rtimes_\varphi \mathcal{P}_n,$$

where

$$\Delta_n^{>0} = \{D = \text{diag}(d_1, \dots, d_n) : d_i > 0\}, \quad \text{and} \quad (3)$$

$$\Delta_n^{\pm 1} = \{D = \text{diag}(d_1, \dots, d_n) : d_i \in \{-1, 1\}\} \quad (4)$$

are two subgroups of Δ_n .

4 Weight Spaces and Monomial Matrix Group Actions on Weight Spaces

In this section, we formulate the general structure of the weight spaces of FCNNs and CNNs. We then determine the group action on these weight spaces using monomial matrices. The activation function σ using on the considered FCNNs and CNNs are assumed to be ReLU or sin or Tanh.

4.1 Weight Spaces of FCNNs and CNNs

Weight Spaces of FCNNs. Consider an FCNN with L layers, n_i neurons at the i -th layer, and n_0 and n_L be the input and output dimensions, together with the activation σ , as follows:

$$f(\mathbf{x}; U, \sigma) = W^{(L)} \cdot \sigma \left(\dots \sigma \left(W^{(2)} \cdot \sigma \left(W^{(1)} \cdot \mathbf{x} + b^{(1)} \right) + b^{(2)} \right) \dots \right) + b^{(L)}. \quad (5)$$

Here, $U = (W, b)$ is the parameters with the weights $W = \{W^{(i)} \in \mathbb{R}^{n_i \times n_{i-1}}\}_{i=1}^L$ and the biases $b = \{b^{(i)} \in \mathbb{R}^{n_i \times 1}\}_{i=1}^L$. The pair $U = (W, b)$ belongs to the weight space $\mathcal{U} = \mathcal{W} \times \mathcal{B}$, where:

$$\mathcal{W} = \mathbb{R}^{n_L \times n_{L-1}} \times \dots \times \mathbb{R}^{n_2 \times n_1} \times \mathbb{R}^{n_1 \times n_0}, \quad (6)$$

$$\mathcal{B} = \mathbb{R}^{n_L \times 1} \times \dots \times \mathbb{R}^{n_2 \times 1} \times \mathbb{R}^{n_1 \times 1}. \quad (7)$$

Weight Spaces of CNNs. Consider a CNN with L convolutional layers, ending with an average pooling layer then fully connected layers, together with activation σ . Let n_i and w_i be the number of channels and the size of the convolutional kernel at the i^{th} convolutional layer. We will only take account of the L convolutional layers, since the weight space of the fully connected layers are already considered above, and the pooling layer has no learnable parameters:

$$f(\mathbf{x}; U, \sigma) = \sigma \left(W^{(L)} * \sigma \left(\dots \sigma \left(W^{(2)} * \sigma \left(W^{(1)} * \mathbf{x} + b^{(1)} \right) + b^{(2)} \right) \dots \right) + b^{(L)} \right) \quad (8)$$

Here, $U = (W, b)$ is the learnable parameters with the weights $W = \{W^{(i)} \in \mathbb{R}^{w_i \times n_i \times n_{i-1}}\}_{i=1}^L$ and the biases $b = \{b^{(i)} \in \mathbb{R}^{1 \times n_i \times 1}\}_{i=1}^L$. The convolutional operator $*$ is defined depending on the purpose of the model, and adding b means adding $b_j^{(i)}$ to all entries of j -th channel at i^{th} layer. The pair $U = (W, b)$ belongs to the weight space $\mathcal{U} = \mathcal{W} \times \mathcal{B}$, where:

$$\mathcal{W} = \mathbb{R}^{w_L \times n_L \times n_{L-1}} \times \dots \times \mathbb{R}^{w_2 \times n_2 \times n_1} \times \mathbb{R}^{w_1 \times n_1 \times n_0}, \quad (9)$$

$$\mathcal{B} = \mathbb{R}^{1 \times n_L \times 1} \times \dots \times \mathbb{R}^{1 \times n_2 \times 1} \times \mathbb{R}^{1 \times n_1 \times 1}. \quad (10)$$

Remark 4.1. See in Appendix. C.2 for concrete descriptions of weight spaces of FCNNs and CNNs.

4.2 Monomial Matrix Group Action on Weight Spaces

The **weight space** \mathcal{U} of an FCNN or CNN with L layers and n_i channels at i^{th} layer has the general form $\mathcal{U} = \mathcal{W} \times \mathcal{B}$, where:

$$\mathcal{W} = \mathbb{R}^{w_L \times n_L \times n_{L-1}} \times \dots \times \mathbb{R}^{w_2 \times n_2 \times n_1} \times \mathbb{R}^{w_1 \times n_1 \times n_0}, \quad (11)$$

$$\mathcal{B} = \mathbb{R}^{b_L \times n_L \times 1} \times \dots \times \mathbb{R}^{b_2 \times n_2 \times 1} \times \mathbb{R}^{b_1 \times n_1 \times 1}. \quad (12)$$

Here, n_i is the number of channels at the i^{th} layer, in particular, n_0 and n_L are the number of channels of input and output; w_i is the dimension of weights and b_i is the dimension of the biases in each channel at the i -th layer. The dimension of the weight space \mathcal{U} is:

$$\dim \mathcal{U} = \sum_{i=1}^L (w_i \times n_i \times n_{i-1} + b_i \times n_i \times 1). \quad (13)$$

Notation. When working with weight matrices in \mathcal{W} , the space $\mathbb{R}^{w_i \times n_i \times n_{i-1}} = (\mathbb{R}^{w_i})^{n_i \times n_{i-1}}$ at the i^{th} layer will be considered as the space of $n_i \times n_{i-1}$ matrices, whose entries are real vectors in \mathbb{R}^{w_i} . In particular, the symbol $W^{(i)}$ denotes a matrix in $\mathbb{R}^{w_i \times n_i \times n_{i-1}} = (\mathbb{R}^{w_i})^{n_i \times n_{i-1}}$, while $W_{jk}^{(i)} \in \mathbb{R}^{w_i}$ denotes the entry at row j and column k of $W^{(i)}$. Similarly, the notion $b^{(i)}$ denotes a bias column vector in $\mathbb{R}^{b_i \times n_i \times 1} = (\mathbb{R}^{b_i})^{n_i \times 1}$, while $b_j^{(i)} \in \mathbb{R}^{b_i}$ denotes the entry at row j of $b^{(i)}$.

To define the **group action** of \mathcal{U} using monomial matrices, denote $\mathcal{G}_{\mathcal{U}}$ as the group:

$$\mathcal{G}_{\mathcal{U}} := \mathcal{G}_{n_L} \times \dots \times \mathcal{G}_{n_1} \times \mathcal{G}_{n_0}.$$

Ideally, each monomial matrix group \mathcal{G}_{n_i} will act on the weights and the biases at the i^{th} layer of the network. Each element of $\mathcal{G}_{\mathcal{U}}$ will be of the form $g = (g^{(L)}, \dots, g^{(0)})$, where:

$$g^{(i)} = D^{(i)} \cdot P_{\pi_i} = \text{diag} \left(d_1^{(i)}, \dots, d_{n_i}^{(i)} \right) \cdot P_{\pi_i} \in \mathcal{G}_{n_i} \quad (14)$$

for some invertible diagonal matrix $D^{(i)}$ and permutation matrix P_{π_i} . The action of $\mathcal{G}_{\mathcal{U}}$ on \mathcal{U} is defined formally as follows.

Definition 4.2 (Group action on weight spaces). With the notation as above, the *group action* of $\mathcal{G}_{\mathcal{U}}$ on \mathcal{U} is defined to be the map $\mathcal{G}_{\mathcal{U}} \times \mathcal{U} \rightarrow \mathcal{U}$ with $(g, U) \mapsto gU = (gW, gb)$, where:

$$(gW)^{(i)} := \left(g^{(i)} \right) \cdot W^{(i)} \cdot \left(g^{(i-1)} \right)^{-1} \quad \text{and} \quad (gb)^{(i)} := \left(g^{(i)} \right) \cdot b^{(i)}. \quad (15)$$

In concrete:

$$(gW)_{jk}^{(i)} := \frac{d_j^{(i)}}{d_k^{(i-1)}} \cdot W_{\pi_i^{-1}(j)\pi_{i-1}^{-1}(k)}^{(i)} \quad \text{and} \quad (gb)_j^{(i)} := d_j^{(i)} \cdot b_{\pi_i^{-1}(j)}^{(i)}. \quad (16)$$

Remark 4.3. The group $\mathcal{G}_{\mathcal{U}}$ is determined only by the number of layers L and the numbers of channels n_i , not by the dimensions of weights w_i and biases b_i at each channel.

The group $\mathcal{G}_{\mathcal{U}}$ has nice behaviors when acting on the weight spaces of FCNNs given in Eq. (5) and CNNs given in Eq. (8). In particular, depending on the specific choice of the activation σ , the function f built by the given FCNN or CNN is invariant under the action of a subgroup G of $\mathcal{G}_{\mathcal{U}}$, as we will see in the following proposition.

Proposition 4.4 (G -equivariance of neural functionals). Let $f = f(\cdot; U, \sigma)$ be an FCNN given in Eq. (5) or CNN given in Eq. (8) with the weight space $U \in \mathcal{U}$ and an activation $\sigma \in \{\text{ReLU}, \text{Tanh}, \sin\}$. Let us defined a subgroup G of $\mathcal{G}_{\mathcal{U}}$ as follows:

- (i) If $\sigma = \text{ReLU}$, we set $G = \{\text{id}_{\mathcal{G}_{n_L}}\} \times \mathcal{G}_{n_{L-1}}^{>0} \times \dots \times \mathcal{G}_{n_1}^{>0} \times \{\text{id}_{\mathcal{G}_{n_0}}\}$.
- (ii) If $\sigma = \sin$ or \tanh , then we set $G = \{\text{id}_{\mathcal{G}_{n_L}}\} \times \mathcal{G}_{n_{L-1}}^{\pm 1} \times \dots \times \mathcal{G}_{n_1}^{\pm 1} \times \{\text{id}_{\mathcal{G}_{n_0}}\}$.

Then f is G -invariant under the action of G on its weight space, i.e.

$$f(\mathbf{x}; U, \sigma) = f(\mathbf{x}; gU, \sigma) \quad (17)$$

for all $g \in G$, $U \in \mathcal{U}$ and $\mathbf{x} \in \mathbb{R}^{n_0}$.

Remark 4.5 (Maximality of G). The proof of Proposition 4.4 can be found in Appendix C.2. The group G defined above is even proved to be the maximal choice in the case:

- $\sigma = \text{ReLU}$ and $n_L \geq \dots \geq n_2 \geq n_1 > n_0 = 1$ (see [12, 25]), or
- $\sigma = \tanh$ (see [14, 22]).

Here, G is maximal in the sense that: if U' is another element in \mathcal{U} with $f(\cdot; U, \sigma) = f(\cdot; U', \sigma)$, then there exists an element $g \in G$ such that $U' = gU$. It is natural to ask whether the group G is still maximal in the other case. This question still remains open and we leave it for future exploration.

According to Proposition 4.4, the symmetries of the weight space of an FCNN or CNN must include not only permutation matrices but also other types of monomial matrices resulting from scaling (for ReLU networks) or sign flipping (for sin and tanh networks) the weights. Recent works on NFN design only take into account the permutation symmetries of the weight space. Therefore, it is necessary to design a new class of NFNs that incorporates these missing symmetries. We will introduce such a class in the next section.

5 Monomial Matrix Group Equivariant and Invariant NFNs

In this section, we introduce a new family of NFNs, called Monomial-NFNs, by incorporating symmetries arising from monomial matrix groups which have been clarified in Proposition 4.4. The main components of Monomial-NFNs are the monomial matrix group equivariant and invariant linear layers between two weight spaces which will be presented in Subsections 5.1 and 5.2, respectively. We will only consider the case of ReLU activation. Network architectures with other activations will be considered in detail in Appendices A and B.

In the following, $\mathcal{U} = (\mathcal{W}, \mathcal{B})$ is the weight space with L layers, n_i channels at i^{th} layer, and the dimensions of weight and bias are w_i and b_i , respectively (see Eqs. (11) and (12)). Since we consider ReLU network architectures, according to Proposition 4.4, the symmetries of the weight space is given by the subgroup $G = \{\text{id}_{\mathcal{G}_{n_L}}\} \times \mathcal{G}_{n_{L-1}}^{>0} \times \dots \times \mathcal{G}_{n_1}^{>0} \times \{\text{id}_{\mathcal{G}_{n_0}}\}$ of $\mathcal{G}_{\mathcal{U}}$.

Table 1: Number of parameters in a linear equivariant layer $E: \mathcal{U} \rightarrow \mathcal{U}'$ with respect to permutation matrix groups in [71], and monomial matrix groups. Here, $c = \max\{w_i, b_j\}$ and $c' = \max\{w'_i, b'_j\}$.

Subgroups of $\mathcal{G}_{\mathcal{U}}$		Number of parameters of E
$\mathcal{P}_{n_L} \times \mathcal{P}_{n_{L-1}} \times \dots \times \mathcal{P}_{n_1} \times \mathcal{P}_{n_0}$	([71])	$\mathcal{O}(cc'L^2)$
$\{\text{id}_{\mathcal{G}_{n_L}}\} \times \mathcal{P}_{n_{L-1}} \times \dots \times \mathcal{P}_{n_1} \times \{\text{id}_{\mathcal{G}_{n_0}}\}$	([71])	$\mathcal{O}(cc'(L + n_0 + n_L)^2)$
$\{\text{id}_{\mathcal{G}_{n_L}}\} \times \mathcal{G}_{n_{L-1}}^{>0} \times \dots \times \mathcal{G}_{n_1}^{>0} \times \{\text{id}_{\mathcal{G}_{n_0}}\}$	(Ours)	$\mathcal{O}(cc'(L + n_0 + n_L))$
$\{\text{id}_{\mathcal{G}_{n_L}}\} \times \mathcal{G}_{n_{L-1}}^{\pm 1} \times \dots \times \mathcal{G}_{n_1}^{\pm 1} \times \{\text{id}_{\mathcal{G}_{n_0}}\}$	(Ours)	$\mathcal{O}(cc'(L + n_0 + n_L))$

5.1 Equivariant Layers

We now construct a linear G -equivariant layer between weight spaces. These layers form the fundamental building blocks for our Monomimal-NFNs. Let \mathcal{U} and \mathcal{U}' be two weight spaces of the same network architecture described in Eqs. (11) and (12), i.e. they have the same number of layers as well as the same number of channels at each layer. Denote the dimension of weights and biases in each channel at the i -th layer of \mathcal{U}' as w'_i and b'_i , respectively. Note that, in this case, we have $\mathcal{G}_{\mathcal{U}} = \mathcal{G}_{\mathcal{U}'}$. We construct G -equivariant affine maps $E: \mathcal{U} \rightarrow \mathcal{U}'$ with $\mathbf{x} \mapsto \mathbf{a}\mathbf{x} + \mathbf{b}$, where $\mathbf{a} \in \mathbb{R}^{\dim \mathcal{U}' \times \dim \mathcal{U}}$ and $\mathbf{b} \in \mathbb{R}^{\dim \mathcal{U}' \times 1}$ are learnable parameters.

