
Any-dimensional equivariant neural networks

Eitan Levin

California Institute of Technology

Mateo Díaz

Johns Hopkins University

Abstract

Traditional supervised learning aims to learn an unknown mapping by fitting a function to a set of input-output pairs with a *fixed dimension*. The fitted function is then defined on inputs of the same dimension. However, in many settings, the unknown mapping takes inputs in *any dimension*; examples include graph parameters defined on graphs of any size and physics quantities defined on an arbitrary number of particles. We leverage a newly-discovered phenomenon in algebraic topology, called representation stability, to define equivariant neural networks that can be trained with data in a fixed dimension and then extended to accept inputs in any dimension. Our approach is black-box and user-friendly, requiring only the network architecture and the groups for equivariance, and can be combined with any training procedure. We provide a simple open-source implementation of our methods and offer preliminary numerical experiments.

1 INTRODUCTION

Researchers are often interested in learning mappings that are defined for inputs of different sizes. We call such objects “*free*” as they are dimension-agnostic. Free objects are pervasive in a range of scientific and engineering fields. For example, in physics, quantities like electromagnetic fields are defined for any number of particles. In signal processing, problems such as deconvolution are well-defined for signals of any length, while image and text classification tasks are meaningful regardless of the input size. In mathematics, norms of vectors and matrices are defined for any dimension; sorting algorithms handle arrays of any length, and

graph parameters, such as the max-cut value, are defined for graphs of any size.

In contrast, most existing supervised learning methods aim to learn mappings by fitting a parametrized function $\hat{f}: \mathbb{R}^d \rightarrow \mathbb{R}^k$ to a set of input-output pairs $\mathcal{S} = \{(X_1, y_1), \dots, (X_n, y_n)\} \subseteq \mathbb{R}^d \times \mathbb{R}^k$ with fixed dimensions d and k . Naturally, the function \hat{f} is only defined on inputs of dimension d , and *a priori* cannot be applied to inputs in other dimensions. Practitioners use ad-hoc heuristics, such as downsampling, to resize the dimension of new inputs and match that of the learned map (Hashemi, 2019). The lack of systematic methodology for applying learned functions to inputs of different sizes motivates the main question of this work:

How can we learn mappings that accept inputs in any dimension?

We answer this question for mappings that are invariant or equivariant under the action of groups. Such mappings are ubiquitous since many application domains exhibit symmetry, e.g., physical quantities are equivariant under translations and rotations because they only depend on the relative position of particles (Villar et al., 2021), while graph parameters are equivariant under permutations because they only depend on the underlying topology rather than on its labelling (Chandrasekaran et al., 2012). We propose to learn such mappings using free equivariant neural networks. Formally, we consider a sequence of neural networks and a sequence of groups, such that a network at any level in the sequence is equivariant with respect to the group at the same level; see Figure 1. We train an equivariant network with data \mathcal{S} in a fixed dimension and leverage symmetry to extend it to other dimensions.

Core contributions. The question posed above presents three key challenges. First, how do we encode infinite sequences of neural networks, one in each dimension, using only finitely-many parameters? Doing so is necessary if we want to learn from a finite amount of data. Second, how do we ensure that the networks we learn from data in a fixed dimension generalize well to higher dimensions? Third, is there a user-friendly

procedure for learning these networks? We proceed to tackle these challenges.

Free equivariant neural networks. Equivariance is the key ingredient in tackling the first challenge. The authors of (Maron et al., 2019) showed that the dimension of permutation-equivariant linear layers between tensors is independent of the size of the tensors, and gave free bases for such linear layers. For example, the space of permutation-equivariant linear layers for graph adjacency matrices is 15-dimensional for any n , and its basis contains the identity, taking a transpose and extracting the diagonal, among others. The authors of that paper use this fact to derive finite parameterizations of linear layers

$$W = \alpha_1 A_1^{(n)} + \dots + \alpha_{15} A_{15}^{(n)} \quad (1)$$

where $\{A_1^{(n)}, \dots, A_{15}^{(n)}\}$ is a basis of free equivariant maps for graphs on n nodes. This parameterization allows them to instantiate linear layers in any dimension using only the fifteen parameters α_i , which can be learned from data in a fixed dimension. Our **first contribution** is showing that this is not a coincidence but rather stems from a general, recently-identified phenomenon known as *representation stability*. We utilize this phenomenon to show that the dimension of equivariant linear layers stabilizes for a number of group actions, including those induced by (signed) permutations, the orthogonal groups $O(n)$, rotations $SO(n)$, and the Lorentz groups $O(1, n)$. We leverage this observation to derive finite parametrizations of infinite sequences of equivariant neural networks.

Generalization across dimensions. The authors of (Maron et al., 2019) observed that neural networks obtained using (1) do not always generalize well to dimensions different from the one used for training. Intuitively, this is due to overparametrization, as there are many ways to represent the same function in a fixed dimension, but not all of them extend correctly to other dimensions. Our **second contribution** is to introduce a compatibility condition relating the mappings in different dimensions, which has a regularization effect that often leads to better generalization. Formally, it entails to the commutativity of the diagrams in Figure 1. We further explain how to impose this condition on the network architecture.

Computational recipe. The authors of (Maron et al., 2019) manually found the free basis elements in (1). However, for more complicated groups and spaces of linear layers, manually finding a free basis becomes prohibitive. In a fixed dimension, one can circumvent this issue using an algorithm proposed in (Finzi et al., 2021) to compute a basis. For our **third contribution**, we extend this basis to any other dimension by solving a sparse linear system, enabling us to obtain free bases

computationally.

We provide proof-of-concept experiments demonstrating the ability of our approach to learn functions defined for any dimension. We emphasize, however, that the primary contribution of this paper is its conceptual novelty. To our knowledge, our method is the first *black-box approach* for training neural networks that can handle inputs in any dimension. Instead of customizing architectures for specific domains and data types, as has been done in the prior work discussed below, our strategy adapts itself to the application domain by selecting appropriate group actions and nested sequences of vector spaces. The downside of the more general approach is that we do not expect it to outperform application-specific methods. Thus, we envision our framework serving as an initial step when approaching novel and poorly understood problems, or serving as a baseline to evaluate the performance of application-specific architectures.

Notation. Given two finite-dimensional vector spaces U and V we use $\mathcal{L}(V, U)$ to denote the set of linear mappings between them. For a group G acting on both U and V , we denote invariant elements as $U^G = \{u \in U \mid g \cdot u = u\}$, and equivariant linear maps as $\mathcal{L}(V, U)^G = \{A \in \mathcal{L}(V, U) \mid gAg^{-1} = A\}$. The map $\mathcal{P}_W: U \rightarrow U$ denotes the projection onto a subspace $W \subseteq U$. The symbols $\mathbb{1}_n$ and I_n represent the all-ones vector and the identity in dimension n , respectively. The symbol $\mathbb{R}[G]$ denotes the set of finite linear combinations of elements in G . To avoid cluttering notation, we use bold font symbols whenever possible to denote the n th element in a sequence, e.g., the symbol \mathbf{G} denotes $G^{(n)}$. We add a “+” superscript to denote the $(n+1)$ th element, e.g., $G^{(n+1)}$ is denoted by \mathbf{G}^+ .

Outline. The remainder of this section focuses on related work. Section 2 formally defines free neural networks. Section 3 introduces a compatibility condition that ensures good generalization across different dimensions. In Section 4, we describe our computational recipe to learn free neural networks from data in a fixed dimension and extend to other dimensions. Section 5 provides numerical experiments. We close the paper in Section 7 with conclusions, limitations, and future work. We defer several technical details to the appendix.

1.1 Related work

Equivariant learning. The benefits of equivariant architectures first became apparent with the success of convolutional neural networks (CNNs) in computer vision (Krizhevsky et al., 2017). Since then, equivariance has been applied to a range of applications. Examples include DeepSets (Zaheer et al., 2017) and graph neural

networks (Wu et al., 2021; Maron et al., 2019) using permutation equivariance to process sets and graphs; AlphaFold 2 (Jumper et al., 2021) and ARES (Townshend et al., 2021) using SE(3)-equivariance for protein and RNA structure prediction and assessment; steerable (Cohen and Welling, 2017) and spherical (Kondor et al., 2018) CNNs using rotation-equivariance to classify images; and physics-informed neural networks are equivariant under the symmetries of the corresponding physical systems (Karniadakis et al., 2021). Many of these architectures have been shown to implement a generalized notion of convolution over the groups at hand (Kondor and Trivedi, 2018; Cohen and Welling, 2016). Under additional assumptions, equivariant architectures have been derived using invariant theory to explicitly parametrize polynomial equivariant maps (Villar et al., 2021). We refer the reader to (Lim and Nelson, 2023; Bronstein et al., 2021) for an introduction to equivariant deep learning.

Dimension-free learning. Certain architectures processing data of a particular type are defined for inputs of different sizes. Many architectures processing graph-based data update the features at each vertex by applying the same function of the features of the vertex’s neighbors, hence can be applied to graphs of any size (Gilmer et al., 2017; Vignac et al., 2020; Wu et al., 2021). Convolutional neural networks (CNNs) processing signals and images convolve their inputs with filters of constant size, hence can also be applied to inputs of different sizes. Nevertheless, it has been observed that naïvely applying CNNs to inputs of size different from that used during training leads to artifacts, and several downsampling techniques have been proposed to resize inputs to a CNN (Hashemi, 2019). Networks processing natural language embed words as vectors of the same length but are defined for arbitrarily long sequences of such vectors. Recurrent neural networks process such sequences one-by-one, using cycles in their architectures to sequentially combine a new input vector with a function of the previous inputs (Mikolov et al., 2011). Recursive neural networks apply the same weights recursively to pairs of input vectors to combine them into one, until the entire sequence is reduced to one vector (Socher et al., 2011). Notably, all these architectures must process the input in particular ways motivated by the specific application at hand, whereas we only need to assume that the network processes its input equivariantly.