To make E to be G -equivariant, \mathbf{a} and \mathbf{b} have to satisfy a system of constraints (usually called *parameter sharing*), which are induced from the condition $E(gU) = gE(U)$ for all $g \in G$ and $U \in \mathcal{U}$. We show in details what are these constraints and how to derive the concrete formula of E in Appendix A. The formula of E is presented as follows: For $U = (W, b) \in \mathcal{U}$, the image $E(U) = (W', b') \in \mathcal{U}'$ is computed by:

$$\begin{aligned}
 W'_{jk}{}^{(1)} &= \sum_{q=1}^{n_0} \mathfrak{p}_{1jq}^{1jk} W_{jq}{}^{(1)} + \mathfrak{q}_{1j}^{1jk} b_j^{(1)}, & b_j^{(1)} &= \sum_{q=1}^{n_0} \mathfrak{r}_{1jq}^{1j} W_{jq}{}^{(1)} + \mathfrak{s}_{1j}^{1j} b_j^{(1)}, \\
 W'_{jk}{}^{(i)} &= \mathfrak{p}_{ijk}^{ijk} W_{jk}{}^{(i)}, & b_j^{(i)} &= \mathfrak{s}_{ij}^{ij} b_j^{(i)}, \quad 1 < i < L, \\
 W'_{jk}{}^{(L)} &= \sum_{p=1}^{n_L} \mathfrak{p}_{Lpk}^{Ljk} W_{pk}{}^{(L)}, & b_j^{(L)} &= \sum_{p=1}^{n_L} \mathfrak{s}_{Lp}^{Lj} b_p^{(L)} + \mathfrak{t}^{Lj}.
 \end{aligned} \tag{18}$$

Here, $(\mathfrak{p}, \mathfrak{q}, \mathfrak{r}, \mathfrak{s}, \mathfrak{t})$ is the hyperparameter of E . We discuss in detail the dimensions and sharing information between these parameters in Appendix A.1. Note that, we also show that all linear G -equivariant functional are in this form in Appendix A. To conclude, we have:

Theorem 5.1. *With notation as above, the linear functional map $E: \mathcal{U} \rightarrow \mathcal{U}'$ defined by Eq. (18) is G -equivariant. Moreover, every G -equivariant linear functional map from \mathcal{U} to \mathcal{U}' are in that form.*

Number of parameters and comparison previous works. The number of parameters in our layer is linear in L, n_0, n_L , which is significantly smaller than the number of parameters in layers described in [71], where it is quadratic in L, n_0, n_L (see Table 1). This reduction in parameter count means that our model is suitable for weight spaces of large-scale networks and deep NFNs. Intuitively, the advantage of our layer arises because the group G acting on the weight spaces in our setting is much larger, resulting in a significantly smaller number of orbits in the quotient space \mathcal{U}/G . Since the number of orbits is equal to the number of parameters, this leads to a more compact representation. Additionally, the presence of the group $\Delta_*^{>0}$ forces many coefficients of the linear layer E to be zero, further contributing to the efficiency of our model.

5.2 Invariant Layers

We will construct an G -invariant layer $I: \mathcal{U} \rightarrow \mathbb{R}^d$ for a fixed integer $d > 0$. In order to do that, we will seek a map I in the form:

$$I = \text{MLP} \circ I_{\mathcal{P}} \circ I_{\Delta^{>0}}, \tag{19}$$

where $I_{\Delta^{>0}}: \mathcal{U} \rightarrow \mathcal{U}$ is an $\Delta_*^{>0}$ -invariance and \mathcal{P}_* -equivariance map, $I_{\mathcal{P}}: \mathcal{U} \rightarrow \mathbb{R}^{\dim \mathcal{U}}$ is an \mathcal{P}_* -invariant map, and $\text{MLP}: \mathbb{R}^{\dim \mathcal{U}} \rightarrow \mathbb{R}^d$ is an arbitrary multilayer perceptron to adjust the output dimension. Since $G = \mathcal{G}_*^{>0} = \Delta_*^{>0} \rtimes_{\varphi} \mathcal{P}_*$ (see Remark 3.5), the composition $I = \text{MLP} \circ I_{\mathcal{P}} \circ I_{\Delta^{>0}}$ is clearly G -invariant as expected. The construction of $I_{\Delta^{>0}}$ and $I_{\mathcal{P}}$ will be presented below.

Table 2: CNN prediction on Tanh subset of Small CNN Zoo with original and augmented data.

	STATNN	NP	HNP	Monomial-NFN (ours)	Gap
Original	0.913 \pm 0.001	0.925 \pm 0.001	0.933 \pm 0.002	0.939 \pm 0.001	0.006
Augmented	0.914 \pm 0.001	0.928 \pm 0.001	0.935 \pm 0.001	0.943 \pm 0.001	0.008

Construct $I_{\Delta>0}$. To capture $\Delta_*^{>0}$ -invariance, we recall the notion of positively homogeneous of degree zero maps. For $n > 0$, a map α from \mathbb{R}^n is called *positively homogeneous of degree zero* if for all $\lambda > 0$ and $(x_1, \dots, x_n) \in \mathbb{R}^n$, we have $\alpha(\lambda x_1, \dots, \lambda x_n) = \alpha(x_1, \dots, x_n)$. We construct $I_{\Delta>0}: \mathcal{U} \rightarrow \mathcal{U}$ by taking collections of positively homogeneous of degree zero functions $\{\alpha_{jk}^{(i)}: \mathbb{R}^{w_i} \rightarrow \mathbb{R}^{w_i}\}$ and $\{\alpha_j^{(i)}: \mathbb{R}^{b_i} \rightarrow \mathbb{R}^{b_i}\}$, each one corresponds to weight and bias of \mathcal{U} . The maps $I_{\Delta>0}: \mathcal{U} \rightarrow \mathcal{U}$ that $(W, b) \mapsto (W', b')$ is defined by simply applying these functions on each weight and bias entries as follows:

$$W'_{jk}{}^{(i)} = \alpha_{jk}^{(i)}(W_{jk}^{(i)}) \quad \text{and} \quad b'_j{}^{(i)} = \alpha_j^{(i)}(b_j^{(i)}). \quad (20)$$

$I_{\Delta>0}$ is $\Delta_*^{>0}$ -invariant by homogeneity of the α functions. To make it become \mathcal{P}_* -equivariant, some α functions have to be shared across any axis that have permutation symmetry. We derive this relation in Appendix B. Some candidates for positively homogeneous of degree zero functions are also presented in Appendix B. They can be fixed or learnable.

Construct $I_{\mathcal{P}}$. To capture \mathcal{P}_* -invariance, we simply take summing or averaging the weight and bias across any axis that have permutation symmetry as in [71]. In concrete, we have $I_{\mathcal{P}}: \mathcal{U} \rightarrow \mathbb{R}^{\dim \mathcal{U}}$ is computed as follows:

$$I_{\mathcal{P}}(U) = \left(W_{\star, \star}^{(1)}, W_{\star, \star}^{(L)}, W_{\star, \star}^{(2)}, \dots, W_{\star, \star}^{(L-1)}; v_{\star}^{(L)}, v_{\star}^{(1)}, \dots, v_{\star}^{(L-1)} \right). \quad (21)$$

Here, \star denotes summation or averaging over the rows or columns of the weight and bias.

Remark 5.2. In our experiments, we use averaging operator since it is empirically more stable.

Finally we compose an MLP before $I_{\mathcal{P}}$ and $I_{\Delta>0}$ to obtain an G -invariant map. We summarize the above construction as follows.

Theorem 5.3. The functional map $I: \mathcal{U} \rightarrow \mathbb{R}^d$ defined by Eq. (19) is G -invariant.

5.3 Monomial Matrix Group Equivariant Neural Functionals (Monomial-NFNs)

We build Monomial-NFNs by the constructed equivariant and invariant functional layers, with activations and additional layers discussed below. The equivariant NFN is built by simply stacking G -equivariant layers. For the invariant counterpart, we follow the construction in [71]. In particular, we first stack some G -equivariant layers, then a $\Delta_*^{>0}$ -invariant and \mathcal{P}_* -equivariant layer. This makes our NFN to be $\Delta_*^{>0}$ -invariant and \mathcal{P}_* -equivariant. Then we finish the construction by stacking a \mathcal{P}_* -invariant layer and the end. This process makes the whole NFN to be G -invariant as expected.

Activations of G -equivariant functionals. Dealing with equivariance under action of \mathcal{P}_* only requires activation of the NFN is enough, since \mathcal{P}_* acts on only the order of channels in each channel of the weight space. For our G -equivariant NFNs, between each layer that is $\Delta_*^{>0}$ -equivariant, we have to use the same type of activations as the activation in the network input (i.e. either ReLU, sin or tanh in our consideration) to maintain the equivariance of the NFN.

Fourier Features and Positional Embedding. As mentioned in [35, 71, 72], Fourier Features [30, 62] and (sinusoidal) position embedding play a significant role in the performance of their functionals. Also, in [71], position embedding breaks the symmetry at input and output neurons, and allows us to use equivariant layers that act on input and output neurons. In our G -equivariant layers, we do not consider action on input and output neurons as mentioned. Also, using Fourier Features does not maintain $\Delta_*^{>0}$, so we can not use this Fourier layer for our equivariant Monomial-NFNs, and in our invariant Monomial-NFNs, we only can use Fourier layer after the $\Delta_*^{>0}$ -invariant layer. This can be considered as a limitation of Monomial-NFNs.

6 Experimental Results

In this session, we empirically demonstrate the performance of our Monomial Matrix Group Equivariant Neural Functional Networks (Monomial-NFNs) on various tasks that are either invariant




Figure 1: CNN prediction on ReLU subset of Small CNN Zoo with different ranges of augmentations. Here the x-axis is the augment upper scale, presented in log scale. The metric used is Kendall’s τ .

Table 3: Classification train and test accuracies (%) for implicit neural representations of MNIST, FashionMNIST, and CIFAR-10. Uncertainties indicate standard error over 5 runs.

	Monomial-NFN (ours)	NP	HNP	MLP
CIFAR-10	34.23 ± 0.33	33.74 ± 0.26	31.61 ± 0.22	10.48 ± 0.74
MNIST-10	<u>68.43 ± 0.51</u>	69.82 ± 0.42	66.02 ± 0.51	10.62 ± 0.54
FashionMNIST	61.15 ± 0.55	<u>58.21 ± 0.31</u>	57.43 ± 0.46	9.95 ± 0.36

(predicting CNN generalization from weights and classifying INR representations of images) or equivariant (weight space style editing). We aim to establish two key points. First, our model exhibits more stable behavior when the input undergoes transformations from the monomial matrix groups. Second, our model, Monomial-NFN, achieves competitive performance compared to other baseline models. Our results are averaged over 5 runs. Hyperparameter settings and the number of parameters can be found in Appendix D.

6.1 Predicting CNN Generalization from Weights

Experiment Setup. In this experiment, we evaluate how our Monomial-NFN predicts the generalization of pretrained CNN networks. We employ the Small CNN Zoo [64], which consists of multiple network weights trained with different initialization and hyperparameter settings, together with activations Tanh or ReLU. Since Monomial-NFNs depend on activations of network inputs, we divide the Small CNN Zoo into two smaller datasets based on their activations. The ReLU dataset considers the group $\mathcal{G}_n^{>0}$, while the Tanh dataset considers the group $\mathcal{G}_n^{\pm 1}$.

We construct the dataset with additional weights that undergo random hidden vector permutation and scaling based on their monomial matrix group. For the ReLU dataset with the group $\mathcal{G}_n^{>0}$, we uniformly sample the diagonal indices of D (see Eq. 14) for various ranges: $[1, 10]$, $[1, 1 \times 10^2]$, \dots , $[1, 1 \times 10^6]$, while belonging to $\{-1, 1\}$ in the case of Tanh dataset with the group $\mathcal{G}_n^{\pm 1}$. For both datasets, we compare our model with STATNN [65], and with two permutation equivariant neural functional networks from [71], referred to as HNP and NP. To compare the performance of all models, we use Kendall’s τ rank correlation metric [33].

Results. We demonstrate the results of all models on the ReLU subset in Figure 1, showing that our model attains stable Kendall’s τ when the scale operators are sampled from different ranges. Specifically, when the log of augmentation upper scale is 0, i.e. the data remains unaltered, our model performs as well as the HNP model. However, as the weights undergo more extensive scaling and permutation, the performance of the HNP and STATNN models drops significantly, indicating their lack of scaling symmetry. The NP model exhibits a similar trend, albeit to a lesser extent. In contrast, our model maintains stable performance throughout.

Table 2 illustrates the performance of all models on both the original and augmented Tanh subsets of CNN Zoo. Our model achieves the highest performance among all models and shows the greatest improvement after training with the augmented dataset. The gap between our model and the second-best model (HNP) increases from 0.006 to 0.008. Additionally, in both experiments, our model utilizes significantly fewer parameters than the baseline models, using only up to 50% of the parameters compared to HNP.

6.2 Classifying implicit neural representations of images

Experiment Setup. In this experiment, our focus is on extracting the original data information encoded within the weights of implicit neural representations (INRs). We utilize the dataset from [71],

Table 4: Test mean squared error (lower is better) between weight-space editing methods and ground-truth image-space transformations. Uncertainties indicate standard error over 5 runs.

	Monomial-NFN (ours)	NP	HNP	MLP
Contrast (CIFAR-10)	0.020 ± 0.001	0.020 ± 0.002	0.021 ± 0.002	0.031 ± 0.001
Dilate (MNIST)	<u>0.069 ± 0.002</u>	0.068 ± 0.002	0.071 ± 0.001	0.306 ± 0.001

which comprises pretrained INR networks [58] that encode images from the CIFAR-10 [36], MNIST [39], and FashionMNIST [69] datasets. Each pretrained INR network is designed to map image coordinates (x, y) to color pixel values - 3-dimensional RGB values for CIFAR-10 and 1-dimensional grayscale values for MNIST and FashionMNIST.

Results. We compare our model with NP, HNP, and MLP baselines. The results in Table 3 demonstrate that our model outperforms the second-best baseline, NP, for the FashionMNIST and CIFAR-10 datasets by 2.94% and 0.49%, respectively. For the MNIST dataset, our model also obtains comparable performance.

6.3 Weight space style editing.

Experiment setup. In this experiment, we explore altering the weights of the pretrained SIREN model [58] to change the information encoded within the network. We use the network weights provided in the HNP paper for the pretrained SIREN networks on MNIST and CIFAR-10 images. Our focus is on two tasks: the first involves modifying the network to dilate digits from the MNIST dataset, and the second involves altering the SIREN network weights to enhance the contrast of CIFAR-10 images. The objective is to minimize the mean squared error (MSE) training loss between the generated image from the edited SIREN network and the dilated/enhanced contrast image.

Results. Table 4 shows that our model performs on par with the best-performing model for increasing the contrast of CIFAR-10 images. For the MNIST digit dilation task, our model also achieves competitive performance compared to the NP baseline. Additionally, Figure 2 presents random samples of the digits that each model encodes for the dilation and contrast tasks, demonstrating that our model’s results are visually comparable to those of HNP and NP in both tasks.

7 Conclusion

In this paper, we formally describe a group of monomial matrices that preserves FCNNs and CNNs while acting on their weight spaces. For ReLU networks, this group includes permutation and scaling symmetries, while for networks with sin or Tanh activations, it encompasses permutation and sign-flipping symmetries. We introduce Monomial-NFNs, a first-of-a-kind class of NFNs that incorporates these scaling or sign-flipping symmetries in weight spaces. We demonstrate that the low number of trainable parameters in our equivariant linear layer of Monomial-NFNs compared to previous works on NFNs, highlighting their capability to efficiently process weight spaces of deep networks. Our NFNs exhibit competitive generalization performance and efficiency compared to existing models across several benchmarks.

One limitation of our model is that, due to the large size of the group considered, the resulting linear layers can be limited in terms of expressivity. For example, a weight corresponding to an edge between two neurons will be updated based only on its previous value, ignoring other edges across the same or other layers. To resolve this issue, it is necessary to construct an equivariant nonlinear layer to encode further relations between these weights, thus enhancing the expressivity. Another limitation is that we are uncertain about the maximality of the group G acting on the weight space of the ReLU network. Therefore, other types of symmetries may exist in the weight space beyond neuron permutation and weight scaling, and our model is not equivariant with respect to these symmetries. We leave the problem of identifying such a maximal group for future research.

Acknowledgments and Disclosure of Funding

This research / project is supported by the National Research Foundation Singapore under the AI Singapore Programme (AISG Award No: AISG2-TC-2023-012-SGIL). This research / project is supported by the Ministry of Education, Singapore, under the Academic Research Fund Tier 1 (FY2023) (A-8002040-00-00, A-8002039-00-00). This research / project is also supported by the NUS Presidential Young Professorship Award (A-0009807-01-00).

References

- [1] Francesca Albertini and Eduardo D Sontag. For neural networks, function determines form. *Neural networks*, 6(7):975–990, 1993.
- [2] Francesca Albertini and Eduardo D Sontag. Identifiability of discrete-time neural networks. In *Proc. European Control Conference*, pages 460–465. Springer Berlin, 1993.
- [3] Zeyuan Allen-Zhu, Yuanzhi Li, and Zhao Song. A convergence theory for deep learning via over-parameterization. In *International conference on machine learning*, pages 242–252. PMLR, 2019.
- [4] Bruno Andreis, Bedionita Soro, and Sung Ju Hwang. Set-based neural network encoding. *CoRR*, abs/2305.16625, 2023.
- [5] Marcin Andrychowicz, Misha Denil, Sergio Gomez, Matthew W Hoffman, David Pfau, Tom Schaul, Brendan Shillingford, and Nando De Freitas. Learning to learn by gradient descent by gradient descent. *Advances in neural information processing systems*, 29, 2016.
- [6] Maor Ashkenazi, Zohar Rimon, Ron Vainshtein, Shir Levi, Elad Richardson, Pinchas Mintz, and Eran Treister. Nern: Learning neural representations for neural networks. In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net, 2023.
- [7] Vijay Badrinarayanan, Bamdev Mishra, and Roberto Cipolla. Understanding symmetries in deep networks. *CoRR*, abs/1511.01029, 2015.
- [8] Bowen Baker, Otkrist Gupta, Ramesh Raskar, and Nikhil Naik. Accelerating neural architecture search using performance prediction. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Workshop Track Proceedings*. OpenReview.net, 2018.
- [9] Matthias Bauer, Emilien Dupont, Andy Brock, Dan Rosenbaum, Jonathan Schwarz, and Hyunjik Kim. Spatial functa: Scaling functa to imagenet classification and generation. *CoRR*, abs/2302.03130, 2023.
- [10] Mikhail Belkin, Daniel Hsu, Siyuan Ma, and Soumik Mandal. Reconciling modern machine-learning practice and the classical bias–variance trade-off. *Proceedings of the National Academy of Sciences*, 116(32):15849–15854, 2019.
- [11] Samy Bengio, Yoshua Bengio, Jocelyn Cloutier, and Jan Gecsei. On the optimization of a synaptic learning rule. In *Optimality in Biological and Artificial Networks?*, pages 265–287. Routledge, 2013.
- [12] Phuong Bui Thi Mai and Christoph Lampert. Functional vs. parametric equivalence of relu networks. In *8th International Conference on Learning Representations*, 2020.
- [13] Nicola De Cao, Wilker Aziz, and Ivan Titov. Editing factual knowledge in language models. In Marie-Francine Moens, Xuanjing Huang, Lucia Specia, and Scott Wen-tau Yih, editors, *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing, EMNLP 2021, Virtual Event / Punta Cana, Dominican Republic, 7-11 November, 2021*, pages 6491–6506. Association for Computational Linguistics, 2021.
- [14] An Mei Chen, Haw-minn Lu, and Robert Hecht-Nielsen. On the geometry of feedforward neural network error surfaces. *Neural computation*, 5(6):910–927, 1993.
- [15] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: pre-training of deep bidirectional transformers for language understanding. In Jill Burstein, Christy Doran, and Thamar Solorio, editors, *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, pages 4171–4186. Association for Computational Linguistics, 2019.

- [16] Simon Du, Jason Lee, Haochuan Li, Liwei Wang, and Xiyu Zhai. Gradient descent finds global minima of deep neural networks. In *International conference on machine learning*, pages 1675–1685. PMLR, 2019.
- [17] David Steven Dummit and Richard M Foote. *Abstract algebra*, volume 3. Wiley Hoboken, 2004.
- [18] Emilien Dupont, Hyunjik Kim, S. M. Ali Eslami, Danilo Jimenez Rezende, and Dan Rosenbaum. From data to functa: Your data point is a function and you can treat it like one. In Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvári, Gang Niu, and Sivan Sabato, editors, *International Conference on Machine Learning, ICML 2022, 17-23 July 2022, Baltimore, Maryland, USA*, volume 162 of *Proceedings of Machine Learning Research*, pages 5694–5725. PMLR, 2022.
- [19] Emilien Dupont, Yee Whye Teh, and Arnaud Doucet. Generative models as distributions of functions. In Gustau Camps-Valls, Francisco J. R. Ruiz, and Isabel Valera, editors, *International Conference on Artificial Intelligence and Statistics, AISTATS 2022, 28-30 March 2022, Virtual Event*, volume 151 of *Proceedings of Machine Learning Research*, pages 2989–3015. PMLR, 2022.
- [20] Gabriel Eilertsen, Daniel Jönsson, Timo Ropinski, Jonas Unger, and Anders Ynnerman. Classifying the classifier: Dissecting the weight space of neural networks. In Giuseppe De Giacomo, Alejandro Catalá, Bistra Dilkina, Michela Milano, Senén Barro, Alberto Bugarín, and Jérôme Lang, editors, *ECAI 2020 - 24th European Conference on Artificial Intelligence, 29 August-8 September 2020, Santiago de Compostela, Spain, August 29 - September 8, 2020 - Including 10th Conference on Prestigious Applications of Artificial Intelligence (PAIS 2020)*, volume 325 of *Frontiers in Artificial Intelligence and Applications*, pages 1119–1126. IOS Press, 2020.
- [21] Ziya Erkoç, Fangchang Ma, Qi Shan, Matthias Nießner, and Angela Dai. Hyperdiffusion: Generating implicit neural fields with weight-space diffusion. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 14300–14310, 2023.
- [22] Charles Fefferman and Scott Markel. Recovering a feed-forward net from its output. *Advances in neural information processing systems*, 6, 1993.
- [23] Jonathan Frankle and Michael Carbin. The lottery ticket hypothesis: Finding sparse, trainable neural networks. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net, 2019.
- [24] Charles Godfrey, Davis Brown, Tegan Emerson, and Henry Kvinge. On the symmetries of deep learning models and their internal representations. *Advances in Neural Information Processing Systems*, 35:11893–11905, 2022.
- [25] Elisenda Grigsby, Kathryn Lindsey, and David Rolnick. Hidden symmetries of relu networks. In *International Conference on Machine Learning*, pages 11734–11760. PMLR, 2023.
- [26] Jean Harb, Tom Schaul, Doina Precup, and Pierre-Luc Bacon. Policy evaluation networks. *arXiv preprint arXiv:2002.11833*, 2020.
- [27] Kaiming He, X. Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778, 2015.
- [28] Robert Hecht-Nielsen. On the algebraic structure of feedforward network weight spaces. In *Advanced Neural Computers*, pages 129–135. Elsevier, 1990.
- [29] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural Comput.*, 9(8):1735–1780, 1997.
- [30] Arthur Jacot, Franck Gabriel, and Clément Hongler. Neural tangent kernel: Convergence and generalization in neural networks. *Advances in neural information processing systems*, 31, 2018.

- [31] John M. Jumper, Richard O. Evans, Alexander Pritzel, Tim Green, Michael Figurnov, Olaf Ronneberger, Kathryn Tunyasuvunakool, Russell Bates, Augustin Židek, Anna Potapenko, Alex Bridgland, Clemens Meyer, Simon A. A. Kohl, Andrew J. Ballard, Andrew Cowie, Bernardino Romera-Paredes, Stanislav Nikolov, R. D. Jain, Jonas Adler, Trevor Back, Stig Petersen, David Reiman, Ellen Clancy, Michal Zielinski, Martin Steinegger, Michalina Pacholska, Tamas Berghammer, Sebastian Bodenstern, David L. Silver, Oriol Vinyals, Andrew W. Senior, Koray Kavukcuoglu, Pushmeet Kohli, and Demis Hassabis. Highly accurate protein structure prediction with alphafold. *Nature*, 596(7873):583–589, July 2021. <https://europepmc.org/article/MED/34265844> ; <https://www.nature.com/articles/s41586-021-03819-2> ; <https://www.mendeley.com/catalogue/bde88f33-525c-3af0-823a-3bb305a93020/> ; <https://facultyopinions.com/prime/740477161> ; <https://pubmed.ncbi.nlm.nih.gov/34265844/> ; <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC8371605/> ; <https://www.nature.com/articles/s41586-021-03819-2.pdf> ; <https://europepmc.org/abstract/MED/34265844> ; https://econpapers.repec.org/RePEc:nat:nature:v:596:y:2021:i:7873:d:10.1038_s41586-021-03819-2 ; <https://www.scienceopen.com/document?vid=b1f93136-f0f0-4c78-8e45-27bd037e9bb9> ; <http://rucweb.tsg211.com/http/77726476706e69737468656265737421e7e056d229317c456c0dc7af9758/articles/s41586-021-03819-2> ; <http://pubmed02.keyan123.cn/34265844/>.
- [32] Ioannis Kalogeropoulos, Giorgos Bouritsas, and Yannis Panagakis. Scale equivariant graph metanetworks. *arXiv preprint arXiv:2406.10685*, 2024.
- [33] M. G. KENDALL. A NEW MEASURE OF RANK CORRELATION. *Biometrika*, 30(1-2):81–93, 06 1938.
- [34] Boris Knyazev, Michal Drozdal, Graham W Taylor, and Adriana Romero Soriano. Parameter prediction for unseen deep architectures. *Advances in Neural Information Processing Systems*, 34:29433–29448, 2021.
- [35] Miltiadis Kofinas, Boris Knyazev, Yan Zhang, Yunlu Chen, Gertjan J. Burghouts, Efstratios Gavves, Cees G. M. Snoek, and David W. Zhang. Graph neural networks for learning equivariant representations of neural networks. In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*. OpenReview.net, 2024.
- [36] Alex Krizhevsky and Geoffrey Hinton. Learning multiple layers of features from tiny images. Technical Report 0, University of Toronto, Toronto, Ontario, 2009.
- [37] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E. Hinton. Imagenet classification with deep convolutional neural networks. In Peter L. Bartlett, Fernando C. N. Pereira, Christopher J. C. Burges, Léon Bottou, and Kilian Q. Weinberger, editors, *Advances in Neural Information Processing Systems 25: 26th Annual Conference on Neural Information Processing Systems 2012. Proceedings of a meeting held December 3-6, 2012, Lake Tahoe, Nevada, United States*, pages 1106–1114, 2012.
- [38] Vera Kurkova and Paul C Kainen. Functionally equivalent feedforward neural networks. *Neural Computation*, 6(3):543–558, 1994.
- [39] Yann LeCun and Corinna Cortes. The mnist database of handwritten digits. 2005.
- [40] Derek Lim, Haggai Maron, Marc T. Law, Jonathan Lorraine, and James Lucas. Graph metanetworks for processing diverse neural architectures. In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*. OpenReview.net, 2024.
- [41] Luca De Luigi, Adriano Cardace, Riccardo Spezialetti, Pierluigi Zama Ramirez, Samuele Salti, and Luigi Di Stefano. Deep learning on implicit neural representations of shapes. In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net, 2023.
- [42] Luke Metz, James Harrison, C Daniel Freeman, Amil Merchant, Lucas Beyer, James Bradbury, Naman Agrawal, Ben Poole, Igor Mordatch, Adam Roberts, et al. Velo: Training versatile learned optimizers by scaling up. *arXiv preprint arXiv:2211.09760*, 2022.

- [43] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM*, 65(1):99–106, 2021.
- [44] Eric Mitchell, Charles Lin, Antoine Bosselut, Chelsea Finn, and Christopher D. Manning. Fast model editing at scale. In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*. OpenReview.net, 2022.
- [45] Aviv Navon, Aviv Shamsian, Idan Achituve, Ethan Fetaya, Gal Chechik, and Haggai Maron. Equivariant architectures for learning in deep weight spaces. In *International Conference on Machine Learning*, pages 25790–25816. PMLR, 2023.
- [46] Behnam Neyshabur, Russ R Salakhutdinov, and Nati Srebro. Path-sgd: Path-normalized optimization in deep neural networks. *Advances in neural information processing systems*, 28, 2015.
- [47] Roman Novak, Yasaman Bahri, Daniel A. Abolafia, Jeffrey Pennington, and Jascha Sohl-Dickstein. Sensitivity and generalization in neural networks: an empirical study. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net, 2018.
- [48] William Peebles, Ilija Radosavovic, Tim Brooks, Alexei A Efros, and Jitendra Malik. Learning to learn with generative models of neural network checkpoints. *arXiv preprint arXiv:2209.12892*, 2022.
- [49] M. Raissi, P. Perdikaris, and G.E. Karniadakis. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*, 378:686–707, 2019.
- [50] Joseph J Rotman. *An introduction to the theory of groups*, volume 148. Springer Science & Business Media, 2012.
- [51] David E. Rumelhart, Geoffrey E. Hinton, and Ronald J. Williams. Learning internal representations by error propagation. 1986.
- [52] Thomas Philip Runarsson and Magnus Thor Jonsson. Evolution and design of distributed learning rules. In *2000 IEEE Symposium on Combinations of Evolutionary Computation and Neural Networks. Proceedings of the First IEEE Symposium on Combinations of Evolutionary Computation and Neural Networks (Cat. No. 00*, pages 59–63. IEEE, 2000.
- [53] Konstantin Schürholt, Boris Knyazev, Xavier Giró-i Nieto, and Damian Borth. Hyperrepresentations as generative models: Sampling unseen neural network weights. *Advances in Neural Information Processing Systems*, 35:27906–27920, 2022.
- [54] Konstantin Schürholt, Dimche Kostadinov, and Damian Borth. Self-supervised representation learning on neural network weights for model characteristic prediction. *Advances in Neural Information Processing Systems*, 34:16481–16493, 2021.
- [55] Konstantin Schürholt, Diyar Taskiran, Boris Knyazev, Xavier Giró-i Nieto, and Damian Borth. Model zoos: A dataset of diverse populations of neural network models. *Advances in Neural Information Processing Systems*, 35:38134–38148, 2022.
- [56] Anton Sinitsin, Vsevolod Plokhotnyuk, Dmitry V. Pyrkin, Sergei Popov, and Artem Babenko. Editable neural networks. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net, 2020.
- [57] Vincent Sitzmann, Julien Martel, Alexander Bergman, David Lindell, and Gordon Wetzstein. Implicit neural representations with periodic activation functions. *Advances in neural information processing systems*, 33:7462–7473, 2020.
- [58] Vincent Sitzmann, Julien N. P. Martel, Alexander W. Bergman, David B. Lindell, and Gordon Wetzstein. Implicit neural representations with periodic activation functions. In Hugo Larochelle, Marc’Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin, editors, *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020.