Representation stability. Representation stability considers nested sequences of groups and their representation, and implies that such sequences of representations often stabilize. Specifically, there is a labelling of the irreducibles of these groups such that the decomposition of the representations in the sequence contain the same irreducibles with the same multiplicities. This




Figure 1: Equivariant free neural networks. We use bold font to denote the n th element in a sequence; see Notation in Section 1 for details.

phenomenon has been formalized in (Church and Farb, 2013) and further studied in (Church et al., 2015; Wilson, 2014; Gadish, 2017; Sam and Snowden, 2015, 2016, 2017). It has been applied to study free convex sets and algebraic varieties (Levin and Chandrasekaran, 2023; Van Le and Römer, 2021; Draisma, 2014; Chiu et al., 2022; Alexandr et al., 2023), though to our knowledge, this paper is the first to apply it to equivariant deep learning. We refer the reader to (Farb, 2014; Wilson, 2018; Sam, 2020) for introductions to this area.

2 FREE NEURAL NETWORKS

In this section, we introduce the concept of free neural networks, i.e., networks that can be instantiated in every dimension. To set the stage, let us recall the classical notion of a neural network (NN). A NN is a mapping $f = f_L \circ \dots \circ f_1$ where $f_i: V_i \rightarrow V_{i+1}$ is a composition $f_i(x) = \sigma_i(W_i x + b_i)$ of an affine map $x \mapsto W_i x + b_i$ and an activation map $\sigma_i: U_i \rightarrow V_{i+1}$. This yields a family of mappings parametrized by the weights $W_i \in \mathcal{L}(V_i, U_i)$ and biases $b_i \in U_i$. In turn, these parameters are chosen to minimize a prescribed loss function.

In several settings, we want f to respect symmetries in its inputs. Formally, we require f to be equivariant with respect to the action of a group G , i.e., $f \circ g = g \circ f$ for all $g \in G$. A natural way to obtain equivariance is by ensuring that each building block of f is equivariant, i.e., $W_i \in \mathcal{L}(V_i, U_i)^G$, $b_i \in U_i^G$, and σ_i is G -equivariant. NNs satisfying these properties are called *equivariant*. Equivariant NNs can be trained by computing a basis for the space of weights and biases and optimizing over the coefficients in this basis (Finzi et al., 2021).

In this work, we seek equivariant NNs that extend to inputs in *any dimension*. To this end, we consider sequences $\{U_i\}$ and $\{V_i\}$ of nested¹ vector spaces $U_i \subseteq U_i^+$ and $V_i \subseteq V_i^+$, with actions of a nested sequence of groups $\{G\}$, see Figure 1. For example, we have inclusions $\mathbb{R}^n \subseteq \mathbb{R}^{n+1}$ by zero-padding on which the group of permutations on n letters $G = S_n$ acts by

¹Formally, we have embeddings $V \hookrightarrow V^+$ and $U \hookrightarrow U^+$.

permuting coordinates. These define a sequence of NNs whose weights $\mathbf{W}_i: \mathbf{V}_i \rightarrow \mathbf{U}_i$ and biases $\mathbf{b}_i \in \mathbf{U}_i$ increase in size. The activation functions σ are often defined for every dimension, as the next section shows. Thus, we focus here on extending the linear layers.

Free equivariant networks. Because the dimensions of the space of linear layers $\mathcal{L}(\mathbf{U}_i, \mathbf{V}_i)$ and the vector spaces \mathbf{U}_i usually grow, sequences of *general* NNs are not finitely-parameterizable and cannot be learned from a finite amount of data. The situation radically simplifies when considering equivariant networks since the dimensions of $\mathcal{L}(\mathbf{U}_i, \mathbf{V}_i)^{\mathbf{G}}$ and $\mathbf{U}_i^{\mathbf{G}}$ often stabilize. This was previously observed in the context of simple graph NNs by (Maron et al., 2019), and follows from a general phenomenon known as representation stability (Church and Farb, 2013). To illustrate it with a concrete example, consider the nested sequence $\mathbf{U} = \mathbf{V} = \mathbb{R}^n$ with the action of $\mathbf{G} = \mathbb{S}_n$ as above. Then $\mathbf{V}^{\mathbf{G}}$ is one-dimensional and spanned by $\mathbb{1}_n$ while $\mathcal{L}(\mathbf{U}, \mathbf{V})^{\mathbf{G}}$ is two-dimensional and spanned by $I_n, \mathbb{1}_n \mathbb{1}_n^{\top}$, which are basis elements defined for every n . Similarly, (Maron et al., 2019) obtained explicit free bases for spaces of invariant tensors.

Interestingly, all the above basis elements project onto each other, e.g., the orthogonal projection of $\mathbb{1}_n$ onto \mathbb{R}^{n-1} is $\mathbb{1}_{n-1}$, and similarly $I_n, \mathbb{1}_n \mathbb{1}_n^{\top}$ project onto $I_{n-1}, \mathbb{1}_{n-1} \mathbb{1}_{n-1}^{\top}$. Motivated by this observation, we say that a sequence of equivariant NNs is *free* if the weights $\mathbf{W}_i: \mathbf{V}_i \rightarrow \mathbf{U}_i$ and biases $\mathbf{b}_i \in \mathbf{U}_i$ satisfy

$$\mathbf{W}_i = \mathcal{P}_{\mathbf{U}_i} \mathbf{W}_i^+ |_{\mathbf{V}_i}, \quad \text{and} \quad \mathbf{b}_i = \mathcal{P}_{\mathbf{U}_i} \mathbf{b}_i^+. \quad (\text{FreeNN})$$

For many sequences $\{\mathbf{V}_i\}, \{\mathbf{U}_i\}$, equation (FreeNN) uniquely determines $\mathbf{W}_i^+, \mathbf{b}_i^+$ from $\mathbf{W}_i, \mathbf{b}_i$, allowing us to uniquely extend a network to accept larger inputs.

Representation stability. To understand for which sequences (FreeNN) allows us to extend, we introduce some key definitions from the representation stability literature (Church and Farb, 2013; Church et al., 2015).

Definition 2.1 (Consistent sequences). Fix a sequence of compact groups $\mathcal{G} = \{\mathbf{G}\}_{n \in \mathbb{N}}$ such that $\mathbf{G} \subseteq \mathbf{G}^+$. The family $\mathcal{V} = \{(\mathbf{V}, \varphi)\}_{n \in \mathbb{N}}$ is a *consistent sequence* of \mathcal{G} -representations if the following hold true for all n :

- (a) (**Representations**) The set \mathbf{V} is an orthogonal \mathbf{G} -representation;
- (b) (**Equivariant isometries**) The map $\varphi: \mathbf{V} \hookrightarrow \mathbf{V}^+$ is a \mathbf{G} -equivariant isometry.

We will identify \mathbf{V} with its image $\varphi(\mathbf{V}) \subseteq \mathbf{V}^+$, and omit the inclusions φ unless needed. Importantly, consistent sequences can be combined to form more complex sequences. In particular, if $\mathcal{V} = \{(\mathbf{V}, \varphi)\}$

and $\mathcal{U} = \{(\mathbf{U}, \psi)\}$ are consistent sequences of $\{\mathbf{G}\}$ -representations, then so are their sum and tensor product:

$$\mathcal{V} \oplus \mathcal{U} = \{(\mathbf{V} \oplus \mathbf{U}, \varphi \oplus \psi)\}, \quad \mathcal{V} \otimes \mathcal{U} = \{(\mathbf{V} \otimes \mathbf{U}, \varphi \otimes \psi)\}.$$

This follows directly from Definition 2.1. We use $\mathcal{V}^{\oplus k}$ (resp. $\mathcal{V}^{\otimes k}$) to denote the direct sum (resp., tensor product) of \mathcal{V} taken k times. As a direct by-product of this observation, we obtain that the sequence of linear maps $\mathcal{L}(\mathcal{V}, \mathcal{U}) = \{(\mathcal{L}(\mathbf{V}, \mathbf{U}), \varphi \otimes \psi)\}$ is consistent since it is isomorphic to $\mathcal{V} \otimes \mathcal{U}$. The following parameter controls the complexity of consistent sequences and ensures that their spaces of invariants are eventually isomorphic.

Definition 2.2 (Generation degree). A consistent sequence $\{\mathbf{V}\}$ is *generated in degree d* if $\mathbb{R}[\mathbf{G}]V^{(d)} = \mathbf{V}$ for all $n \geq d$. The generation degree is the smallest such $d \in \mathbb{N}$. The sequence is *finitely-generated* if the generation degree is finite.

In words, a consistent sequence $\{\mathbf{V}\}$ is generated in degree d if for all $n \geq d$, \mathbf{V} is equal to the span of linear combinations of elements of \mathbf{G} applied to $V^{(d)}$. For instance, if $\mathbf{V} = \mathbb{R}^d$ with the action of permutations $\mathbf{G} = \mathbb{S}_n$, then its generation degree is one, as any $v \in \mathbf{V}$ can be obtained from $e_1 = (1, 0, \dots, 0) \in V^{(1)} \subseteq \mathbf{V}$ via $v = \sum v_i g_i \cdot e_1$ where g_i swaps the first and i th components. The next proposition gives an isomorphism between spaces of invariants in a finitely-generated consistent sequence. This result was first established in the representation stability literature, albeit in a different form. We include proof for completion.

Proposition 2.3 (Isomorphism of invariants). *If $\{\mathbf{V}\}$ is generated in degree d , then the orthogonal projections $\mathcal{P}_{\mathbf{V}}: (\mathbf{V}^+)^{\mathbf{G}^+} \rightarrow \mathbf{V}^{\mathbf{G}}$ are injective for all $n \geq d$, and isomorphisms for all large n .*

Proof. First, the projection $\mathcal{P}_{\mathbf{V}}$ is \mathbf{G} -equivariant because \mathbf{G} acts orthogonally. Second, it maps \mathbf{G}^+ -invariants in \mathbf{V}^+ to \mathbf{G} -invariants in \mathbf{V} because $\mathbf{G} \subseteq \mathbf{G}^+$. Third, we prove injectivity for $n \geq d$. Suppose $\mathcal{P}_{\mathbf{V}}v = 0$ for some $v \in (\mathbf{V}^+)^{\mathbf{G}^+}$. For any $u \in \mathbf{V}^+$, write $u = \sum_i g_i u_i$ where $u_i \in \mathbf{V}$ and $g_i \in \mathbf{G}^+$. We then have $\langle v, u \rangle = \langle v, \sum_i u_i \rangle = 0$ because v is invariant and $\sum_i u_i \in \mathbf{V}$. Since $u \in \mathbf{V}^+$ was arbitrary, we conclude that $v = 0$. The injectivity of $\mathcal{P}_{\mathbf{V}}$ shows that $\dim(\mathbf{V})^{\mathbf{G}} \geq \dim(\mathbf{V}^+)^{\mathbf{G}^+}$ for all $n \geq d$, hence the sequence of dimensions $\dim(\mathbf{V})^{\mathbf{G}}$ eventually stabilizes, at which point the projections become isomorphisms. \square

Note that (FreeNN) precisely requires \mathbf{W}^+ and \mathbf{b}^+ to project to \mathbf{W}, \mathbf{b} inside the consistent sequences $\mathcal{L}(\mathcal{V}_i, \mathcal{U}_i), \mathcal{U}_i$, respectively. Thus, Proposition 2.3 implies that for all large n , we can parameterize infinite sequences of equivariant linear layers with a finite set

of parameters $\alpha_1, \dots, \alpha_\ell$, just as in (1). This result enables us to train free NNs as we describe in Section 4.