- [59] Samuel Sokota, Hengyuan Hu, David J Wu, J Zico Kolter, Jakob Nicolaus Foerster, and Noam Brown. A fine-tuning approach to belief state modeling. In *International Conference on Learning Representations*, 2021.
- [60] Kenneth O Stanley. Compositional pattern producing networks: A novel abstraction of development. *Genetic programming and evolvable machines*, 8:131–162, 2007.
- [61] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott E. Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. Going deeper with convolutions. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2015, Boston, MA, USA, June 7-12, 2015*, pages 1–9. IEEE Computer Society, 2015.
- [62] Matthew Tancik, Pratul Srinivasan, Ben Mildenhall, Sara Fridovich-Keil, Nithin Raghavan, Utkarsh Singhal, Ravi Ramamoorthi, Jonathan Barron, and Ren Ng. Fourier features let networks learn high frequency functions in low dimensional domains. *Advances in neural information processing systems*, 33:7537–7547, 2020.
- [63] Viet-Hoang Tran, Thieu N Vo, An Nguyen The, Tho Tran Huu, Minh-Khoi Nguyen-Nhat, Thanh Tran, Duy-Tung Pham, and Tan Minh Nguyen. Equivariant neural functional networks for transformers. *arXiv preprint arXiv:2410.04209*, 2024.
- [64] Thomas Unterthiner, Daniel Keysers, Sylvain Gelly, Olivier Bousquet, and Ilya Tolstikhin. Predicting neural network accuracy from weights. *arXiv preprint arXiv:2002.11448*, 2020.
- [65] Thomas Unterthiner, Daniel Keysers, Sylvain Gelly, Olivier Bousquet, and Ilya Tolstikhin. Predicting neural network accuracy from weights, 2021.
- [66] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In Isabelle Guyon, Ulrike von Luxburg, Samy Bengio, Hanna M. Wallach, Rob Fergus, S. V. N. Vishwanathan, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 5998–6008, 2017.
- [67] Thieu N Vo, Viet-Hoang Tran, Tho Tran Huu, An Nguyen The, Thanh Tran, Minh-Khoi Nguyen-Nhat, Duy-Tung Pham, and Tan Minh Nguyen. Equivariant polynomial functional networks. *arXiv preprint arXiv:2410.04213*, 2024.
- [68] Jeffrey Wood and John Shawe-Taylor. Representation theory and invariant neural networks. *Discrete applied mathematics*, 69(1-2):33–60, 1996.
- [69] Han Xiao, Kashif Rasul, and Roland Vollgraf. Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms. *CoRR*, abs/1708.07747, 2017.
- [70] Allan Zhou, Chelsea Finn, and James Harrison. Universal neural functionals. *CoRR*, abs/2402.05232, 2024.
- [71] Allan Zhou, Kaien Yang, Kaylee Burns, Adriano Cardace, Yiding Jiang, Samuel Sokota, J Zico Kolter, and Chelsea Finn. Permutation equivariant neural functionals. *Advances in Neural Information Processing Systems*, 36, 2024.
- [72] Allan Zhou, Kaien Yang, Yiding Jiang, Kaylee Burns, Winnie Xu, Samuel Sokota, J Zico Kolter, and Chelsea Finn. Neural functional transformers. *Advances in Neural Information Processing Systems*, 36, 2024.

Supplement to “Monomial Matrix Group Equivariant Neural Functional Networks”

Table of Contents

A Construction of Monomial Matrix Group Equivariant Layers	16
A.1 ReLU activation	19
A.2 Sin or Tanh activation	20
B Construction of Monomial Matrix Group Invariant Layers	21
B.1 ReLU activation	21
B.2 Sin or Tanh activation	22
C Proofs of Theoretical Results	23
C.1 Proof of Proposition 3.4	23
C.2 Proof of Proposition 4.4	24
D Additional experimental details	27
D.1 Runtime and Memory Consumption	27
D.2 Comparison of Monomial-NFNs and GNN-based NFNs	27
D.3 Predicting generalization from weights	28
D.4 Classifying implicit neural representations of images	29
D.5 Weight space style editing	30
D.6 Ablation Regarding Design Choices	31

A Construction of Monomial Matrix Group Equivariant Layers

In this appendix, we present how we constructed Monomial Matrix Group Equivariant Layers. We adopt the idea of notation in [71] to derive the formula of linear functional layers. For two weight spaces \mathcal{U} and \mathcal{U}' with the same number of layers L as well as the same number of channels at i -th layer n_i :

$$\begin{aligned} \mathcal{U} &= \mathcal{W} \times \mathcal{B} \text{ where:} & (22) \\ \mathcal{W} &= \mathbb{R}^{w_L \times n_L \times n_{L-1}} \times \dots \times \mathbb{R}^{w_2 \times n_2 \times n_1} \times \mathbb{R}^{w_1 \times n_1 \times n_0}, \\ \mathcal{B} &= \mathbb{R}^{b_L \times n_L \times 1} \times \dots \times \mathbb{R}^{b_2 \times n_2 \times 1} \times \mathbb{R}^{b_1 \times n_1 \times 1}, \end{aligned}$$

and

$$\begin{aligned} \mathcal{U}' &= \mathcal{W}' \times \mathcal{B}' \text{ where:} & (23) \\ \mathcal{W}' &= \mathbb{R}^{w'_L \times n_L \times n_{L-1}} \times \dots \times \mathbb{R}^{w'_2 \times n_2 \times n_1} \times \mathbb{R}^{w'_1 \times n_1 \times n_0}, \\ \mathcal{B}' &= \mathbb{R}^{b'_L \times n_L \times 1} \times \dots \times \mathbb{R}^{b'_2 \times n_2 \times 1} \times \mathbb{R}^{b'_1 \times n_1 \times 1}, \end{aligned}$$

our equivariant layer $E: \mathcal{U} \rightarrow \mathcal{U}'$ will has the form as follows:

$$E : (W, b) = U \mapsto U' = (W', b') \text{ where:} \quad (24)$$

$$W'_{jk}{}^{(i)} := \sum_{s=1}^L \sum_{p=1}^{n_s} \sum_{q=1}^{n_{s-1}} \mathfrak{p}_{spq}^{ijk} W_{pq}^{(s)} + \sum_{s=1}^L \sum_{p=1}^{n_s} \mathfrak{q}_{sp}^{ijk} b_p^{(s)} + \mathfrak{t}^{ijk} \quad (25)$$

$$b'_j{}^{(i)} := \sum_{s=1}^L \sum_{p=1}^{n_s} \sum_{q=1}^{n_{s-1}} \mathfrak{r}_{spq}^{ij} W_{pq}^{(s)} + \sum_{s=1}^L \sum_{p=1}^{n_s} \mathfrak{s}_{sp}^{ij} b_p^{(s)} + \mathfrak{t}^{ij} \quad (26)$$

Here, the map E is parameterized by hyperparameter $\theta = (\mathfrak{p}, \mathfrak{q}, \mathfrak{s}, \mathfrak{r}, \mathfrak{t})$ with dimensions of each component as follows:

- $\mathfrak{p}_{spq}^{ijk} \in \mathbb{R}^{w'_i \times w_s}$ represents the contribution of $W_{pq}^{(s)}$ to $W'_{jk}{}^{(i)}$,

- $\mathbf{q}_{sp}^{ijk} \in \mathbb{R}^{w_i' \times b_s}$ represents the contribution of $b_p^{(s)}$ to $W_{jk}'^{(i)}$,
- $\mathbf{t}^{ijk} \in \mathbb{R}^{w_i'}$ is the bias of the layer for $W_{jk}'^{(i)}$;
- $\mathbf{r}_{spq}^{ij} \in \mathbb{R}^{b_i' \times w_s}$ represents the contribution of $W_{pq}^{(s)}$ to $b_j'^{(i)}$,
- $\mathbf{s}_{sp}^{ij} \in \mathbb{R}^{b_i' \times b_s}$ represents the contribution of $b_p^{(s)}$ to $b_j'^{(i)}$,
- $\mathbf{t}^{ij} \in \mathbb{R}^{b_i'}$ is the bias of the layer for $b_j'^{(i)}$.

We want to see how an element of the group \mathcal{G}_U acts on input and output of layer E . Let

$$g = (g^{(L)}, \dots, g^{(0)}) \in \mathcal{G}_{n_L} \times \dots \times \mathcal{G}_{n_0} = \mathcal{G}_U, \quad (27)$$

where

$$g^{(i)} = D^{(i)} \cdot P_{\pi_i} = \text{diag}(d_1^{(i)}, \dots, d_{n_i}^{(i)}) \cdot P_{\pi_i} \in \mathcal{G}_{n_i}. \quad (28)$$

Recall the definition of the group action $gU = (gW, gb)$ where:

$$(gW)^{(i)} := (g^{(i)}) \cdot W^{(i)} \cdot (g^{(i-1)})^{-1} \quad \text{and} \quad (gb)^{(i)} := (g^{(i)}) \cdot b^{(i)}, \quad (29)$$

or in term of entries:

$$(gW)_{jk}^{(i)} := \frac{d_j^{(i)}}{d_k^{(i-1)}} \cdot W_{\pi_i^{-1}(j)\pi_{i-1}^{-1}(k)}^{(i)} \quad \text{and} \quad (gb)_j^{(i)} := d_j^{(i)} \cdot b_{\pi_i^{-1}(j)}^{(i)}. \quad (30)$$

$gE(U) = gU' = (gW', gb')$ is computed as follows:

$$(gW')_{jk}^{(i)} = \frac{d_j^{(i)}}{d_k^{(i-1)}} \cdot W_{\pi_i^{-1}(j)\pi_{i-1}^{-1}(k)}'^{(i)} \quad (31)$$

$$= \frac{d_j^{(i)}}{d_k^{(i-1)}} \cdot \left(\sum_{s=1}^L \sum_{p=1}^{n_s} \sum_{q=1}^{n_{s-1}} \mathbf{p}_{spq}^{i\pi_i^{-1}(j)\pi_{i-1}^{-1}(k)} W_{pq}^{(s)} + \right. \quad (32)$$

$$\left. \sum_{s=1}^L \sum_{p=1}^{n_s} \mathbf{q}_{sp}^{i\pi_i^{-1}(j)\pi_{i-1}^{-1}(k)} b_p^{(s)} + \mathbf{t}^{i\pi_i^{-1}(j)\pi_{i-1}^{-1}(k)} \right) \quad (33)$$

$$(gb')_j^{(i)} = d_j^{(i)} \cdot b_{\pi_i^{-1}(j)}'^{(i)} \quad (34)$$

$$= d_j^{(i)} \cdot \left(\sum_{s=1}^L \sum_{p=1}^{n_s} \sum_{q=1}^{n_{s-1}} \mathbf{s}_{spq}^{i\pi_i^{-1}(j)} W_{pq}^{(s)} + \right. \quad (35)$$

$$\left. \sum_{s=1}^L \sum_{p=1}^{n_s} \mathbf{r}_{sp}^{i\pi_i^{-1}(j)} b_p^{(s)} + \mathbf{t}^{i\pi_i^{-1}(j)} \right). \quad (36)$$

$E(gU) = (gU)' = ((gW)', (gb)')$ is computed as follows:

$$(gU)_{jk}^{(i)} = \sum_{s=1}^L \sum_{p=1}^{n_s} \sum_{q=1}^{n_{s-1}} \mathfrak{p}_{spq}^{ijk} \cdot \frac{d_p^{(s)}}{d_q^{(s-1)}} \cdot W_{\pi_s^{-1}(p)\pi_{s-1}^{-1}(q)}^{(s)} + \sum_{s=1}^L \sum_{p=1}^{n_s} \mathfrak{q}_{sp}^{ijk} \cdot d_p^{(s)} \cdot b_{\pi_s^{-1}(p)}^{(s)} + \mathfrak{t}^{ijk} \quad (37)$$

$$= \sum_{s=1}^L \sum_{p=1}^{n_s} \sum_{q=1}^{n_{s-1}} \mathfrak{p}_{s\pi_s(p)\pi_{s-1}(q)}^{ijk} \cdot \frac{d_{\pi_s(p)}^{(s)}}{d_{\pi_{s-1}(q)}^{(s-1)}} \cdot W_{pq}^{(s)} + \sum_{s=1}^L \sum_{p=1}^{n_s} \mathfrak{q}_{s\pi_s(p)}^{ijk} \cdot d_{\pi_s(p)}^{(s)} \cdot b_p^{(s)} + \mathfrak{t}^{ijk} \quad (38)$$

$$(gb)_j^{(i)} = \sum_{s=1}^L \sum_{p=1}^{n_s} \sum_{q=1}^{n_{s-1}} \mathfrak{r}_{spq}^{ij} \cdot \frac{d_p^{(s)}}{d_q^{(s-1)}} \cdot W_{\pi_s^{-1}(p)\pi_{s-1}^{-1}(q)}^{(s)} + \sum_{s=1}^L \sum_{p=1}^{n_s} \mathfrak{s}_{sp}^{ij} \cdot d_p^{(s)} \cdot b_{\pi_s^{-1}(p)}^{(s)} + \mathfrak{t}^{ij} \quad (39)$$

$$= \sum_{s=1}^L \sum_{p=1}^{n_s} \sum_{q=1}^{n_{s-1}} \mathfrak{r}_{s\pi_s(p)\pi_{s-1}(q)}^{ij} \cdot \frac{d_{\pi_s(p)}^{(s)}}{d_{\pi_{s-1}(q)}^{(s-1)}} \cdot W_{pq}^{(s)} + \sum_{s=1}^L \sum_{p=1}^{n_s} \mathfrak{s}_{s\pi_s(p)}^{ij} \cdot d_{\pi_s(p)}^{(s)} \cdot b_p^{(s)} + \mathfrak{t}^{ij}. \quad (40)$$

We need E is G -equivariant under the action of subgroups of \mathcal{G}_U as in Theorem 4.4. From the above computation, if $gE(U) = E(gU)$, the hyperparameter $\theta = (\mathfrak{p}, \mathfrak{q}, \mathfrak{r}, \mathfrak{s}, \mathfrak{t})$ have to satisfy the system of constraints as follows:

$$\frac{d_j^{(i)}}{d_k^{(i-1)}} \cdot \mathfrak{p}_{spq}^{i\pi_i^{-1}(j)\pi_{i-1}^{-1}(k)} = \mathfrak{p}_{s\pi_s(p)\pi_{s-1}(q)}^{ijk} \cdot \frac{d_{\pi_s(p)}^{(s)}}{d_{\pi_{s-1}(q)}^{(s-1)}} \quad (41)$$

$$\frac{d_j^{(i)}}{d_k^{(i-1)}} \cdot \mathfrak{q}_{sp}^{i\pi_i^{-1}(j)\pi_{i-1}^{-1}(k)} = \mathfrak{q}_{s\pi_s(p)} \cdot d_{\pi_s(p)}^{(s)} \quad (42)$$

$$d_j^{(i)} \cdot \mathfrak{r}_{spq}^{i\pi_i^{-1}(j)} = \mathfrak{r}_{s\pi_s(p)\pi_{s-1}(q)}^{ij} \cdot \frac{d_{\pi_s(p)}^{(s)}}{d_{\pi_{s-1}(q)}^{(s-1)}} \quad (43)$$

$$d_j^{(i)} \cdot \mathfrak{s}_{sp}^{i\pi_i^{-1}(j)} = \mathfrak{s}_{s\pi_s(p)}^{ij} \cdot d_{\pi_s(p)}^{(s)} \quad (44)$$

$$\frac{d_j^{(i)}}{d_k^{(i-1)}} \cdot \mathfrak{t}^{i\pi_i^{-1}(j)\pi_{i-1}^{-1}(k)} = \mathfrak{t}^{ijk} \quad (45)$$

$$d_j^{(i)} \cdot \mathfrak{t}^{i\pi_i^{-1}(j)} = \mathfrak{t}^{ij}. \quad (46)$$

for all possible tuples $((i, j, k), (s, p, q))$ and all $g \in G$. Since the two subgroups G considered in Theorem 4.4 satisfy that: $G \cap \mathcal{P}_i$ is trivial (for $i = 0$ or $i = L$) or the whole \mathcal{P}_i (for $0 < i < L$), so we can simplify the above system of constraints by moving all the permutation π 's to LHS, then replacing π^{-1} by π . The system, denoted as (*), now is written as follows:

$$\frac{d_j^{(i)}}{d_k^{(i-1)}} \cdot \mathfrak{p}^{i\pi_i(j)\pi_{i-1}(k)} = \mathfrak{p}_{spq}^{ijk} \cdot \frac{d_p^{(s)}}{d_q^{(s-1)}} \quad (*1)$$

$$\frac{d_j^{(i)}}{d_k^{(i-1)}} \cdot \mathfrak{q}^{i\pi_i(j)\pi_{i-1}(k)} = \mathfrak{q}_{sp}^{ijk} \cdot d_p^{(s)} \quad (*2)$$

$$d_j^{(i)} \cdot \mathfrak{r}^{i\pi_i(j)} = \mathfrak{r}_{spq}^{ij} \cdot \frac{d_p^{(s)}}{d_q^{(s-1)}} \quad (*3)$$

$$d_j^{(i)} \cdot \mathfrak{s}^{i\pi_i(j)} = \mathfrak{s}_{sp}^{ij} \cdot d_p^{(s)} \quad (*4)$$

$$\frac{d_j^{(i)}}{d_k^{(i-1)}} \cdot \mathfrak{t}^{i\pi_i^{-1}(j)\pi_{i-1}^{-1}(k)} = \mathfrak{t}^{ijk} \quad (*5)$$

$$d_j^{(i)} \cdot \mathfrak{t}^{i\pi_i^{-1}(j)} = \mathfrak{t}^{ij} \quad (*6)$$

We treat each case of activation separately.

Table 5: Hyperparameter of Equivariant Layers with ReLU activation. *Left* presents all possible case of tuple $((i, j, k), (s, p, q))$, and *Right* presents the parameter at the corresponding position. Here, we have three types of notations: 0 means the parameter equal to 0; equations with π 's in LHS means the equation holds for all possible π ; and a single term with no further information means the term can be arbitrary.

Tuple $((i, j, k), (s, p, q))$			Hyperparameter $(\mathbf{p}, \mathbf{q}, \mathbf{r}, \mathbf{s})$			
i and s	j and p	k and q	\mathbf{p}_{spq}^{ijk}	\mathbf{q}_{sp}^{ijk}	\mathbf{p}_{spq}^{ij}	\mathbf{p}_{sp}^{ij}
$i = s = 1$	$j \neq p$		0	0	0	0
	$j = p$		$\mathbf{p}_{1\pi(j)q}^{1\pi(j)k} = \mathbf{p}_{1jq}^{1jk}$	$\mathbf{q}_{1\pi(j)}^{1\pi(j)k} = \mathbf{q}_{1j}^{1jk}$	$\mathbf{r}_{1\pi(j)q}^{1\pi(j)} = \mathbf{r}_{1jq}^{1j}$	$\mathbf{s}_{1\pi(j)}^{1\pi(j)} = \mathbf{s}_{1j}^{1j}$
$i = s = L$		$k \neq q$	0	0	0	\mathbf{s}_{Lp}^{Lj}
		$k = q$	$\mathbf{p}_{Lp\pi(k)}^{Lj\pi(k)} = \mathbf{p}_{Ljq}^{Ljk}$	0	0	\mathbf{s}_{Lp}^{Lj}
$1 < i = s < L$	$j \neq p$		0	0	0	0
	$j = p$	$k \neq q$	0	0	0	$\mathbf{s}_{i\pi(j)}^{i\pi(j)} = \mathbf{s}_{ij}^{ij}$
		$k = q$	$\mathbf{p}_{i\pi(j)\pi'(k)}^{i\pi(j)\pi'(k)} = \mathbf{p}_{ijk}^{ijk}$	0	0	$\mathbf{s}_{i\pi(j)}^{i\pi(j)} = \mathbf{s}_{ij}^{ij}$
$i \neq s$			0	0	0	0

Table 6: Construction of equivariant functional layer with ReLU activation. Note that all parameters have to satisfy the conditions presented in Table 5.

Layer	Equivariant layer $E : (W, b) \mapsto (W', b')$	
	$W'_{jk}^{(i)}$	$b'_j^{(i)}$
$i = 1$	$\sum_{q=1}^{n_0} \mathbf{p}_{1jq}^{1jk} W_{jq}^{(1)} + \mathbf{q}_{1j}^{1jk} b_j^{(1)}$	$\sum_{q=1}^{n_0} \mathbf{r}_{1jq}^{1j} W_{jq}^{(1)} + \mathbf{s}_{1j}^{1j} b_j^{(1)}$
$1 < i < L$	$\mathbf{p}_{ijk}^{ijk} W_{jk}^{(i)}$	$\mathbf{s}_{ij}^{ij} b_j^{(i)}$
$i = L$	$\sum_{p=1}^{n_L} \mathbf{p}_{Lpk}^{Ljk} W_{pk}^{(L)}$	$\sum_{p=1}^{n_L} \mathbf{s}_{Lp}^{Lj} b_p^{(L)} + \mathbf{t}^{Lj}$

A.1 ReLU activation

Recall that, in this case:

$$G := \{\text{id}_{\mathcal{G}_{n_L}}\} \times \mathcal{G}_{n_{L-1}}^{>0} \times \dots \times \mathcal{G}_{n_1}^{>0} \times \{\text{id}_{\mathcal{G}_{n_0}}\}. \quad (47)$$

So the system of constraints (*) holds for:

1. all possible tuples $((i, j, k), (s, p, q))$,
2. all $\pi_i \in \mathcal{P}_i$ for $0 < i < L$, all $d_j^{(i)} > 0$ for $0 < i < L, 1 \leq j \leq n_i$,
3. $\pi_i = \text{id}_{\mathcal{G}_{n_i}}$ and $d_j^{(i)} = 1$ for $i = 0$ or $i = L$.

By treat each case of tuples $((i, j, k), (s, p, q))$, we solve Eq. *1, Eq. *2, Eq. *3, Eq. *4 in the system (*) for hyperparameter $(\mathbf{p}, \mathbf{q}, \mathbf{r}, \mathbf{s})$ as in Table 5. For \mathbf{t}^{ijk} and \mathbf{t}^{ij} , by Eq. *5, Eq. *6, we have $\mathbf{t}^{ijk} = 0$ for all (i, j, k) , $\mathbf{t}^{ij} = 0$ if $i < L$, and \mathbf{t}^{Lj} is arbitrary for all $1 \leq j \leq n_L$. In conclusion, the formula of equivariant layers E in case of activation ReLU is presented as in Table 6.

Example A.1. Let us consider a two-hidden-layers MLP with activation $\sigma = \text{ReLU}$. Assume that $n_0 = n_1 = n_2 = n_3 = 2$, i.e., all layers have two neurons. This MLP defines a function $f : \mathbb{R}^2 \rightarrow \mathbb{R}^2$ given by

$$f(x) = W^{(3)} \sigma \left(W^{(2)} \sigma \left(W^{(1)} x + b^{(1)} \right) + b^{(2)} \right) + b^{(3)},$$

where $W^{(i)} = \begin{pmatrix} W_{11}^{(i)} & W_{12}^{(i)} \\ W_{21}^{(i)} & W_{22}^{(i)} \end{pmatrix}$ is a 2×2 matrix and $b^{(i)} = \begin{bmatrix} b_1^{(i)} \\ b_2^{(i)} \end{bmatrix}$ for each $i = 1, 2, 3$. In this case, the weight space \mathcal{U} consists of the tuples

$$U = (W^{(1)}, W^{(2)}, W^{(3)}, b^{(1)}, b^{(2)}, b^{(3)})$$

and it has dimension 18.

According to Eq. (27), an equivariant layer E over \mathcal{U} has the form

$$E(U) = \left(W'^{(1)}, W'^{(2)}, W'^{(3)}, b'^{(1)}, b'^{(2)}, b'^{(3)} \right),$$

where

$$\begin{aligned} W'_{jk}{}^{(1)} &= \mathfrak{p}_{1j_1}^{1jk} W_{j_1 1}^{(1)} + \mathfrak{p}_{1j_2}^{1jk} W_{j_2 2}^{(1)} + \mathfrak{q}_{1j}^{1jk} b_j^{(1)}, & b'_j{}^{(1)} &= \mathfrak{r}_{j_1}^{1j} W_{j_1 1}^{(1)} + \mathfrak{r}_{j_2}^{1j} W_{j_2 2}^{(1)} + \mathfrak{s}_{1j}^{1j} b_j^{(1)}, \\ W'_{jk}{}^{(2)} &= \mathfrak{p}_{2j}^{2jk} W_{jk}^{(2)}, & b'_j{}^{(2)} &= \mathfrak{s}_{2j}^{2j} b_j^{(2)}, \\ W'_{jk}{}^{(3)} &= \mathfrak{p}_{3k_1}^{3jk} W_{3k}^{(3)} + \mathfrak{p}_{3k_2}^{3jk} W_{2k}^{(3)}, & b'_j{}^{(3)} &= \mathfrak{s}_{3j_1}^{3j} b_1^{(3)} + \mathfrak{s}_{3j_2}^{3j} b_2^{(3)} + \mathfrak{r}_j^3. \end{aligned}$$

These equations can be written in a friendly matrix form as follows.

$$\begin{aligned} \begin{bmatrix} W'_{11}{}^{(1)} \\ W'_{12}{}^{(1)} \\ W'_{21}{}^{(1)} \\ W'_{22}{}^{(1)} \\ b'_1{}^{(1)} \\ b'_2{}^{(1)} \end{bmatrix} &= \begin{bmatrix} \mathfrak{p}_{111}^{111} & \mathfrak{p}_{112}^{111} & 0 & 0 & \mathfrak{q}_{111}^{111} & 0 \\ \mathfrak{p}_{112}^{111} & \mathfrak{p}_{112}^{112} & 0 & 0 & \mathfrak{q}_{111}^{112} & 0 \\ 0 & 0 & \mathfrak{p}_{121}^{121} & \mathfrak{p}_{122}^{121} & 0 & \mathfrak{q}_{112}^{121} \\ 0 & 0 & \mathfrak{p}_{121}^{122} & \mathfrak{p}_{122}^{122} & 0 & \mathfrak{q}_{112}^{122} \\ \mathfrak{r}_{111}^{111} & \mathfrak{r}_{112}^{111} & 0 & 0 & \mathfrak{s}_{111}^{111} & 0 \\ 0 & 0 & \mathfrak{r}_{121}^{121} & \mathfrak{r}_{122}^{122} & 0 & \mathfrak{s}_{112}^{112} \end{bmatrix} \begin{bmatrix} W_{11}^{(1)} \\ W_{12}^{(1)} \\ W_{21}^{(1)} \\ W_{22}^{(1)} \\ b_1^{(1)} \\ b_2^{(1)} \end{bmatrix}, \\ \begin{bmatrix} W'_{11}{}^{(2)} \\ W'_{12}{}^{(2)} \\ W'_{21}{}^{(2)} \\ W'_{22}{}^{(2)} \\ b'_1{}^{(2)} \\ b'_2{}^{(2)} \end{bmatrix} &= \begin{bmatrix} \mathfrak{p}_{211}^{211} & 0 & 0 & 0 & 0 & 0 \\ 0 & \mathfrak{p}_{212}^{212} & 0 & 0 & 0 & 0 \\ 0 & 0 & \mathfrak{p}_{221}^{221} & 0 & 0 & 0 \\ 0 & 0 & 0 & \mathfrak{p}_{222}^{222} & 0 & 0 \\ 0 & 0 & 0 & 0 & \mathfrak{s}_{211}^{211} & 0 \\ 0 & 0 & 0 & 0 & 0 & \mathfrak{s}_{222}^{222} \end{bmatrix} \begin{bmatrix} W_{11}^{(2)} \\ W_{12}^{(2)} \\ W_{21}^{(2)} \\ W_{22}^{(2)} \\ b_1^{(2)} \\ b_2^{(2)} \end{bmatrix}, \\ \begin{bmatrix} W'_{11}{}^{(3)} \\ W'_{12}{}^{(3)} \\ W'_{21}{}^{(3)} \\ W'_{22}{}^{(3)} \\ b'_1{}^{(3)} \\ b'_2{}^{(3)} \end{bmatrix} &= \begin{bmatrix} \mathfrak{p}_{311}^{311} & 0 & \mathfrak{p}_{321}^{311} & 0 & 0 & 0 \\ 0 & \mathfrak{p}_{312}^{312} & 0 & \mathfrak{p}_{322}^{312} & 0 & 0 \\ \mathfrak{p}_{312}^{321} & 0 & \mathfrak{p}_{321}^{321} & 0 & 0 & 0 \\ 0 & \mathfrak{p}_{312}^{322} & 0 & \mathfrak{p}_{322}^{322} & 0 & 0 \\ 0 & 0 & 0 & 0 & \mathfrak{s}_{311}^{311} & \mathfrak{s}_{312}^{312} \\ 0 & 0 & 0 & 0 & \mathfrak{s}_{321}^{321} & \mathfrak{s}_{322}^{322} \end{bmatrix} \begin{bmatrix} W_{11}^{(3)} \\ W_{12}^{(3)} \\ W_{21}^{(3)} \\ W_{22}^{(3)} \\ b_1^{(3)} \\ b_2^{(3)} \end{bmatrix} + \begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \\ \mathfrak{r}_1^3 \\ \mathfrak{r}_2^3 \end{bmatrix}. \end{aligned}$$

A.2 Sin or Tanh activation

Recall that, in this case:

$$G := \{\text{id}_{\mathcal{G}_{n_L}}\} \times \mathcal{G}_{n_{L-1}}^{\pm 1} \times \dots \times \mathcal{G}_{n_1}^{\pm 1} \times \{\text{id}_{\mathcal{G}_{n_0}}\}. \quad (48)$$

So the system of constraints (*) holds for:

1. all possible tuples $((i, j, k), (s, p, q))$,
2. all $\pi_i \in \mathcal{P}_i$ for $0 < i < L$, all $d_j^{(i)} \in \{\pm 1\}$ for $0 < i < L$, $1 \leq j \leq n_i$,
3. $\pi_i = \text{id}_{\mathcal{G}_{n_i}}$ and $d_j^{(i)} = 1$ for $i = 0$ or $i = L$.

We assume $L \geq 3$, the case $L \leq 2$ can be solved similarly. By treat each case of tuples $((i, j, k), (s, p, q))$, we solve Eq. *1, Eq. *2, Eq. *3, Eq. *4 in the system (*) for hyperparameter $(\mathfrak{p}, \mathfrak{q}, \mathfrak{r}, \mathfrak{s})$ as in Table 7. For t^{ijk} and t^{ij} , by Eq. *5, Eq. *6, we have $t^{ijk} = 0$ for all (i, j, k) , $t^{ij} = 0$ if $i < L$, and t^{Lj} is arbitrary for all $1 \leq j \leq n_L$. In conclusion, the formula of equivariant layers E in case of sin or Tanh activation is presented as in Table 8.

Table 7: Hyperparameter of Equivariant Layers with sin or Tanh activation. *Left* presents all possible case of tuple $((i, j, k), (s, p, q))$, and *Right* presents the parameter at the corresponding position. Here, we have three types of notations: 0 means the parameter equal to 0; equations with π 's in LHS means the equation holds for all possible π ; and a single term with no further information means the term can be arbitrary.