Examples. We collect a few crucial examples and include additional ones in the appendix.

Scalar sequence. The sequence $\mathcal{S} = \{\mathbb{R}\}$ together with $\varphi(x) = x$ yields a consistent sequence for any sequence of groups acting trivially, i.e., $g \cdot x = x$. It is generated in degree one. By setting the last layer $\mathbf{V}_{L+1} = \mathbb{R}$ in the free NN architecture of Figure 1, we obtain a sequence of invariant functions \mathbf{f} .

Permutation sequences. Let $\mathcal{R} = \{\mathbb{R}^n\}$ be the consistent sequence with zero-padding embeddings $\varphi(x) = (x^\top, 0)^\top$ and the action of the symmetric group $\mathbf{G} = \mathbb{S}_n$ by permuting coordinates. Consider $\mathcal{V}_k = \mathcal{U}_{k-1} = \mathcal{R}^{\otimes k}$, which is generated in degree k . The dimensions of the space of weights and biases for low-order tensors are eventually

$$\begin{aligned} \dim \mathcal{L}(\mathbf{V}_1, \mathbf{U}_1)^{\mathbf{G}} &= 4, & \dim \mathbf{U}_1^{\mathbf{G}} &= 2, \\ \dim \mathcal{L}(\mathbf{V}_2, \mathbf{U}_2)^{\mathbf{G}} &= 52, & \dim \mathbf{U}_2^{\mathbf{G}} &= 5. \end{aligned}$$

Graph(on) sequences. There are two ways to identify small graphs with larger ones, namely, appending isolated vertices and identifying graphs with their associated graphons. These two embeddings yield two consistent sequences. For the first embedding, let $\mathbf{V} = \mathbb{S}^n$ with the Frobenius inner product and embeddings $\varphi(X) = \text{blkdiag}(X, 0)$ padding X by a zero row and zero column. We endow \mathbf{V} with the action of $\mathbf{G} = \mathbb{S}_n$ by simultaneously permuting rows and columns (so $g \cdot X = gXg^\top$ for all $g \in \mathbf{G}$ and $X \in \mathbf{V}$). This sequence is generated in degree 2. For the second embedding, let $\mathbf{V} = \mathbb{S}^{2^n}$ with the normalized inner product $\langle X, Y \rangle = 2^{-2^n} \text{Tr}(XY)$, embeddings $\varphi(X) = X \otimes \mathbb{1}_{2 \times 2}$, and the same action of the symmetric group $\mathbf{G} = \mathbb{S}_{2^n}$ as above. This sequence arises in the theory of graphons Lovász (2012), where adjacency matrices X are associated with a step function (the graphon) $W_X : [0, 1]^2 \rightarrow \mathbb{R}$ in such a way that X and $X \otimes \mathbb{1}_{2 \times 2}$ represent the same graphon and $\langle X, Y \rangle = \int_{[0, 1]^2} W_X(x, y) W_Y(x, y) dx dy$. The graphon consistent sequence is generated in degree 2 by (Levin and Chandrasekaran, 2023, Prop. 3.1).

When does stability kick in? Understanding the exact level at which the projections become isomorphisms is important, yet, an exact characterization is rather technical. Stabilization occurs at the so-called **presentation degree** of the sequence, which might be larger than the generation degree in general. However, they agree for many relevant examples, such as in the sequence $\mathcal{R}^{\otimes k}$ above. We include a formal definition and a discussion in Appendix B.

3 GENERALIZATION ACROSS DIMENSIONS AND COMPATIBILITY

Free NNs often do not generalize correctly to other dimensions. In Section 5, we provide examples where a free NN yields excellent test error in its trained dimension while exhibiting errors 10^{10} times worst in other dimensions. This discrepancy arises due to over-parametrization, as there are numerous ways to encode the same function in a fixed dimension, and yet not every encoding extends seamlessly. In this section, we propose a regularization strategy that often leads to better generalization across dimensions.

Specifically, inspired by a similar condition for free convex sets (Levin and Chandrasekaran, 2023), we introduce the following compatibility condition.

Definition 3.1 (Compatible networks). A sequence of maps $\{\mathbf{f}: \mathbf{V} \rightarrow \mathbf{U}\}$ is (intersection-) *compatible* if $\mathbf{f}^+|_{\mathbf{V}} = \mathbf{f}$. A sequence of equivariant networks (Figure 1) is compatible if

$$\mathbf{W}_i^+|_{\mathbf{V}_i} = \mathbf{W}_i, \quad \mathbf{b}_i^+ = \mathbf{b}_i, \quad \text{and} \quad \sigma_i^+|_{\mathbf{U}_i} = \sigma_i. \quad (\text{CompNN})$$

Intuitively, this condition ensures that applying a function in a fixed dimension and then extending to a higher dimension is equivalent to first extending and then applying the function. Layer-wise compatibility (CompNN) guarantees that the sequence of NN $\{\mathbf{f}: \mathbf{V}_1 \rightarrow \mathbf{V}_{L+1}\}$, given by Figure 1, is compatible, or equivalently, that the diagram in that figure commutes. The condition (CompNN) is strictly stronger than (FreeNN), as compatible NNs are free, but free NNs are not compatible in general.

Compatibility in the wild. Compatibility is a natural condition satisfied by sequences of equivariant maps $\{\mathbf{f}\}$ arising in many problems of interest. Here we list some examples.

Graph(on) parameters. Graph parameters are functions of graphs (typically of any size) that do not depend on the labeling of their vertices. In other words, these are sequences of \mathbb{S}_n -invariant functions $\{\mathbf{f}: \mathbb{S}^n \rightarrow \mathbb{R}\}$. Several graph invariants are also compatible with zero-padding, which is equivalent to adding an isolated vertex. For example, the max-cut value, the number of triangles, and cycles do not change if we add an isolated vertex, hence they are compatible. Other graph invariants only depend on a graph via its associated graphon, and hence are compatible with the embeddings in the graphon sequence (see example at the end of Section 2). These include graph homomorphism densities, which play a central role in extremal combinatorics Lovász (2012).

Matrix mappings. Linear algebra operations are often compatible. For instance, the experiments in (Maron et al., 2019) aimed to learn the following matrix mappings: $X \mapsto \frac{1}{2}(X+X^\top)$, $X \mapsto \text{diag}(\text{diag}(X))$, $X \mapsto \text{tr}(X)$, and $X \mapsto \text{argmax}_{\|v\|=1} \|Xv\|$. All of these are compatible sequences with the zero-padding embedding described in Section 2, and S_n -equivariant.

Orthogonal invariance. The papers (Finzi et al., 2021; Villar et al., 2021) consider the task of learning the function

$$\mathbf{f}(x_1, x_2) = \sin(\|x_1\|) - \frac{\|x_2\|^3}{2} + \frac{x_1^\top x_2}{\|x_1\| \|x_2\|}, \quad (2)$$

defined on $(\mathbb{R}^n)^{\oplus 2}$ with $n = 5$ from evaluation data. This function is $O(n)$ -invariant if we let rotations g act by $(g \cdot x_1, g \cdot x_2)$. Embedding $(\mathbb{R}^n)^{\oplus 2}$ into $(\mathbb{R}^{n+1})^{\oplus 2}$ by zero-padding each vector, we see that $\{\mathbf{f}\}$ is a compatible sequence of invariant functions.

Orthogonal equivariance. The $O(3)$ -equivariant task in (Finzi et al., 2021; Villar et al., 2021) consists of taking as input $n = 5$ particles in space, given by their masses and position vectors $(m_i, x_i)_{i=1}^n \in (\mathbb{R} \oplus \mathbb{R}^3)^{\oplus n}$, and outputting their moment of inertia matrix $\mathbf{f}(m_i, x_i) = \sum_{i=1}^n m_i(x_i^\top x_i I_3 - x_i x_i^\top)$. Embedding $\mathbf{V} = (\mathbb{R} \oplus \mathbb{R}^3)^{\oplus n}$ into \mathbf{V}^+ by zero-padding and letting $\mathbf{U} = \mathbb{S}^3$, we see that $\{\mathbf{f}\}$ is a compatible sequence of maps. Furthermore, let $\mathbf{G} = S_n \times O(3)$ act on \mathbf{V} by $(\pi, g) \cdot (m_i, x_i)_{i=1}^n = (m_{\pi^{-1}(i)}, g x_{\pi^{-1}(i)})_{i=1}^n$, and on \mathbf{U} by $(\pi, g) \cdot X = g X g^{-1}$ for $(\pi, g) \in \mathbf{G}$ (so permutations act trivially). Then each \mathbf{f} is \mathbf{G} -equivariant.

There are many more examples of compatible mappings in the literature (Levin and Chandrasekaran, 2023). Therefore, we proceed to study compatible NNs. We derive conditions ensuring the compatibility of the linear layers and show that several standard activation functions are compatible.

Compatible linear layers. Since we only have data in some finite level n_0 , we ask when do fixed-dimensional weights $W_i^{(n_0)}$ extend to a sequence satisfying (CompNN). The set of such $W_i^{(n_0)}$ forms a linear space, and the next theorem characterizes a subspace of it and enables us to find a basis for that subspace in Section 4. We defer its proof to Appendix C.

Assumption 1. The sequences $\mathcal{V} = \{\mathbf{V}\}$, $\mathcal{U} = \{\mathbf{U}\}$ are obtained from direct sums and tensor products of the same sequence \mathcal{V}_0 . The sequence \mathcal{V} is generated in degree d_g and presented in degree d_p . The mapping $W^{(n_0)} \in \mathcal{L}(V^{(n_0)}, U^{(n_0)})^{G^{(n_0)}}$ at level $n_0 \geq d_p$ satisfies

$$W^{(n_0)}(V^{(m)}) \subseteq U^{(m)} \text{ for } m \leq d_g.$$

Theorem 3.2. *Suppose that the sequences $\mathcal{V} = \{\mathbf{V}\}$, $\mathcal{U} = \{\mathbf{U}\}$ and linear map $W^{(n_0)} \in$*

$\mathcal{L}(V^{(n_0)}, U^{(n_0)})^{G^{(n_0)}}$ satisfy Assumption 1. Then there is a unique extension $\{\mathbf{W}\}$ of $W^{(n_0)}$ to a sequence of equivariant linear maps satisfying (CompNN).

Recall that we define the presentation degree in the appendix. In words, Theorem 3.2 says that if $W^{(n_0)}$ is equivariant and its restrictions to lower-dimensional subspaces satisfy (CompNN), then it uniquely extends to higher-dimensional weights satisfying (CompNN). In Section 4, we use this result to find a basis for such $W^{(n_0)}$, and use it to train compatible equivariant NNs.

Compatible activation functions. The majority of equivariant nonlinearities proposed in the literature are compatible, such as the following few examples.

Entrywise activations and permutations. Let $\sigma: \mathbb{R} \rightarrow \mathbb{R}$ be any nonlinear function. Define $\boldsymbol{\sigma}: (\mathbb{R}^n)^{\otimes k} \rightarrow (\mathbb{R}^n)^{\otimes k}$ by applying σ to each entry of a tensor. Then each $\boldsymbol{\sigma}$ is $\mathbf{G} = S_n$ -equivariant, and $\{\boldsymbol{\sigma}\}$ is a compatible sequence with respect to the zero-padding embeddings if and only if $\sigma(0) = 0$. On the other hand, if $\boldsymbol{\sigma}: (\mathbb{S}^{2^n})^{\otimes k} \rightarrow (\mathbb{S}^{2^n})^{\otimes k}$ is again applied entrywise and we use the graphon embedding (see example at the end of Section 2), then $\{\boldsymbol{\sigma}\}$ is a compatible sequence for any σ .

Bilinear layers. The authors of (Finzi et al., 2021) map $(\mathbf{V} \otimes \mathbf{U}) \oplus \mathbf{V} \rightarrow \mathbf{U}$ by sending $(v \otimes u, v') \mapsto \langle v, v' \rangle u$, which is equivariant for any group action on \mathbf{V}, \mathbf{U} . This sequence of maps is also compatible.

Gated nonlinearities. The authors of (Weiler et al., 2018) map $\mathbf{V} \oplus \mathbb{R} \rightarrow \mathbf{V}$ by $(v, \alpha) \mapsto v\sigma(\alpha)$ where $\sigma: \mathbb{R} \rightarrow \mathbb{R}$ is a nonlinearity. These are equivariant maps for any group acting trivially on the \mathbb{R} component, and they form a compatible sequence.

4 A COMPUTATIONAL RECIPE FOR LEARNING FREE NEURAL NETWORKS

In this section, we describe an algorithm to train free and compatible NNs. We do so in two stages. First, we train our NN at a large enough level n_0 by finding bases for the weights and biases at that level and optimizing over the coefficients in this basis. Second, for any higher level $n > n_0$, we extend the trained NN at level n_0 by solving linear systems for the higher-dimensional weights and biases; for any lower level $n < n_0$ we project the NN using (FreeNN). Throughout, we fix consistent sequences $\mathcal{V}_i, \mathcal{U}_i$ and sequences of compatible equivariant nonlinearities $\{\boldsymbol{\sigma}_i\}$. We summarize our procedure in Algorithm 1. In Appendix D, we show that our procedure provably generates free and compatible NN.

Finding bases for free NN. To train free networks (**FreeNN**), we fix n_0 *exceeding the presentation degrees* of $\mathcal{V}_i \otimes \mathcal{U}_i$ and \mathcal{U}_i for all $i = 1, \dots, L$, and find a basis for equivariant weights and invariant biases at level n_0 using the algorithm of (Finzi et al., 2021). Specifically, Theorem 1 in (Finzi et al., 2021) states that $b \in (U_i^{(n_0)})^{G^{(n_0)}}$ if and only if b satisfies that

$$\begin{aligned} Bb &= 0 && \left(\text{for all } B \in \mathcal{B}_i^{(n_0)}\right), \\ (D - I)b &= 0 && \left(\text{for all } D \in \mathcal{D}_i^{(n_0)}\right), \end{aligned} \quad (3)$$

where $\mathcal{B}_i^{(n_0)}$ is a basis for the Lie algebra of $G^{(n_0)}$ and $\mathcal{D}_i^{(n_0)}$ are discrete generators for $G^{(n_0)}$, which are finite sets. Here B and D are represented as matrices acting on $U_i^{(n_0)}$. Since equivariant linear maps $W: V_i^{(n_0)} \rightarrow U_i^{(n_0)}$ are just the invariants in $\mathcal{L}(V_i^{(n_0)}, U_i^{(n_0)})$, they constitute the kernel of the analogously-defined set of equations. Thus, finding a basis for equivariant weights and invariant biases reduces to finding a basis for the kernel of a matrix (3), which is typically large and sparse. This can be tackled either using a Krylov-subspace method (Finzi et al., 2021) or a sparse LU decomposition (Gotsman and Toledo, 2008; Kowal, 2006).

Finding bases for Compatible NN. If we rather want to find basis for layers satisfying (**CompNN**), we fix n_0 *exceeding the presentation degrees* of \mathcal{V}_i and let d_i be its generation degree. For most representations, (**CompNN**) implies zero bias, so we focus on the weights. We find a basis for weights $W_i^{(n_0)}$ satisfying the hypotheses of Theorem 3.2 by noting that Assumption 1 holds if and only if $W_i^{(n_0)}$ is $G^{(n_0)}$ -equivariant and satisfies

$$\left[\mathcal{P}_{V_i^{(m)}} \otimes (I - \mathcal{P}_{U_i^{(m)}}) \right] \text{vec} \left(W_i^{(n_0)} \right) = 0,$$

for all $m \leq d_i$, where vec denotes the vectorization of a matrix by stacking its columns. This characterization allows us to find a basis for weights by finding a basis for the kernel of a sparse matrix.

Extending to arbitrary dimensions. For any $n > n_0$, we extend our trained network at level n_0 by finding the unique equivariant weights and biases at level n projecting onto the trained ones as in (**FreeNN**). Formally, we find the unique $W_i^{(n)}$ and $b_i^{(n)}$ satisfying (3) with n_0 replaced by n , and

$$\begin{aligned} \left[\mathcal{P}_{V_i^{(n_0)}} \otimes \mathcal{P}_{U_i^{(n_0)}} \right] \text{vec} \left(W_i^{(n)} \right) &= \text{vec} \left(W_i^{(n_0)} \right), \\ \mathcal{P}_{U_i^{(n_0)}} b_i^{(n)} &= b_i^{(n_0)}. \end{aligned} \quad (\text{ExtSyst})$$

This amounts to solving a linear system, which is again typically sparse. It can be solved via iterative methods, e.g., stochastic gradient descent or LSQR (Paige and

Algorithm 1 Train a freely-described (resp., compatible) neural network.

Input: Architecture $\mathcal{V}_i, \mathcal{U}_i$, group sequence $\{\mathbf{G}\}$, and training data at level n_0 .

Stage 1 Find free (resp., compatible) bases of the linear layers at level n_0 .

Optimize over the coefficients in the bases to train the network at level n_0 .

Stage 2 For any $n > n_0$, solve (**ExtSyst**) to extend the trained network to level n .

For any $n < n_0$, project via (**FreeNN**) to restrict the trained network to level n .

Saunders, 1982). For any $n < n_0$, we set $W_i^{(n)}$ and $b_i^{(n)}$ via (**FreeNN**), which is equivalent to (**ExtSyst**) with n and n_0 swapped.

5 NUMERICAL EXPERIMENTS

In this section, we use Algorithm 1 to learn some of the compatible mappings from Section 3. Our implementation is based on that of (Finzi et al., 2021), and is available at

<https://github.com/mateodd25/free-nets>.

For each experiment, we aim to learn a sequence of mappings $\{f: \mathbf{V}_1 \rightarrow \mathbf{V}_{L+1}\}$. To do so, we fix a level n_0 and randomly generate data $\{(X_i, f^{(n_0)}(X_i))\} \subseteq V_1^{(n_0)} \times V_{L+1}^{(n_0)}$. To fit the data, we consider the architectures described in the Appendix E and compute bases for weights and biases as we describe in Section 4. We then optimize over the coefficients in those bases using ADAM (Kingma and Ba, 2015) starting from random initialization. Finally, we extend the trained network at level n_0 to several levels n . We evaluate our extended network on random test data generated at each level from the ground-truth sequence of maps.

We consider five examples from Section 3: the trace $f(X) = \text{tr}(X)$, diagonal extraction $f(x) = \text{diag}(\text{diag}(X))$, symmetric projection $f(X) = (X + X^\top)/2$, the top right-singular vector $f(X) = \text{argmax}_{\|v\|=1} \|Xv\|$, and the function defined in (2). The first and last examples are invariant functions; the rest are equivariant mappings. For the first four examples, we set $\mathbf{G} = S_n$, and for the last one, we set $\mathbf{G} = O(n)$. We train with 3000 data points generated randomly and evaluate the test error at each dimension on 1000 fresh samples.