Tuple $((i, j, k), (s, p, q))$			Hyperparameter $(p, q, \tau, \mathfrak{s})$			
i and s	j and p	k and q	\mathfrak{p}_{spq}^{ijk}	\mathfrak{q}_{sp}^{ijk}	\mathfrak{r}_{spq}^{ij}	\mathfrak{s}_{sp}^{ij}
$i = s = 1$	$j \neq p$		0	0	0	0
	$j = p$		$\mathfrak{p}_{1\pi(j)q}^{1\pi(j)k} = \mathfrak{p}_{1jq}^{1jk}$	$\mathfrak{q}_{1\pi(j)}^{1\pi(j)k} = \mathfrak{q}_{1j}^{1jk}$	$\mathfrak{r}_{1\pi(j)q}^{1\pi(j)} = \mathfrak{r}_{1jq}^{1j}$	$\mathfrak{s}_{1\pi(j)}^{1\pi(j)} = \mathfrak{s}_{1j}^{1j}$
$i = s = L$		$k \neq q$	0	0	0	\mathfrak{s}_{Lp}^{Lj}
		$k = q$	$\mathfrak{p}_{Lp\pi(k)}^{Lj\pi(k)} = \mathfrak{p}_{Ljq}^{Ljk}$	0	0	\mathfrak{s}_{Lp}^{Lj}
$1 < i = s < L$	$j \neq p$		0	0	0	0
	$j = p$	$k \neq q$	0	0	0	$\mathfrak{s}_{i\pi(j)}^{i\pi(j)} = \mathfrak{s}_{ij}^{ij}$
		$k = q$	$\mathfrak{p}_{i\pi(j)\pi'(k)}^{i\pi(j)\pi'(k)} = \mathfrak{p}_{ijk}^{ijk}$	0	0	$\mathfrak{s}_{i\pi(j)}^{i\pi(j)} = \mathfrak{s}_{ij}^{ij}$
$(i, s) = (L-1, L)$	$j = q$	0	0	$\mathfrak{r}_{Lp\pi(j)}^{(L-1)\pi(j)} = \mathfrak{r}_{Lpj}^{(L-1)j}$	0	
$(i, s) = (L, L-1)$		$k = p$	0	$\mathfrak{q}_{(L-1)\pi(k)}^{Lj\pi(k)} = \mathfrak{q}_{(L-1)k}^{Ljk}$	0	0
otherwise			0	0	0	0

Table 8: Construction of equivariant functional layer with sin or Tanh activation. Note that all parameters have to satisfy the conditions presented in Table 5.

Layer	Equivariant layer $E : (W, b) \mapsto (W', b')$	
	$W_{jk}'^{(i)}$	$b_j'^{(i)}$
$i = 1$	$\sum_{q=1}^{n_0} \mathfrak{p}_{1jq}^{1jk} W_{jq}^{(1)} + \mathfrak{q}_{1j}^{1jk} b_j^{(1)}$	$\sum_{q=1}^{n_0} \mathfrak{r}_{1jq}^{1j} W_{jq}^{(1)} + \mathfrak{s}_{1j}^{1j} b_j^{(1)}$
$1 < i < L-1$	$\mathfrak{p}_{ijk}^{ijk} W_{jk}^{(i)}$	$\mathfrak{s}_{ij}^{ij} b_j^{(i)}$
$i = L-1$	$\mathfrak{p}_{(L-1)jk}^{(L-1)jk} W_{jk}^{(L-1)}$	$\sum_{p=1}^{n_L} \mathfrak{r}_{Lpj}^{(L-1)j} W_{pj}^{(L)} + \mathfrak{s}_{(L-1)j}^{(L-1)j} b_j^{(L-1)}$
$i = L$	$\sum_{p=1}^{n_L} \mathfrak{p}_{Lpk}^{Ljk} W_{pk}^{(L)} + \mathfrak{q}_{(L-1)k}^{Ljk} b_k^{(L-1)}$	$\sum_{p=1}^{n_L} \mathfrak{s}_{Lp}^{Lj} b_p^{(L)} + \mathfrak{t}^{Lj}$

B Construction of Monomial Matrix Group Invariant Layers

In this appendix, we present how we constructed Monomial Matrix Group Invariant Layers. Let \mathcal{U} be a weight spaces with the number of layers L as well as the number of channels at i -th layer n_i . We want to construct G -invariant layers $I : \mathcal{U} \rightarrow \mathbb{R}^d$ for some $d > 0$. We treat each case of activations separately.

B.1 ReLU activation

Recall that, in this case:

$$G := \{\text{id}_{\mathcal{G}_{n_L}}\} \times \mathcal{G}_{n_{L-1}}^{\pm 1} \times \dots \times \mathcal{G}_{n_1}^{\pm 1} \times \{\text{id}_{\mathcal{G}_{n_0}}\}. \quad (49)$$

Since $\mathcal{G}_*^{>0}$ is the semidirect product of $\Delta_*^{>0}$ and \mathcal{P}_* with $\Delta_*^{>0}$ is the normal subgroup, we will treat these two actions consecutively, $\Delta_*^{>0}$ first then \mathcal{P}_* . We denote these layers by $I_{\Delta_*^{>0}}$ and $I_{\mathcal{P}_*}$. Note that, since $I_{\Delta_*^{>0}}$ comes before $I_{\mathcal{P}_*}$, $I_{\Delta_*^{>0}}$ is required to be $\Delta_*^{>0}$ -invariant and \mathcal{P}_* -equivariant, and $I_{\mathcal{P}_*}$ is required to be \mathcal{P}_* -invariant.

$\Delta_*^{>0}$ -invariance and \mathcal{P}_* -equivariance. To capture $\Delta_*^{>0}$ -invariance, we recall the notion of positively homogeneous of degree zero maps. For $n > 0$, a map α from \mathbb{R}^n is called *positively*

Table 9: Constraints of α component in invariant functional layer with ReLU, sin, Tanh activations.

Layer	$I_{\Delta>0} : (W, b) \mapsto (W', b')$	
	$\alpha_{jk}^{(i)} : W_{jk}^{(i)} \mapsto W_{jk}'^{(i)}$	$\alpha_j^{(i)} : b_j^{(i)} \mapsto b_j'^{(i)}$
$i = 1$	$\alpha_{\pi(j)k}^{(i)} = \alpha_{jk}^{(i)}$	$\alpha_{\pi(j)}^{(i)} = \alpha_j^{(i)}$
$1 < i < L$	$\alpha_{\pi(j)\pi'(k)}^{(i)} = \alpha_{jk}^{(i)}$	$\alpha_{\pi(j)}^{(i)} = \alpha_j^{(i)}$
$i = L$	$\alpha_{j\pi(k)}^{(i)} = \alpha_{jk}^{(i)}$	$\alpha_j^{(i)}$

homogeneous of degree zero if

$$\alpha(\lambda x_1, \dots, \lambda x_n) = \alpha(x_1, \dots, x_n). \quad (50)$$

for all $\lambda > 0$ and $(x_1, \dots, x_n) \in \mathbb{R}^n$. We construct $I_{\Delta>0} : \mathcal{U} \rightarrow \mathcal{U}$ by taking collections of positively homogeneous of degree zero functions $\{\alpha_{jk}^{(i)} : \mathbb{R}^{w_i} \rightarrow \mathbb{R}^{w_i}\}$ and $\{\alpha_j^{(i)} : \mathbb{R}^{b_i} \rightarrow \mathbb{R}^{b_i}\}$, each one corresponds to weight and bias of \mathcal{U} . The maps $I_{\Delta>0} : \mathcal{U} \rightarrow \mathcal{U}$ that $(W, b) \mapsto (W', b')$ is defined by simply applying these functions on each weight and bias entries as follows:

$$W_{jk}'^{(i)} = \alpha_{jk}^{(i)}(W_{jk}^{(i)}) \quad \text{and} \quad b_j'^{(i)} = \alpha_j^{(i)}(b_j^{(i)}). \quad (51)$$

$I_{\Delta>0}$ is $\Delta_*^{>0}$ -invariant by homogeneity of the α functions. To make it become \mathcal{P}_* -equivariant, some α functions have to be shared across any axis that have permutation symmetry, presented in Table 9.

Candidates of function α . We simply choose positively homogeneous of degree zero function $\alpha : \mathbb{R}^n \rightarrow \mathbb{R}^n$ by taking $\alpha(0) = 0$ and:

$$\alpha(x_1, \dots, x_n) = \beta \left(\frac{x_1^2}{x_1^2 + \dots + x_n^2}, \dots, \frac{x_n^2}{x_1^2 + \dots + x_n^2} \right). \quad (52)$$

where $\beta : \mathbb{R}^n \rightarrow \mathbb{R}^n$ is an arbitrary function. The function β can be fixed or parameterized to make α to be fixed or learnable.

\mathcal{P}_* -invariance. To capture \mathcal{P}_* -invariance, we simply take summing or averaging the weight and bias across any axis that have permutation symmetry as in [71]. In concrete, some $d > 0$, we have $I_{\mathcal{P}} : \mathcal{U} \rightarrow \mathbb{R}^d$ is computed as follows:

$$I_{\mathcal{P}}(U) = \left(W_{\star, \cdot}^{(1)}, W_{\cdot, \star}^{(L)}, W_{\star, \star}^{(2)}, \dots, W_{\star, \star}^{(L-1)}; v^{(L)}, v_{\star}^{(1)}, \dots, v_{\star}^{(L-1)} \right). \quad (53)$$

Here, \star denotes summation or averaging over the rows or columns of the weight and bias.

G -invariance. Now we simply compose $I_{\mathcal{P}} \circ I_{\Delta>0}$ to get an G -invariant map. We use an MLP to complete constructing an G -invariant layer with output dimension d as desired:

$$I = \text{MLP} \circ I_{\mathcal{P}} \circ I_{\Delta>0}. \quad (54)$$

B.2 Sin or Tanh activation

Recall that, in this case:

$$G := \{\text{id}_{\mathcal{G}_{n_L}}\} \times \mathcal{G}_{n_{L-1}}^{\pm 1} \times \dots \times \mathcal{G}_{n_1}^{\pm 1} \times \{\text{id}_{\mathcal{G}_{n_0}}\}. \quad (55)$$

Since $\mathcal{G}_*^{\pm 1}$ is the semidirect product of $\Delta_*^{\pm 1}$ and \mathcal{P}_* with $\Delta_*^{\pm 1}$ is the normal subgroup, we will treat these two actions consecutively, $\Delta_*^{\pm 1}$ first then \mathcal{P}_* . We denote these layers by $I_{\Delta^{\pm 1}}$ and $I_{\mathcal{P}}$. Note that, since $I_{\Delta^{\pm 1}}$ comes before $I_{\mathcal{P}}$, $I_{\Delta^{\pm 1}}$ is required to be $\Delta_*^{\pm 1}$ -invariant and \mathcal{P}_* -equivariant, and $I_{\mathcal{P}}$ is required to be \mathcal{P}_* -invariant.

$\Delta_*^{\pm 1}$ -invariance and \mathcal{P}_* -equivariance. To capture $\Delta_*^{\pm 1}$ -invariance, we use even functions, i.e. $\alpha(x) = \alpha(-x)$ for all x . We construct $I_{\Delta^{\pm 1}}: \mathcal{U} \rightarrow \mathcal{U}$ by taking collections of even functions $\{\alpha_{jk}^{(i)}: \mathbb{R}^{w_i} \rightarrow \mathbb{R}^{w_i}\}$ and $\{\alpha_j^{(i)}: \mathbb{R}^{b_i} \rightarrow \mathbb{R}^{b_i}\}$, each one corresponds to weight and bias of \mathcal{U} . The maps $I_{\Delta^{\pm 1}}: \mathcal{U} \rightarrow \mathcal{U}$ that $(W, b) \mapsto (W', b')$ is defined by simply applying these functions on each weight and bias entries as follows:

$$W'_{jk}{}^{(i)} = \alpha_{jk}^{(i)}(W_{jk}^{(i)}) \quad \text{and} \quad b'_j{}^{(i)} = \alpha_j^{(i)}(b_j^{(i)}). \quad (56)$$

$I_{\Delta^{\pm 1}}$ is $\Delta_*^{\pm 1}$ -invariant by design. To make it become \mathcal{P}_* -equivariant, some α functions have to be shared across any axis that have permutation symmetry, presented in Table 9.

Candidates of function α . We simply choose even function $\alpha: \mathbb{R}^n \rightarrow \mathbb{R}^n$ by:

$$\alpha(x_1, \dots, x_n) = \beta(|x_1|, \dots, |x_n|). \quad (57)$$

where $\beta: \mathbb{R}^n \rightarrow \mathbb{R}^n$ is an arbitrary function. The function β can be fixed or parameterized to make α to be fixed or learnable.

\mathcal{P}_* -invariance. To capture \mathcal{P}_* -invariance, we simply take summing or averaging the weight and bias across any axis that have permutation symmetry as in [71]. In concrete, some $d > 0$, we have $I_{\mathcal{P}}: \mathcal{U} \rightarrow \mathbb{R}^d$ is computed as follows:

$$I_{\mathcal{P}}(U) = \left(W_{\star, \cdot}^{(1)}, W_{\cdot, \star}^{(L)}, W_{\star, \star}^{(2)}, \dots, W_{\star, \star}^{(L-1)}; v_{\star}^{(L)}, v_{\star}^{(1)}, \dots, v_{\star}^{(L-1)} \right). \quad (58)$$

Here, \star denotes summation or averaging over the rows or columns of the weight and bias.

G -invariance. Now we simply compose $I_{\mathcal{P}} \circ I_{\Delta^{\pm 1}}$ to get an G -invariant map. We use an MLP to complete constructing an G -invariant layer with output dimension d as desired:

$$I = \text{MLP} \circ I_{\mathcal{P}} \circ I_{\Delta^{\pm 1}}. \quad (59)$$

C Proofs of Theoretical Results

C.1 Proof of Proposition 3.4

Proof. We simply denote the activation ReLU or sin or tanh by σ . Let $A \in \text{GL}(n)$ that satisfies:

$$\sigma(A \cdot \mathbf{x}) = A \cdot \sigma(\mathbf{x}),$$

for all $\mathbf{x} \in \mathbb{R}^n$. This means:

$$\sigma \left(\begin{bmatrix} a_{11} & \dots & a_{1n} \\ \vdots & \ddots & \vdots \\ a_{n1} & \dots & a_{nn} \end{bmatrix} \cdot \begin{bmatrix} x_1 \\ \vdots \\ x_n \end{bmatrix} \right) = \begin{bmatrix} a_{11} & \dots & a_{1n} \\ \vdots & \ddots & \vdots \\ a_{n1} & \dots & a_{nn} \end{bmatrix} \cdot \sigma \left(\begin{bmatrix} x_1 \\ \vdots \\ x_n \end{bmatrix} \right),$$

for all $x_1, \dots, x_n \in \mathbb{R}$. We rewrite this equation as:

$$\sigma \left(\begin{bmatrix} a_{11}x_1 + a_{12}x_2 + \dots + a_{1n}x_n \\ \vdots \\ a_{n1}x_1 + a_{n2}x_2 + \dots + a_{nn}x_n \end{bmatrix} \right) = \begin{bmatrix} a_{11} & \dots & a_{1n} \\ \vdots & \ddots & \vdots \\ a_{n1} & \dots & a_{nn} \end{bmatrix} \cdot \begin{bmatrix} \sigma(x_1) \\ \vdots \\ \sigma(x_n) \end{bmatrix},$$

or equivalently:

$$\begin{bmatrix} \sigma(a_{11}x_1 + a_{12}x_2 + \dots + a_{1n}x_n) \\ \vdots \\ \sigma(a_{n1}x_1 + a_{n2}x_2 + \dots + a_{nn}x_n) \end{bmatrix} = \begin{bmatrix} a_{11}\sigma(x_1) + a_{12}\sigma(x_2) + \dots + a_{1n}\sigma(x_n) \\ \vdots \\ a_{n1}\sigma(x_1) + a_{n2}\sigma(x_2) + \dots + a_{nn}\sigma(x_n) \end{bmatrix}.$$

Thus,

$$\sigma \left(\sum_{j=1}^n a_{ij}x_j \right) = \sum_{j=1}^n a_{ij}\sigma(x_j),$$

for all $x_1, \dots, x_n \in \mathbb{R}$ and $i = 1, \dots, n$. We will consider the case $i = 1$, i.e.

$$\sigma \left(\sum_{j=1}^n a_{1j}x_j \right) = \sum_{j=1}^n a_{1j}\sigma(x_j), \quad (60)$$

and treat the case $i > 1$ similarly. Now we consider the activation σ case by case as follows.

(i) **Case 1.** $\sigma = \text{ReLU}$. We have some observations:

1. Let $x_1 = 1$, and $x_2 = \dots = x_n = 0$. Then from Eq. (60), we have:

$$\sigma(a_{11}) = a_{11},$$

which implies that $a_{11} \geq 0$. Similarly, we also have $a_{12}, \dots, a_{1n} \geq 0$.

2. Since A is an invertible matrix, the entries a_{11}, \dots, a_{1n} in the first row of A can not be simultaneously equal to 0.
3. There is at most only one nonzero number among the entries a_{11}, \dots, a_{1n} . Indeed, assume by the contrary that $a_{11}, a_{12} > 0$. Let $x_3 = \dots = x_n = 0$, from Eq. (60), we have:

$$\sigma(a_{11}x_1 + a_{12}x_2) = a_{11}\sigma(x_1) + a_{12}\sigma(x_2).$$

Let $x_2 = -1$, we have:

$$\sigma(a_{11}x_1 - a_{12}) = a_{11}\sigma(x_1).$$

Now, let $x_1 > 0$ be a sufficiently large number such that $a_{11}x_1 - a_{12} > 0$. (Note that this number exists since $a_{11} > 0$). Then we have:

$$a_{11}x_1 - a_{12} = a_{11}x_1,$$

which implies $a_{12} = 0$, a contradiction.

It follows from these three observations that there is exactly one non-zero element among the entries a_{11}, \dots, a_{1n} . In other words, matrix A has exactly one nonzero entry in the first row. This applies for every row, so A has exactly one non-zero entry in each row. Since A is invertible, each column of A has at least one non-zero entry. Thus A also has exactly one non-zero entry in each column. Hence, A is in \mathcal{G}_n . Moreover, all entries of A are non-negative, so A is in $\mathcal{G}_n^{>0}$.

It is straight forward to check that for all A in $\mathcal{G}_n^{>0}$ we have $\sigma(A \cdot \mathbf{x}) = A \cdot \sigma(\mathbf{x})$.

(ii) **Case 2.** $\sigma = \text{Tanh}$ or $\sigma = \sin$. We have some observations:

1. Let $x_2 = \dots = x_n = 0$. Then from Eq. (60), we have:

$$\sigma(a_{11}x_1) = a_{11}\sigma(x_1),$$

which implies $a_{11} \in \{-1, 0, 1\}$. Similarly, we have $a_{12}, \dots, a_{1n} \in \{-1, 0, 1\}$.

2. Since A is an invertible matrix, the entries a_{11}, \dots, a_{1n} in the first row of A can not be simultaneously equal to 0.
3. There is at most only one nonzero number among the entries a_{11}, \dots, a_{1n} . Indeed, assume by the contrary that $a_{11}, a_{12} \neq 0$. Let $x_3 = \dots = x_n = 0$, from Eq. (60), we have:

$$\sigma(a_{11}x_1 + a_{12}x_2) = a_{11}\sigma(x_1) + a_{12}\sigma(x_2).$$

Note that $a_{11}, a_{12} \in \{-1, 1\}$, so by consider all the cases, we will lead to a contradiction.

It follows from the above three observations that there is exactly one non-zero element among the entries a_{11}, \dots, a_{1n} . In other words, matrix A has exactly one nonzero entry in the first row. This applies for every row, so A has exactly one non-zero entry in each row. Note that, since A is invertible, each column of A has at least one non-zero entry. Therefore, A also has exactly one non-zero entry in each column. Hence, A is in \mathcal{G}_n . Moreover, all entries of A are in $\{-1, 0, 1\}$, so A is in $\mathcal{G}_n^{\pm 1}$.

It is straight forward to check that for all A in $\mathcal{G}_n^{\pm 1}$ we have $\sigma(A \cdot \mathbf{x}) = A \cdot \sigma(\mathbf{x})$.

The proposition is then proved completely. \square

C.2 Proof of Proposition 4.4

Proof. For both Fully Connected Neural Networks case and Convolutional Neural Networks case, we consider a network f with three layers, with n_0, n_1, n_2, n_3 are number of channels at each layer, and its weight space \mathcal{U} . We will show the proof for part (i) where activation σ is ReLU, and part (ii) can be proved similarly. For part (i), we prove f to be G -invariant on its weight space \mathcal{U} , for the group G that is defined by:

$$G = \{\text{id}_{\mathcal{G}_{n_3}}\} \times \mathcal{G}_{n_2}^{>0} \times \mathcal{G}_{n_1}^{>0} \times \{\text{id}_{\mathcal{G}_{n_0}}\} \subset \mathcal{G}_{n_3} \times \mathcal{G}_{n_2} \times \mathcal{G}_{n_1} \times \mathcal{G}_{n_0} = \mathcal{G}_{\mathcal{U}};$$

Case 1. f is a Fully Connected Neural Network with three layers, with n_0, n_1, n_2, n_3 are number of channels at each layer as in Eq. 5:

$$f(\mathbf{x}; U, \sigma) = W^{(3)} \cdot \sigma \left(W^{(2)} \cdot \sigma \left(W^{(1)} \cdot \mathbf{x} + b^{(1)} \right) + b^{(2)} \right) + b^{(3)},$$

Case 2. f is a Convolutional Neural Network with three layers, with n_0, n_1, n_2, n_3 are number of channels at each layer as in Eq. 8:

$$f(\mathbf{x}; U, \sigma) = W^{(3)} * \sigma \left(W^{(2)} * \sigma \left(W^{(1)} * \mathbf{x} + b^{(1)} \right) + b^{(2)} \right) + b^{(3)}$$

We have some observations:

For case 1. For $W \in \mathbb{R}^{m \times n}$, $\mathbf{x} \in \mathbb{R}^n$ and $a > 0$, we have:

$$a \cdot \sigma(W \cdot \mathbf{x} + b) = \sigma((aW) \cdot \mathbf{x} + (ab)).$$

For case 2. For simplicity, we consider $*$ as one-dimensional convolutional operator, and other types of convolutions can be treated similarly. For $W = (w_1, \dots, w_m) \in \mathbb{R}^m$, $b \in \mathbb{R}$ and $\mathbf{x} = (x_1, \dots, x_n) \in \mathbb{R}^n$, we have:

$$W * \mathbf{x} + b = \mathbf{y} = (y_1, \dots, y_{n-m+1}) \in \mathbb{R}^{n-m+1},$$

where:

$$y_i = \sum_{j=1}^m w_j x_{i+j-1} + b.$$

So for $a > 0$, we have:

$$a \cdot \sigma(W * \mathbf{x} + b) = \sigma((aW) * \mathbf{x} + (ab)).$$

With these two observations, we can see the proofs for both cases are similar to each other. We will show the proof for case 2, when f is a convolutional neural network since it is not trivial as case 1. Now we have $U = (W, b)$ with:

$$W = (W^{(3)}, W^{(2)}, W^{(1)}),$$

$$b = (b^{(3)}, b^{(2)}, b^{(1)}).$$

Let g be an element of G :

$$g = (\text{id}_{\mathcal{G}_{n_3}}, g^{(2)}, g^{(1)}, \text{id}_{\mathcal{G}_{n_0}}),$$

where:

$$g^{(2)} = D^{(2)} \cdot P_{\pi_2} = \text{diag}(d_1^{(2)}, \dots, d_{n_2}^{(2)}) \cdot P_{\pi_2} \in \mathcal{G}_{n_2}^{>0},$$

$$g^{(1)} = D^{(1)} \cdot P_{\pi_1} = \text{diag}(d_1^{(1)}, \dots, d_{n_1}^{(1)}) \cdot P_{\pi_1} \in \mathcal{G}_{n_1}^{>0}.$$

We compute gU :

$$gU = (gW, gb),$$

$$gW = ((gW)^{(3)}, (gW)^{(2)}, (gW)^{(1)}),$$

$$gb = ((gb)^{(3)}, (gb)^{(2)}, (gb)^{(1)}).$$

where:

$$\begin{aligned}
(gW)_{jk}^{(3)} &= \frac{1}{d_k^{(2)}} \cdot W_{j\pi_2^{-1}(k)}^{(3)}, \\
(gW)_{jk}^{(2)} &= \frac{d_j^{(2)}}{d_k^{(1)}} \cdot W_{\pi_2^{-1}(j)\pi_1^{-1}(k)}^{(2)}, \\
(gW)_{jk}^{(1)} &= \frac{d_j^{(1)}}{1} \cdot W_{\pi_1^{-1}(j)k}^{(1)},
\end{aligned}$$

and,

$$\begin{aligned}
(gb)_j^{(3)} &= b_j^{(3)}, \\
(gb)_j^{(2)} &= d_j^{(2)} \cdot b_{\pi_2^{-1}(j)}^{(2)}, \\
(gb)_j^{(1)} &= d_j^{(1)} \cdot b_{\pi_1^{-1}(j)}^{(1)}.
\end{aligned}$$

Now we show that $f(\mathbf{x}; U, \sigma) = f(\mathbf{x}; gU, \sigma)$ for all $\mathbf{x} = (x_1, \dots, x_{n_0}) \in \mathbb{R}^{n_0}$. For $1 \leq i \leq n_3$, we compute the i -th entry of $f(\mathbf{x}; gU, \sigma)$ as follows:

$$\begin{aligned}
&f(\mathbf{x}; gU, \sigma)_i \\
&= \sum_{j_2=1}^{n_2} (gW)_{ij_2}^{(3)} * \sigma \left(\sum_{j_1=1}^{n_1} (gW)_{j_2j_1}^{(2)} * \right. \\
&\quad \left. \sigma \left(\sum_{j_0=1}^{n_0} (gW)_{j_1j_0}^{(1)} * x_{j_0} + (gb)_{j_1}^{(1)} \right) + (gb)_{j_2}^{(2)} \right) + (gb)_i^{(3)} \\
&= \sum_{j_2=1}^{n_2} \frac{1}{d_{j_2}^{(2)}} \cdot W_{i\pi_2^{-1}(j_2)}^{(3)} * \sigma \left(\sum_{j_1=1}^{n_1} \frac{d_{j_2}^{(2)}}{d_{j_1}^{(1)}} \cdot W_{\pi_2^{-1}(j_2)\pi_1^{-1}(j_1)}^{(2)} * \right. \\
&\quad \left. \sigma \left(\sum_{j_0=1}^{n_0} \frac{d_{j_1}^{(1)}}{1} \cdot W_{\pi_1^{-1}(j_1)j_0}^{(1)} * x_{j_0} + d_{j_1}^{(1)} \cdot b_{\pi_1^{-1}(j_1)}^{(1)} \right) + d_{j_2}^{(2)} \cdot b_{\pi_2^{-1}(j_2)}^{(2)} \right) + b_i^{(3)} \\
&= \sum_{j_2=1}^{n_2} \frac{1}{d_{j_2}^{(2)}} \cdot W_{i\pi_2^{-1}(j_2)}^{(3)} * \sigma \left(\sum_{j_1=1}^{n_1} \frac{d_{j_2}^{(2)}}{d_{j_1}^{(1)}} \cdot W_{\pi_2^{-1}(j_2)\pi_1^{-1}(j_1)}^{(2)} * \right. \\
&\quad \left. \sigma \left(d_{j_1}^{(1)} \cdot \left(\sum_{j_0=1}^{n_0} W_{\pi_1^{-1}(j_1)j_0}^{(1)} * x_{j_0} + b_{\pi_1^{-1}(j_1)}^{(1)} \right) \right) + d_{j_2}^{(2)} \cdot b_{\pi_2^{-1}(j_2)}^{(2)} \right) + b_i^{(3)} \\
&= \sum_{j_2=1}^{n_2} \frac{1}{d_{j_2}^{(2)}} \cdot W_{i\pi_2^{-1}(j_2)}^{(3)} * \sigma \left(\sum_{j_1=1}^{n_1} \frac{d_{j_2}^{(2)}}{d_{j_1}^{(1)}} \cdot W_{\pi_2^{-1}(j_2)\pi_1^{-1}(j_1)}^{(2)} * \right. \\
&\quad \left. d_{j_1}^{(1)} \cdot \sigma \left(\sum_{j_0=1}^{n_0} W_{\pi_1^{-1}(j_1)j_0}^{(1)} * x_{j_0} + b_{\pi_1^{-1}(j_1)}^{(1)} \right) + d_{j_2}^{(2)} \cdot b_{\pi_2^{-1}(j_2)}^{(2)} \right) + b_i^{(3)} \\
&= \sum_{j_2=1}^{n_2} \frac{1}{d_{j_2}^{(2)}} \cdot W_{i\pi_2^{-1}(j_2)}^{(3)} * \sigma \left(\sum_{j_1=1}^{n_1} d_{j_2}^{(2)} \cdot W_{\pi_2^{-1}(j_2)\pi_1^{-1}(j_1)}^{(2)} * \right. \\
&\quad \left. \sigma \left(\sum_{j_0=1}^{n_0} W_{\pi_1^{-1}(j_1)j_0}^{(1)} * x_{j_0} + b_{\pi_1^{-1}(j_1)}^{(1)} \right) + d_{j_2}^{(2)} \cdot b_{\pi_2^{-1}(j_2)}^{(2)} \right) + b_i^{(3)}
\end{aligned}$$

$$\begin{aligned}
&= \sum_{j_2=1}^{n_2} \frac{1}{d_{j_2}^{(2)}} \cdot W_{i\pi_2^{-1}(j_2)}^{(3)} * \sigma \left(d_{j_2}^{(2)} \cdot \left(\sum_{j_1=1}^{n_1} W_{\pi_2^{-1}(j_2)\pi_1^{-1}(j_1)}^{(2)} * \right. \right. \\
&\quad \left. \left. \sigma \left(\sum_{j_0=1}^{n_0} W_{\pi_1^{-1}(j_1)j_0}^{(1)} * x_{j_0} + b_{\pi_1^{-1}(j_1)}^{(1)} \right) + b_{\pi_2^{-1}(j_2)}^{(2)} \right) \right) + b_i^{(3)} \\
&= \sum_{j_2=1}^{n_2} \frac{1}{d_{j_2}^{(2)}} \cdot W_{i\pi_2^{-1}(j_2)}^{(3)} \cdot d_{j_2}^{(2)} * \sigma \left(\sum_{j_1=1}^{n_1} W_{\pi_2^{-1}(j_2)\pi_1^{-1}(j_1)}^{(2)} * \right. \\
&\quad \left. \sigma \left(\sum_{j_0=1}^{n_0} W_{\pi_1^{-1}(j_1)j_0}^{(1)} * x_{j_0} + b_{\pi_1^{-1}(j_1)}^{(1)} \right) + b_{\pi_2^{-1}(j_2)}^{(2)} \right) + b_i^{(3)} \\
&= \sum_{j_2=1}^{n_2} W_{i\pi_2^{-1}(j_2)}^{(3)} * \sigma \left(\sum_{j_1=1}^{n_1} W_{\pi_2^{-1}(j_2)\pi_1^{-1}(j_1)}^{(2)} * \right. \\
&\quad \left. \sigma \left(\sum_{j_0=1}^{n_0} W_{\pi_1^{-1}(j_1)j_0}^{(1)} * x_{j_0} + b_{\pi_1^{-1}(j_1)}^{(1)} \right) + b_{\pi_2^{-1}(j_2)}^{(2)} \right) + b_i^{(3)} \\
&= \sum_{j_2=1}^{n_2} W_{ij_2}^{(3)} * \sigma \left(\sum_{j_1=1}^{n_1} W_{j_2j_1}^{(2)} * \sigma \left(\sum_{j_0=1}^{n_0} W_{j_1j_0}^{(1)} * x_{j_0} + b_{j_1}^{(1)} \right) + b_{j_2}^{(2)} \right) + b_i^{(3)} \\
&= f(\mathbf{x}; U, \sigma)_i.
\end{aligned}$$

End of proof. □

D Additional experimental details

D.1 Runtime and Memory Consumption

We provide the runtime and memory consumption of Monomial-NFNs and the previous NFNs in Tables 10 and 11 to compare the computational and memory costs in the task of predicting CNN generalization (see Section 6.1). It is observable that our model runs faster and consumes significantly less memory than NP/HNP in [71] and GNN-based method in [35]. This highlights the benefits of parameter savings in Monomial-NFN.