We use the mean squared error (MSE) as our loss function for all the experiments except for learning the top singular vector. For that experiment, we use the

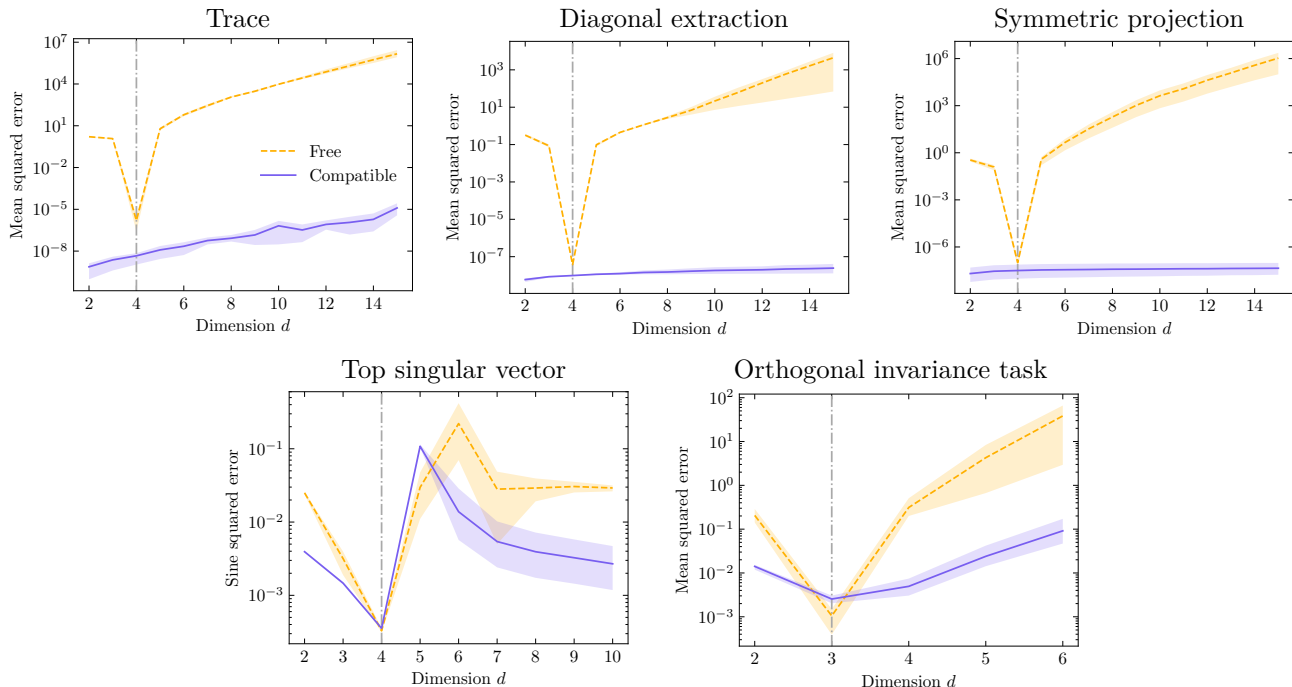


Figure 2: Test errors across dimensions for free and compatible networks. Each experiment is run three times; the lighter bands show the max and min runs, while the bold line shows the average. For diagonal extraction and symmetric projection, we measure the average MSE per entry. Vertical gray lines mark the dimension used for learning.

squared sine loss proposed by (Maron et al., 2019), i.e., $\ell(\hat{y}, y) = 1 - \langle \hat{y}, y \rangle^2 / \|\hat{y}\|^2 \|y\|^2$. The resulting errors in each dimension for both free and compatible NNs are shown in Figure 2. The test errors at the trained dimension are competitive with those obtained in (Maron et al., 2019; Finzi et al., 2021; Villar et al., 2021) for learning the same mappings, see Table 1 for comparison. Moreover, the errors we obtain for higher-dimensional extensions are competitive with errors obtained in the literature for NNs trained at those higher dimensions. For the orthogonally invariant task example, we trained our compatible NN in 3 dimensions using 3000 data points and got an error of $2e-2$ when extending to 5 dimensions. In contrast, the NN of (Villar et al., 2021), which is tailored to orthogonal equivariance, achieves an error of $9e-3$ when trained in 5 dimensions on the same number of points, while the NN of (Finzi et al., 2021) trained in the same dimension achieves $4e-2$.

Due to memory limitations, we set the training level to $n_0 = 3$ for the orthogonal invariance task, whereas Algorithm 1 would require $n_0 = 6$ to uniquely extend free networks² since the presentation degree of the linear maps between the hidden layers equal to six—see Theorem D.1 in the appendix. This highlights an advantage of imposing compatibility—it allows

²We find the minimum norm solution for the weights and biases using LSQR (Paige and Saunders, 1982).

Example	Ours	Previous
trace	$3e-8$	$1e-3$ (Maron et al., 2019)
diag	$1e-8$	$7e-6$ (Maron et al., 2019)
sym	$3e-8$	$7e-6$ (Maron et al., 2019)
svd	$4e-4$	$1.6e-3$ (Maron et al., 2019)
orth.	$2e-3$	$1e-2$ (Finzi et al., 2021); $5e-4$ (Villar et al., 2021)

Table 1: Comparison of errors from the literature

us to uniquely extend a trained network from lower-dimensional data. Moreover, we see that imposing our compatibility condition yields substantially lower errors across dimensions. Remarkably, the test error for free NNs increases by many orders of magnitude for simple functions, a phenomenon that was previously noted by (Maron et al., 2019). This further underscores the importance of compatibility when generalizing to other dimensions.

6 LIMITATIONS

Extension to non-compact groups. We assumed that the group is compact and acts orthogonally so that

the orthogonal projection $\mathcal{P}_V = \varphi^*$ is \mathbf{G} -equivariant, used for instance in the proof of Proposition 2.3. This property holds more generally when \mathbf{G} is linearly reductive. Thus, our framework extends to non-compact groups such as the Lorentz group $O(3, 1)$ and so-called rescaling groups $\mathbb{R}_{>0}^n$ consisting of diagonal matrices with strictly positive entries, both of which are important to applications of equivariant deep learning to physics Bogatskiy et al. (2020); Villar et al. (2023).

Some groups arising in applications are not reductive however, including any group containing translations such as the group of rigid motions $SE(n)$. This difficulty can be circumvented in some cases, such as centering the data when learning an $SE(n)$ -invariant function and imposing only $SO(n)$ -invariance on the architecture. In general, however, more work is needed to address the non-reductive case.

Compatible activation functions and biases.

Our compatibility conditions (**CompNN**) require the activation functions we use to be compatible with our embeddings relating the different dimensions. For example, applying any activation function $\sigma: \mathbb{R} \rightarrow \mathbb{R}$ entrywise always satisfies (**CompNN**) for the graphon embeddings. However, only activations with $\sigma(0) = 0$ satisfy (**CompNN**) for the zero-padding embeddings. On the one hand, some activation functions used in practice, such as the sigmoid $\sigma(x) = (1 + e^{-x})^{-1}$, do not satisfy this condition. Moreover, compatibility requires the biases to be the same across dimensions, which implies that the biases must be multiples of the first canonical basis vector when using the zero-padding embedding, and a multiple of the all-1’s matrix when using the graphon embedding. Therefore, compatibility limits the architectures that we can use. On the other hand, if the sequence of mappings we are trying to learn is known in advance to satisfy compatibility, then it is desirable to use architectures that enforce compatibility as well. In that case the mapping we learn also generalizes correctly to other dimensions, as the numerical experiments in Section 5 demonstrate.

Scalability. We hope to address several limitations of our approach in future work. First, while our compatibility condition improves generalization, it can be restrictive in certain settings. For example, it often yields zero bias. It would be interesting to learn the relation between maps in different dimensions directly from the data rather than assume it is known in advance. Second, for examples like the leading eigenvector, the error increases substantially when extending to higher dimensions, even if we impose compatibility. It would be interesting to improve generalization to higher dimensions in these examples by imposing additional compatibility conditions. Third, extending a trained network to higher dimensions in our examples involved

solving large and sparse systems, which we currently do not fully exploit. Incorporating sparsity in the toolbox of (Finzi et al., 2021) will enable training larger networks.

7 CONCLUSIONS AND FUTURE WORK

We leveraged representation stability to prove that a broad family of equivariant NNs extends to higher dimensions, which enables us to train an infinite sequence of NNs using a finite amount of data. Extending networks to higher dimensions often results in substantial test errors. To improve generalization, we introduced a compatibility condition relating networks across dimensions. We characterized compatibility and developed an algorithm to train free and compatible NNs. Finally, we applied our method to several numerical examples from the literature. In these examples, free NNs trained without imposing compatibility generalize poorly to higher dimensions, even when learning simple linear functions. In contrast, compatible NNs generalize significantly better.

As highlighted in Section 6, our current implementation faces scalability challenges, primarily due to the inadequate support for GPU-accelerated sparse linear algebra operations in Python. Notably, there are already some efforts to close this gap Potapczynski et al. (2023). Moving forward, we are optimistic about enhancing our framework to leverage these advancements effectively. As we also mentioned in Section 6, the compatibility condition constraints the family of functions one could hope to learn. This presents a compelling avenue for future research: exploring other regularization strategies that lead to good generalization across dimensions. Moreover, the concepts introduced in this paper paved the way to studying any-dimensional statistical learning, raising several natural questions: Can we bound the generalization error across dimensions? How do we characterize the notation of statistical complexity for classes of any-dimensional problems? And how do we approach the formulation of lower bounds within this context?

Acknowledgments

We thank Venkat Chandrasekaran for his constant support and encouragement during the development of this work. Both authors were funded by Venkat Chandrasekaran through AFOSR grant FA9550-20-1-0320. MD was partially funded by Joel Tropp through ONR BRC Award N00014-18-1-2363, and NSF FRG Award 1952777.

References

- Yulia Alexandr, Joe Kileel, and Bernd Sturmfels. Moment varieties for mixtures of products. *arXiv preprint arXiv:2301.09068*, 2023.
- Alexander Bogatskiy, Brandon Anderson, Jan Olfemann, Marwah Roussi, David Miller, and Risi Kondor. Lorentz group equivariant neural network for particle physics. In Hal Daumé III and Aarti Singh, editors, *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 992–1002. PMLR, 13–18 Jul 2020. URL <https://proceedings.mlr.press/v119/bogatskiy20a.html>.
- Michael M Bronstein, Joan Bruna, Taco Cohen, and Petar Veličković. Geometric deep learning: Grids, groups, graphs, geodesics, and gauges. *arXiv preprint arXiv:2104.13478*, 2021.
- Venkat Chandrasekaran, Benjamin Recht, Pablo A Parrilo, and Alan S Willsky. The convex geometry of linear inverse problems. *Foundations of Computational mathematics*, 12(6):805–849, 2012.
- Christopher H Chiu, Alessandro Danelon, Jan Draisma, Rob H Eggermont, and Azhar Farooq. Symnoetherianity for powers of gl-varieties. *arXiv preprint arXiv:2212.05790*, 2022.
- Thomas Church and Benson Farb. Representation theory and homological stability. *Advances in Mathematics*, 245:250–314, 2013. ISSN 0001-8708. doi: <https://doi.org/10.1016/j.aim.2013.06.016>. URL <https://www.sciencedirect.com/science/article/pii/S0001870813002259>.
- Thomas Church, Jordan S. Ellenberg, and Benson Farb. FI-modules and stability for representations of symmetric groups. *Duke Mathematical Journal*, 164(9):1833 – 1910, 2015. doi: 10.1215/00127094-3120274. URL <https://doi.org/10.1215/00127094-3120274>.
- Taco Cohen and Max Welling. Group equivariant convolutional networks. In Maria Florina Balcan and Kilian Q. Weinberger, editors, *Proceedings of The 33rd International Conference on Machine Learning*, volume 48 of *Proceedings of Machine Learning Research*, pages 2990–2999, New York, New York, USA, 20–22 Jun 2016. PMLR. URL <https://proceedings.mlr.press/v48/cohenc16.html>.
- Taco S. Cohen and Max Welling. Steerable CNNs. In *International Conference on Learning Representations*, 2017. URL <https://openreview.net/forum?id=rJQKYt511>.
- Jan Draisma. *Noetherianity up to Symmetry*, pages 33–61. Springer International Publishing, Cham, 2014. ISBN 978-3-319-04870-3. doi: 10.1007/978-3-319-04870-3_2. URL https://doi.org/10.1007/978-3-319-04870-3_2.
- Benson Farb. Representation stability. *arXiv preprint arXiv:1404.4065*, 2014.
- Marc Finzi, Max Welling, and Andrew Gordon Gordon Wilson. A practical method for constructing equivariant multilayer perceptrons for arbitrary matrix groups. In Marina Meila and Tong Zhang, editors, *Proceedings of the 38th International Conference on Machine Learning Research*, volume 139 of *Proceedings of Machine Learning Research*, pages 3318–3328. PMLR, 18–24 Jul 2021. URL <https://proceedings.mlr.press/v139/finzi21a.html>.
- Nir Gadish. Categories of FI type: A unified approach to generalizing representation stability and character polynomials. *Journal of Algebra*, 480:450–486, 2017. ISSN 0021-8693. doi: <https://doi.org/10.1016/j.jalgebra.2017.03.010>. URL <https://www.sciencedirect.com/science/article/pii/S0021869317301849>.
- Justin Gilmer, Samuel S. Schoenholz, Patrick F. Riley, Oriol Vinyals, and George E. Dahl. Neural message passing for quantum chemistry. In Doina Precup and Yee Whye Teh, editors, *Proceedings of the 34th International Conference on Machine Learning Research*, volume 70 of *Proceedings of Machine Learning Research*, pages 1263–1272. PMLR, 06–11 Aug 2017. URL <https://proceedings.mlr.press/v70/gilmer17a.html>.
- Craig Gotsman and Sivan Toledo. On the computation of null spaces of sparse rectangular matrices. *SIAM Journal on Matrix Analysis and Applications*, 30(2):445–463, 2008. doi: 10.1137/050638369. URL <https://doi.org/10.1137/050638369>.
- Mahdi Hashemi. Enlarging smaller images before inputting into convolutional neural network: zero-padding vs. interpolation. *Journal of Big Data*, 6(1):1–13, 2019.
- John Jumper, Richard Evans, Alexander Pritzel, Tim Green, Michael Figurnov, Olaf Ronneberger, Kathryn Tunyasuvunakool, Russ Bates, Augustin Žídek, Anna Potapenko, et al. Highly accurate protein structure prediction with alphafold. *Nature*, 596(7873):583–589, 2021.
- George Em Karniadakis, Ioannis G Kevrekidis, Lu Lu, Paris Perdikaris, Sifan Wang, and Liu Yang. Physics-informed machine learning. *Nature Reviews Physics*, 3(6):422–440, 2021.
- Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In Yoshua Bengio and Yann LeCun, editors, *3rd International Conference on Learning Representations, ICLR 2015, San Diego*,

- CA, USA, May 7-9, 2015, *Conference Track Proceedings*, 2015. URL <http://arxiv.org/abs/1412.6980>.
- Risi Kondor and Shubhendu Trivedi. On the generalization of equivariance and convolution in neural networks to the action of compact groups. In Jennifer Dy and Andreas Krause, editors, *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 2747–2755. PMLR, 10–15 Jul 2018. URL <https://proceedings.mlr.press/v80/kondor18a.html>.
- Risi Kondor, Zhen Lin, and Shubhendu Trivedi. Clebsch–gordan nets: a fully fourier space spherical convolutional neural network. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc., 2018. URL https://proceedings.neurips.cc/paper_files/paper/2018/file/a3fc981af450752046be179185ebc8b5-Paper.pdf.
- Pawel Kowal. Null space of a sparse matrix. <https://www.mathworks.com/matlabcentral/fileexchange/11120-null-space-of-a-sparse-matrix>, 2006. [Retrieved July 12, 2022].
- Alex Krizhevsky, Ilya Sutskever, and Geoffrey E. Hinton. Imagenet classification with deep convolutional neural networks. *Commun. ACM*, 60(6):84–90, may 2017. ISSN 0001-0782. doi: 10.1145/3065386. URL <https://doi.org/10.1145/3065386>.
- Eitan Levin and Venkat Chandrasekaran. Free descriptions of convex sets. *arXiv preprint arXiv:2307.04230*, 2023.
- Lek-Heng Lim and Bradley J. Nelson. What is an equivariant neural network? *Notices Amer. Math. Soc.*, 70(4):619–625, 2023. ISSN 0002-9920. doi: 10.1090/noti2666. URL <https://doi.org/10.1090/noti2666>.
- László Lovász. *Large networks and graph limits*, volume 60. American Mathematical Soc., 2012.
- Haggai Maron, Heli Ben-Hamu, Nadav Shamir, and Yaron Lipman. Invariant and equivariant graph networks. In *International Conference on Learning Representations*, 2019. URL <https://openreview.net/forum?id=Syx72jC9tm>.
- Tomáš Mikolov, Stefan Kombrink, Lukáš Burget, Jan Černocký, and Sanjeev Khudanpur. Extensions of recurrent neural network language model. In *2011 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 5528–5531, 2011. doi: 10.1109/ICASSP.2011.5947611.
- Christopher C Paige and Michael A Saunders. LSQR: An algorithm for sparse linear equations and sparse least squares. *ACM Transactions on Mathematical Software (TOMS)*, 8(1):43–71, 1982.
- Andres Potapczynski, Marc Finzi, Geoff Pleiss, and Andrew Gordon Wilson. Cola: Exploiting compositional structure for automatic and efficient numerical linear algebra. *arXiv preprint arXiv:2309.03060*, 2023.
- Steven Sam and Andrew Snowden. GL-equivariant modules over polynomial rings in infinitely many variables. *Transactions of the American Mathematical Society*, 368(2):1097–1158, 2016.
- Steven Sam and Andrew Snowden. Gröbner methods for representations of combinatorial categories. *Journal of the American Mathematical Society*, 30(1):159–203, 2017.
- Steven V Sam. Structures in representation stability. *Notices of the American Mathematical Society*, 67(1), 2020.
- Steven V Sam and Andrew Snowden. Stability patterns in representation theory. *Forum of Mathematics, Sigma*, 3:e11, 2015. doi: 10.1017/fms.2015.10.
- Richard Socher, Cliff C Lin, Chris Manning, and Andrew Y Ng. Parsing natural scenes and natural language with recursive neural networks. In *Proceedings of the 28th international conference on machine learning (ICML-11)*, pages 129–136, 2011.
- Raphael J. L. Townshend, Stephan Eismann, Andrew M. Watkins, Ramya Rangan, Maria Karelina, Rhiju Das, and Ron O. Dror. Geometric deep learning of rna structure. *Science*, 373(6558):1047–1051, 2021. doi: 10.1126/science.abe5650. URL <https://www.science.org/doi/abs/10.1126/science.abe5650>.
- Dinh Van Le and Tim Römer. Theorems of Carathéodory, Minkowski-Weyl, and Gordan up to symmetry. *arXiv preprint arXiv:2110.10657*, 2021.
- Clément Vignac, Andreas Loukas, and Pascal Frossard. Building powerful and equivariant graph neural networks with structural message-passing. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 14143–14155. Curran Associates, Inc., 2020. URL https://proceedings.neurips.cc/paper_files/paper/2020/file/a32d7eeaae19821fd9ce317f3ce952a7-Paper.pdf.
- Soledad Villar, David W Hogg, Kate Storey-Fisher, Weichi Yao, and Ben Blum-Smith. Scalars are universal: Equivariant machine learning, structured like classical physics. In M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan,

editors, *Advances in Neural Information Processing Systems*, volume 34, pages 28848–28863. Curran Associates, Inc., 2021. URL https://proceedings.neurips.cc/paper_files/paper/2021/file/f1b0775946bc0329b35b823b86eeb5f5-Paper.pdf.

Soledad Villar, Weichi Yao, David W. Hogg, Ben Blum-Smith, and Bianca Dumitrascu. Dimensionless machine learning: Imposing exact units equivariance. *Journal of Machine Learning Research*, 24(109):1–32, 2023. URL <http://jmlr.org/papers/v24/22-0680.html>.

Maurice Weiler, Mario Geiger, Max Welling, Wouter Boomsma, and Taco S Cohen. 3d steerable cnns: Learning rotationally equivariant features in volumetric data. *Advances in Neural Information Processing Systems*, 31, 2018.