Table 10: Runtime of models.

	NP [71]	HNP [71]	GNN [35]	Monomial-NFN (ours)
Tanh subset	35m34s	29m37s	4h25m17s	18m23s
ReLU subset	36m40s	30m06s	4h27m29s	23m47s

Table 11: Memory consumption.

	NP [71]	HNP [71]	GNN [35]	Monomial-NFN (ours)
Tanh subset	838MB	856MB	6390MB	582MB
ReLU subset	838MB	856MB	6390MB	560MB

D.2 Comparison of Monomial-NFNs and GNN-based NFNs

We provide experimental result to compare the efficiency of our model and a permutation equivariant GNN-based NFN [35] in two scenarios below.

1. Training the model on augmented train data and testing with the augmented test data (see Tables 12 and 13).

Here, we present the experimental results on the original dataset and the results on the augmented dataset. The augmentation levels for the ReLU subset are 1, 2, 3, and 4,

corresponding to augmentation ranges of $[1, 10]$, $[1, 10^2]$, $[1, 10^3]$, $[1, 10^4]$. The augmented dataset for the Tanh subset corresponds to the augmentation range of $[-1, 1]$

Table 12: Predict CNN generalization on ReLU subset (augmented train data)

	Original	1	2	3	4
GNN [35]	0.897	0.892	0.885	0.858	0.851
Monomial-NF (ours)	0.922	0.920	0.919	0.920	0.920

Table 13: Predict CNN generalization on Tanh subset (augmented train data)

	Original	Augmented
GNN [35]	0.893	0.902
Monomial-NFN (ours)	0.939	0.943

The results for GNN exhibit a similar trend as other baselines that do not incorporate the scaling symmetry into their architectures. In contrast, our model has stable performance. A notable observation is that the GNN model uses 5.5M parameters (4 times more than our model), occupies 6000MB of memory, and takes 4 hours to train.

2. Training the model on original train data and testing with the augmented test data (see Tables 14 and 15).

Table 14: Predict CNN generalization on ReLU subset (original train data)

Augment level	1	2	3	4
GNN [35]	0.794	0.679	0.586	0.562
Monomial-NF (ours)	0.920	0.919	0.920	0.920

Table 15: Predict CNN generalization on Tanh subset (original train data)

	Augmented
GNN [35]	0.883
Monomial-NFN (ours)	0.940

In these more challenging scenario, GNN’s performance drops significantly, which highlights the lack of scaling symmetry in the model. Our model maintains consistent performance, matching the case in which we train with the augmented data.

D.3 Predicting generalization from weights

Dataset. The original ReLU subset of the CNN Zoo dataset includes 6050 instances for training and 1513 instances for testing. For the Tanh dataset, it includes 5949 training and 1488 testing instances. For the augmented data, we set the augmentation factor to 2, which means that we augment the original data once, resulting in a new dataset of double the size. The complete size of all datasets is presented in Table 16

Implementation details. Our model follows the same architecture as in [71], comprising three equivariant Monomial-NFN layers with 16, 16, and 5 channels, respectively, each followed by ReLU activation (ReLU dataset) or Tanh activation (Tanh dataset). The resulting weight space features are input into an invariant Monomial-NFN layer with Monomial-NFN pooling (Equation 19) with learnable parameters (ReLU case) or mean pooling (Tanh case). Specifically, the Monomial-NFN pooling layer normalizes the weights across the hidden dimension and takes the average for rows (first layer), columns (last layer), or both (other layers). The output of this invariant Monomial-NFN layer is flattened and projected to \mathbb{R}^{200} (ReLU case) or \mathbb{R}^{1000} (Tanh case). This resulting vector is then passed through an MLP with two hidden layers with ReLU activations. The output is linearly projected to a scalar and then passed through a sigmoid function. We use the Binary Cross Entropy (BCE) loss function and train the model for 50 epochs, with early stopping based on τ on the validation set, which takes 35 minutes to train on an A100 GPU. The hyperparameters for our model are presented in Table 18.

Table 16: Datasets information for predicting generalization task.

Dataset	Train size	Val size
Original ReLU	6050	1513
Original Tanh	5949	1488
Augment ReLU	12100	3026
Augment Tanh	11898	2976

Table 17: Number of parameters of all models for predicting generalization task.

Model	ReLU dataset	Tanh dataset
STATNN	1.06M	1.06M
NP	2.03M	2.03M
HNP	2.81M	2.81M
Monomial-NFN (ours)	0.25M	1.41M

Table 18: Hyperparameters for Monomial-NFN on predicting generalization task.

	ReLU	Tanh
MLP hidden neurons	200	1000
Loss	Binary cross-entropy	Binary cross-entropy
Optimizer	Adam	Adam
Learning rate	0.001	0.001
Batch size	8	8
Epoch	50	50

Table 19: Dataset size for Classifying INRs task.

	Train	Validation	Test
CIFAR-10	45000	5000	10000
MNIST size	45000	5000	10000
Fashion-MNIST	45000	5000	20000

For the baseline models, we follow the original implementations described in [71], using the official code (available at: <https://github.com/AllanYangZhou/nfn>). For the HNP and NP models, there are 3 equivariant layers with 16, 16, and 5 channels, respectively. The features go through an average pooling layer and 3 MLP layers with 1000 hidden neurons. The hyperparameters of our model and the number of parameters for all models in this task can be found in Table 17.

D.4 Classifying implicit neural representations of images

Dataset. We utilize the original INRs dataset provided by [71], with no augmentation. The data is obtained by implementing a single SIREN model for each image in each dataset: CIFAR-10, MNIST, and Fashion-MNIST. The size of training, validation, and test samples for each dataset is provided in Table 19.

Implementation details. In these experiments, our general architecture includes 2 Monomial-NFN layers with sine activation, followed by 1 Monomial-NFN layer with absolute activation. The choice of hidden dimension in the Monomial-NFN layer depends on each dataset and is described in Table 20. The architecture then follows the same design as the NP and HNP models in [71], where a Gaussian Fourier Transformation is applied to encode the input with sine and cosine components, mapping from 1 dimension to 256 dimensions. If the base layer is NP, the features will go through IOSinusoidalEncoding, a positional encoding designed for the NP layer, with a maximum frequency of 10 and 6 frequency bands. After that, the features go through 3 HNP or NP layers with ReLU activation functions. Then, an average pooling is applied, and the output is flattened, and the resulting vector is passed through an MLP with two hidden layers, each containing 1000 units and ReLU activations. Finally, the output is linearly projected to a scalar. For the MNIST dataset, there is an additional Channel Dropout layer after the ReLU activation of each HNP layer and a Dropout layer after the ReLU activation of each MLP layer, both with a dropout rate of 0.1. We use the Binary Cross Entropy (BCE) loss function and train the model for 200,000 steps, which takes 1 hour and 35

Table 20: Hyperparameters of Monomial-NFN for each dataset in Classify INRs task.

	MNIST	Fashion-MNIST	CIFAR-10
Monomial-NFN hidden dimension	64	64	16
Base model	HNP	NP	HNP
Base model hidden dimension	256	256	256
MLP hidden neurons	1000	500	1000
Dropout	0.1	0	0
Learning rate	0.000075	0.0001	0.0001
Batch size	32	32	32
Step	200000	200000	200000
Loss	Binary cross-entropy	Binary cross-entropy	Binary cross-entropy

Table 21: Number of parameters of all models for classifying INRs task.

	CIFAR-10	MNIST	Fashion-MNIST
MLP	2M	2M	2M
NP	16M	15M	15M
HNP	42M	22M	22M
Monomial-NFN (ours)	16M	22M	20M

Table 22: Number of parameters of all models for Weight space style editing task.

Model	Number of parameters
MLP	4.5M
NP	4.1M
HNP	12.8M
Monomial-NFN (ours)	4.1M

Table 23: Hyperparameters for Monomial-NFN on weight space style editing task.

Name	Value
Monomial-NFN hidden dimension	16
NP dimension	128
Optimizer	Adam
Learning rate	0.001
Batch size	32
Steps	50000

minutes on an A100 GPU. For the baseline models, we follow the same architecture in [71], with minor modifications to the model hidden dimension, reducing it from 512 to 256 to avoid overfitting. We use a hidden dimension of 256 for all baseline models and our base model. The number of parameters of all models can be found in Table 21

D.5 Weight space style editing

Dataset. We use the same INRs dataset as used for classification task, which has the size of train, validation and test set described in Table 19.

Implementation details. In these experiments, our general architecture includes 2 Monomial-NFN layers with 16 hidden dimensions. The architecture then follows the same design as the NP model in [71], where a Gaussian Fourier Transformation with a mapping size of 256 is applied. After that, the features go through IOSinusoidalEncoding and then through 3 NP layers, each with 128 hidden dimensions and ReLU activation. Finally, the output goes through an NP layer to project into a scalar and a LearnedScale layer described in the Appendix of [71]. We use the Binary Cross Entropy (BCE) loss function and train the model for 50,000 steps, which takes 35 minutes on an A100 GPU. For the baseline models, we keep the same settings as the official implementation. Specifically, the HNP or NP model will have 3 layers, each with 128 hidden dimensions, followed by a ReLU activation. An NFN of the same type will be applied to map the output to 1 dimension and pass it through a LearnedScale layer. The number of parameters of all models can be found in Table 22. The detailed hyperparameters for our model can be found in Table 23.

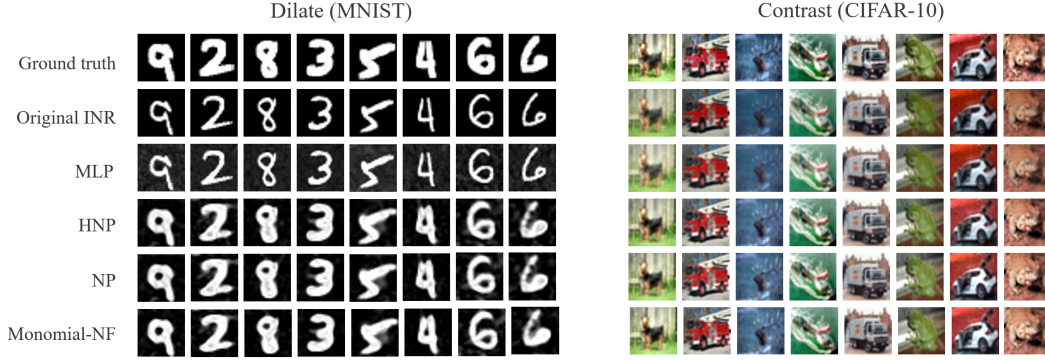


Figure 2: Random qualitative samples of INR editing behavior on the Dilate (MNIST) and Contrast (CIFAR-10) editing tasks.

D.6 Ablation Regarding Design Choices

We provide the ablation study on the choice of architecture for the task Predict CNN Generalization on ReLU subset in Table 24. We denote:

- Monomial Equivariant Functional Layer (Ours): MNF
- Activation: ReLU
- Scaling Invariant and Permutation Equivariant Layer (Ours): Norm
- Hidden Neuron Permutation Invariant Layer (in [71]): HNP
- Permutation Invariant Layer: Avg
- Multilayer Perceptron: MLP

Table 24: Ablation study on design choices for the task Predict CNN generalization on ReLU subset

	Original	1	2	3	4
$(\text{MNF-ReLU}) \times 1 \rightarrow \text{Norm} \rightarrow (\text{HNP-ReLU}) \times 1 \rightarrow \text{Avg} \rightarrow \text{MLP}$	0.917	0.916	0.917	0.917	0.917
$(\text{MNF-ReLU}) \times 2 \rightarrow \text{Norm} \rightarrow (\text{HNP-ReLU}) \times 1 \rightarrow \text{Avg} \rightarrow \text{MLP}$	0.918	0.917	0.917	0.917	0.918
$(\text{MNF-ReLU}) \times 3 \rightarrow \text{Norm} \rightarrow (\text{HNP-ReLU}) \times 1 \rightarrow \text{Avg} \rightarrow \text{MLP}$	0.920	0.919	0.918	0.920	0.920
$(\text{MNF-ReLU}) \times 1 \rightarrow \text{Norm} \rightarrow \text{Avg} \rightarrow \text{MLP}$	0.915	0.914	0.917	0.916	0.914
$(\text{MNF-ReLU}) \times 2 \rightarrow \text{Norm} \rightarrow \text{Avg} \rightarrow \text{MLP}$	0.918	0.919	0.918	0.917	0.918
$(\text{MNF-ReLU}) \times 3 \rightarrow \text{Norm} \rightarrow \text{Avg} \rightarrow \text{MLP}$	0.922	0.920	0.919	0.920	0.920

Among these designs, the architecture incorporating three layers of Monomial-NFN with ReLU activation achieves the best performance.

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification: The claims made in the abstract and introduction are clearly stated in the Contribution in the Section 1. These claims accurately reflect the paper's contributions and scope.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: The limitations are discussed in the Section 7.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory Assumptions and Proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [Yes]

Justification: All theoretical results in the paper are given together with the full set of assumptions and complete/correct proofs (See Appendix C.2 and Appendix C.1).

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental Result Reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: We provide the experiment details in the Implementation details section in the Appendix D of our manuscript. We also provide the source code so that the results in the paper can be easily reproduced.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
 - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: We provide the source code in the supplementary resources with detailed guide to run so that the results in the paper can be easily reproduced. We verify our proposed methods using public benchmarks (See the Section 6 in our manuscript)

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental Setting/Details

Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: We specify all the training and test details necessary to understand the results in the Implementation details section in the Appendix D of our manuscript.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment Statistical Significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [Yes]

Justification: We report error bars suitably and correctly defined of the experiments.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.

- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments Compute Resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: We provide sufficient information on the computer resources for all experiments in our Implementation details in Appendix D.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code Of Ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

Answer: [Yes]

Justification: The research conducted in the paper conforms, in every respect, with the NeurIPS Code of Ethics.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader Impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [Yes]

Justification: We discuss broader impacts in Appendix ??.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.

- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: The paper poses no such risks.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: We cite the githubs we use and the baselines we compare with in the Implementation details part in Appendix D of our manuscript.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.

- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New Assets**

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification: The paper does not release new assets.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and Research with Human Subjects**

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: The paper does not involve crowdsourcing or research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects**

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: The paper does not involve research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.

- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.