Jennifer C.H. Wilson. FI_W-modules and stability criteria for representations of classical Weyl groups. *Journal of Algebra*, 420:269–332, 2014. ISSN 0021-8693. doi: <https://doi.org/10.1016/j.jalgebra.2014.08.010>. URL <https://www.sciencedirect.com/science/article/pii/S0021869314004505>.

Jenny Wilson. An introduction to FI-modules and their generalizations. *Michigan Representation Stability Week*, 2018.

Zonghan Wu, Shirui Pan, Fengwen Chen, Guodong Long, Chengqi Zhang, and Philip S. Yu. A comprehensive survey on graph neural networks. *IEEE Transactions on Neural Networks and Learning Systems*, 32(1):4–24, 2021. doi: 10.1109/TNNLS.2020.2978386.

Manzil Zaheer, Satwik Kottur, Siamak Ravanbakhsh, Barnabas Poczos, Russ R Salakhutdinov, and Alexander J Smola. Deep sets. In I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017. URL https://proceedings.neurips.cc/paper_files/paper/2017/file/f22e4747da1aa27e363d86d40ff442fe-Paper.pdf.

Checklist

1. For all models and algorithms presented, check if you include:
 - (a) A clear description of the mathematical setting, assumptions, algorithm, and/or model. [Yes]
 - (b) An analysis of the properties and complexity (time, space, sample size) of any algorithm. [N/A]
 - (c) (Optional) Anonymized source code, with specification of all dependencies, including external libraries. [No]
2. For any theoretical claim, check if you include:
 - (a) Statements of the full set of assumptions of all theoretical results. [Yes]
 - (b) Complete proofs of all theoretical results. [Yes]
 - (c) Clear explanations of any assumptions. [Yes]
3. For all figures and tables that present empirical results, check if you include:
 - (a) The code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL). [Yes—see Section 5.]
 - (b) All the training details (e.g., data splits, hyperparameters, how they were chosen). [Yes—see Appendix E.]
 - (c) A clear definition of the specific measure or statistics and error bars (e.g., with respect to the random seed after running experiments multiple times). [Yes]
 - (d) A description of the computing infrastructure used. (e.g., type of GPUs, internal cluster, or cloud provider). [Yes]
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets, check if you include:
 - (a) Citations of the creator If your work uses existing assets. [Yes]
 - (b) The license information of the assets, if applicable. [N/A]
 - (c) New assets either in the supplemental material or as a URL, if applicable. [Yes]
 - (d) Information about consent from data providers/curators. [N/A]
 - (e) Discussion of sensible content if applicable, e.g., personally identifiable information or offensive content. [N/A]
5. If you used crowdsourcing or conducted research with human subjects, check if you include:
 - (a) The full text of instructions given to participants and screenshots. [N/A]
 - (b) Descriptions of potential participant risks, with links to Institutional Review Board (IRB) approvals if applicable. [N/A]
 - (c) The estimated hourly wage paid to participants and the total amount spent on participant compensation. [N/A]

A MORE EXAMPLES OF CONSISTENT SEQUENCES

We let $\{\varphi: \mathbb{R}^n \rightarrow \mathbb{R}^{n+1}\}$ denote the sequence of zero-padding embeddings.

Signed permutation sequences. Let B_n (resp. D_n) be the signed permutation group (resp. even-signed permutation group) in n elements. Consider the sequence of representations $\mathbf{V} = \mathbb{R}^n$ with the standard action of permuting and flipping the signs of the coordinates. Then, $\mathcal{V} = \{(\mathbf{V}, \varphi)\}$ is a consistent sequence with generation and presentation degrees equal to one.

Orthogonal rotation sequences. Let $O(n)$ (resp., $SO(n)$) be the orthogonal group (resp., special orthogonal group) in dimension n . Consider the sequence of representations $\mathbf{V} = \mathbb{R}^n$ with the standard rotation action. Then, $\mathcal{V} = \{(\mathbf{V}, \varphi)\}$ is again a consistent sequence with generation degree equal to one. Characterizing the presentation degree is an open question. Yet, numerically, the presentation degree for S_n seems to work well.

Lorentz sequences. Let $O(1, n)$ be the Lorentz group (resp., special orthogonal group) in dimension $n + 1$. Again, consider the sequence of representations $\mathbf{V} = \mathbb{R}^{n+1}$ with the natural action of the Lorentz group. Then, $\mathcal{V} = \{(\mathbf{V}, \varphi)\}$ is once more a consistent sequence with generation degree equal to one.

Tensor products. For any of the above sequences, we can take tensor powers $\mathcal{V}^{\otimes k} = \{(\mathbf{V}^{\otimes k}, \varphi^{\otimes k})\}$ to obtain sequences of tensors. Here $\varphi^{\otimes k}$ zero-pads a $n \times \dots \times n$ tensor $T \in (\mathbb{R}^n)^{\otimes k}$ to a $(n+1) \times \dots \times (n+1)$ tensor, and the group acts on rank-1 tensors by $g \cdot (x_1 \otimes \dots \otimes x_k) = (gx_1) \otimes \dots \otimes (gx_k)$ which extends to an action on general tensors by linearity. For a example, a permutation $g \in S_n$ acts on $T \in (\mathbb{R}^n)^{\otimes k}$ by $(g \cdot T)_{i_1, \dots, i_k} = T_{g(i_1), \dots, g(i_k)}$.

Direct sums. For any of the above sequences, we can take direct sums $\mathcal{V}^{\oplus k} = \{(\mathbf{V}^{\oplus k}, \varphi^{\oplus k})\}$. Here $\varphi^{\oplus k}$ zero-pads each of the k vectors in an element of $(\mathbb{R}^n)^{\oplus k}$, and the group acts on each vector in such an element by $g \cdot (x_1, \dots, x_k) = (g \cdot x_1, \dots, g \cdot x_k)$.

B PRESENTATION DEGREE

In this section, we give a definition of the presentation degree. In what follows, we introduce a series of preliminary definitions. We highlight that the concepts in this section are technical in nature and require a certain familiarity with abstract algebra. We refer the interested reader to (Levin and Chandrasekaran, 2023, §4.1) for a more accessible presentation of this material motivated by the problem of extending an equivariant map to a compatible sequence. For this section, we drop the boldface notation since we need to use additional superscripts.

Definition B.1 (Stabilizing subgroups). If $\mathcal{V} = \{V^{(n)}\}$ is a consistent sequence of $\{G^{(n)}\}$ -representations, define its *stabilizing subgroups* $\{H^{(n,d)}\}_{d \leq n}$ by

$$H^{(n,d)} = \{g \in G^{(n)} : g \cdot v = v \text{ for all } v \in V^{(d)}\}.$$

In the context of $G^{(n)} = S_n$ and $V^{(n)} = \mathbb{R}^n$, the subgroup $H^{(n,d)}$ corresponds to the permutations in n elements that leave fixed the first d components.

Definition B.2 (Modules). Let \mathcal{V} and $\mathcal{U} = \{U^{(n)}\}$ be consistent sequences of $\{G^{(n)}\}$ -representations and let $\{H^{(n,d)}\}$ be the stabilizing subgroups of \mathcal{V} . Then \mathcal{U} is a \mathcal{V} -module if $U^{(d)} \subseteq (U^{(n)})^{H^{(n,d)}}$ for all $d \leq n$.

Recall that if $H \subseteq G$ is a subgroup and U is an H -representation, the induced representation of U to G is $\text{Ind}_H^G U = \mathbb{R}[G] \otimes_{\mathbb{R}[H]} U$. If $G \subseteq G'$ and $H \subseteq H'$ are subgroups satisfying $H' \cap G = H$ and U is an H' -representation, we have an inclusion $\text{Ind}_H^G U \subseteq \text{Ind}_{H'}^{G'} U$.

Definition B.3 (Induction and algebraically free sequences). Let \mathcal{V} by a consistent sequence of $\{G^{(n)}\}$ -representations, and for $d \leq n$ let $H^{(n,d)} \subseteq G^{(n)}$ be its stabilizing subgroups.

(a) Fix $d \in \mathbb{N}$, and a $G^{(d)}$ -representation $U^{(d)}$. Define the \mathcal{V} -induction sequence

$$\text{Ind}_{G^{(d)}}(U^{(d)}) = \left\{ \text{Ind}_{G^{(d)}H^{(n,d)}}^{G^{(n)}} U^{(d)} \right\}_n,$$

where the induced representation is taken to be 0 when $n < d$, and $G^{(d)}H^{(n,d)} = \{gh : g \in G^{(d)}, h \in H^{(n,d)}\}$ is the subgroup generated by $G^{(d)}$ and $H^{(n,d)}$ inside $G^{(n)}$. The \mathcal{V} -induction sequence is a \mathcal{V} -module by construction.

- (b) A consistent sequence \mathcal{F} is an *algebraically free \mathcal{V} -module* if it is a direct sum of \mathcal{V} -induction sequences. The sequence \mathcal{V} itself is algebraically *free* if it is an algebraically free \mathcal{V} -module.

Since $ghg^{-1} \in H^{(n,d)}$ for any $g \in G^{(d)}$ and $h \in H^{(n,d)}$, the subgroup generated by $G^{(d)}$ and $H^{(n,d)}$ inside $G^{(n)}$ is the set of products $G^{(d)}H^{(n,d)}$. As we have inclusions $G^{(d)}H^{(n,d)} \subseteq G^{(d)}H^{(n+1,d)}$ and $G^{(n)} \subseteq G^{(n+1)}$, and moreover $G^{(d)}H^{(n+1,d)} \cap G^{(n)} = G^{(d)}H^{(n,d)}$, we have inclusions $\text{Ind}_{G^{(d)}H^{(n,d)}}^{G^{(n)}} U^{(d)} \subseteq \text{Ind}_{G^{(d)}H^{(n+1,d)}}^{G^{(n+1)}} U^{(d)}$. This shows that $\text{Ind}_{G^{(d)}}(U^{(d)})$ is indeed a consistent sequence.

Example B.4. Let $\mathcal{V} = \{\mathbb{R}^n\}$ with embeddings by zero-padding and the action of $G^{(n)} = S_n$ from Section 2. Then $\mathcal{V} = \text{Ind}_{G^{(1)}}(V^{(1)})$, hence it is algebraically free.

Any consistent sequence is a quotient of an algebraically free one. To formalize this, we first define an appropriate notion of maps between consistent sequences.

Definition B.5 (Morphisms of sequences). If $\mathcal{V} = \{V^{(n)}\}$ and $\mathcal{U} = \{U^{(n)}\}$ are consistent sequences of $\{G^{(n)}\}$ -representations, then a *morphism of sequences* $\mathcal{V} \rightarrow \mathcal{U}$ is a sequence of equivariant linear maps $\{W^{(n)} \in \mathcal{L}(V^{(n)}, U^{(n)})^{G^{(n)}}\}$ satisfying $W^{(n+1)}|_{V^{(n)}} = W^{(n)}$.

Note that (CompNN) precisely requires the sequences of weights in a compatible NN to define a morphism of sequences. To write a general consistent sequence $\mathcal{V} = \{V^{(n)}\}$ as a quotient of a free one, suppose \mathcal{V} is generated in degree d . Define the algebraically free \mathcal{V} -module $\mathcal{F} = \bigoplus_{i=1}^d \text{Ind}_{G^{(i)}}(V^{(i)}) = \{F^{(n)}\}$ and consider the sequence of linear maps $W^{(n)}: F^{(n)} \rightarrow V^{(n)}$ sending $g \otimes v \mapsto g \cdot v$ for each $g \otimes v \in \text{Ind}_{G^{(i)}}(V^{(i)})$. It is easy to check that $\{W^{(n)}\}: \mathcal{F} \rightarrow \mathcal{V}$ is a morphism of sequences and that $W^{(n)}$ is surjective for any $n \geq d$, as the image of $W^{(n)}$ is precisely $\mathbb{R}[G^{(n)}]V^{(\min\{n,d\})} = V^{(n)}$. The sequence of kernels of such a morphism of sequence $\{W^{(n)}\}$ is also a consistent sequence.

Proposition B.6. If $\mathcal{V} = \{(V^{(n)}, \varphi)\}$ and $\mathcal{U} = \{U^{(n)}\}$ are consistent sequences and $\{W^{(n)}\}: \mathcal{V} \rightarrow \mathcal{U}$ is a morphism of sequences, then $\ker\{W^{(n)}\} = \{(\ker W^{(n)}, \varphi)\}$ is a consistent sequence.

Proof. Since $W^{(n)}$ is $G^{(n)}$ -equivariant, its kernel is a $G^{(n)}$ -subrepresentation of $V^{(n)}$. The embeddings φ remain $G^{(n)}$ -equivariant isometries when restricted to $\ker W^{(n)}$. \square

We are now ready to define the presentation degree.

Definition B.7 (Presentation degree). Let \mathcal{V}_0 be a consistent sequence of $\{G^{(n)}\}$ -representations. We say that a \mathcal{V}_0 -module \mathcal{V} is generated in degree d , and presented in degree

$$k = \max\{d, r\}$$

if there exists an algebraically free \mathcal{V}_0 -module \mathcal{F} generated in degree d and a morphism of sequences $\{W^{(n)}\}: \mathcal{F} \rightarrow \mathcal{V}$ such that $W^{(n)}$ is surjective for all n and $\ker\{W^{(n)}\}$ is generated in degree r . The *presentation degree* is the smallest such $k \in \mathbb{N}$.

Note that the presentation degree is always at least as large as the generation degree and that the two are equal for algebraically free sequences. As stated in Section 2, the projections in Proposition 2.3 become isomorphisms starting from the presentation degree. The following proof of this fact, usually stated in the representation stability literature in terms of *coinvariants* (Church et al., 2015, §3), appeared first in (Levin and Chandrasekaran, 2023).

Proposition B.8. Let \mathcal{V}_0 be a consistent sequence of $\{G^{(n)}\}$ -representations and \mathcal{V} be a \mathcal{V}_0 -module presented in degree k . Then, the maps $\mathcal{P}_{U^{(n)}}: (V^{(n+1)})^{G^{(n+1)}} \rightarrow (V^{(n)})^{G^{(n)}}$ are isomorphisms for all $n \geq k$.

Proof. As \mathcal{V} is presented in degree k , there exists an algebraically free \mathcal{V}_0 -module $\mathcal{F} = \{F^{(n)}\}$ and a surjective morphism $\{W^{(n)}\}: \mathcal{F} \rightarrow \mathcal{V}$ such that both its kernel $\mathcal{K} = \{K^{(n)}\}$ and \mathcal{F} itself are generated in degree k . Because each map $F^{(n)} \rightarrow V^{(n)}$ is a $G^{(n)}$ -equivariant surjection with kernel $K^{(n)}$, its restriction to invariants $(F^{(n)})^{G^{(n)}} \rightarrow (V^{(n)})^{G^{(n)}}$ is surjective with kernel $(K^{(n)})^{G^{(n)}}$.

As \mathcal{F} is an algebraically free \mathcal{V}_0 -module, there exist integers d_j and $G^{(d_j)}$ -representations $U^{(d_j)}$ satisfying $\mathcal{F} = \bigoplus_j \text{Ind}_{G^{(d_j)}} U^{(d_j)}$. Such \mathcal{F} has generation degree $\max_j d_j \leq k$. Therefore, letting $\{H^{(n,d)}\}$ be the stabilizing

subgroups of \mathcal{V}_0 , we have for $n \geq k$

$$\begin{aligned} (F^{(n)})^{G^{(n)}} &= \bigoplus_j \left(\text{Ind}_{G^{(d_j)} H^{(n,d_j)}}^{G^{(n)}} (U^{(d_j)}) \right)^{G^{(n)}} \\ &\cong \bigoplus_j (U^{(d_j)})^{G^{(d_j)}}. \end{aligned}$$

The last isomorphism follows from the fact that for any groups $H \subseteq G$ and any H -representation U , we have an isomorphism $(\text{Ind}_H^G(U))^G \cong U^H$ given by sending $u \in U^H$ to $\sum_{i=1}^k g_i \otimes u \in (\text{Ind}_H^G(U))^G$ where $g_1 = \text{id}, g_2, \dots, g_k$ are coset representatives for G/H . Thus, $\dim(F^{(n)})^{G^{(n)}}$ is constant for $n \geq k$. Moreover, by Proposition 2.3 and the fact that \mathcal{K} and \mathcal{U} are generated in degree k , we have $\dim(K^{(n)})^{G^{(n)}} \geq \dim(K^{(n+1)})^{G^{(n+1)}}$ and $\dim(V^{(n)})^{G^{(n)}} \geq \dim(V^{(n+1)})^{G^{(n+1)}}$ for all $n \geq k$.

By the rank-nullity theorem, we have $\dim(V^{(n)})^{G^{(n)}} = \dim(F^{(n)})^{G^{(n)}} - \dim(K^{(n)})^{G^{(n)}}$. As $\dim(F^{(n)})^{G^{(n)}}$ is constant while both $\dim(V^{(n)})^{G^{(n)}}$ and $\dim(K^{(n)})^{G^{(n)}}$ are nonincreasing for $n \geq k$, we conclude that they are all constant for $n \geq k$. To conclude, note that $\mathcal{P}_{V^{(n)}}$ is injective for all $n \geq k$ by Proposition 2.3. \square

C PROOF OF THEOREM 3.2

The following proof first appeared in (Levin and Chandrasekaran, 2023). Since \mathcal{V}, \mathcal{U} are obtained from direct sums and tensor products of \mathcal{V}_0 , they are both \mathcal{V}_0 -modules. Let $\{H^{(n,d)}\}$ be the stabilizing subgroups of \mathcal{V}_0 . Suppose first that $\mathcal{V} = \mathcal{F} = \bigoplus_j \text{Ind}_{G^{(d_j)} H^{(n,d_j)}} F^{(d_j)}$ is free. Note that it is generated in degree $\max_j d_j \leq d_g$. Let $W_{d_j} = W^{(n_0)}|_{F^{(d_j)}}$ and fix $n \geq d_j$. Because \mathcal{U} is a \mathcal{V}_0 -module, we have $U^{(d_j)} \subseteq (U^{(n)})^{H^{n,d_j}}$, so we can view $U^{(d_j)}$ as a representation of $G^{(d_j)} H^{(n,d_j)}$ on which $H^{(n,d_j)}$ acts trivially.

In general, if $H \subseteq G$ is a subgroup, if V, U are H -representations, and if $W \in \mathcal{L}(V, U)^H$, we can define $\text{Ind}(W): \text{Ind}_H^G V \rightarrow \text{Ind}_H^G U$ by sending $g \otimes v \mapsto g \otimes Wv$ for $g \in G$ and $v \in V$. Also, if $V \subseteq U$ and V is furthermore a G -representation, then there is an equivariant map $\text{Ind}_H^G V \rightarrow U$ sending $g \otimes v \mapsto g \cdot v$. As $W_{d_j}(F^{(d_j)}) \subseteq U^{(d_j)}$ and is $G^{(d_j)} H^{(n,d_j)}$ -equivariant, we can combine the above two maps to obtain the following equivariant composition

$$\begin{aligned} W_j^{(n)}: \text{Ind}_{G^{(d_j)} H^{(n,d_j)}}^{G^{(n)}} F^{(d_j)} &\xrightarrow{\text{Ind}(W_{d_j})} \text{Ind}_{G^{(d_j)} H^{(n,d_j)}}^{G^{(n)}} U^{(d_j)} \\ &\xrightarrow{g \otimes v \mapsto g \cdot v} U^{(n)}. \end{aligned}$$

Note that $W_j^{(n_0)} = W^{(n_0)}|_{\text{Ind}_{G^{(d_j)} H^{(n_0,d_j)}} F^{(d_j)}}$, since $W_j^{(n_0)}(g \otimes f) = g \cdot W^{(n_0)}f$ for all $g \in G^{(n_0)}$ and $f \in F^{(d_j)}$. Also, $\{W_j^{(n)}\}$ defines a morphism $\text{Ind}_{G^{(d_j)} H^{(n,d_j)}} F^{(d_j)} \rightarrow \mathcal{U}$. Therefore, the desired extension of $W^{(n_0)}$ to a morphism of sequences $\{W^{(n)}\}$ is given by $W^{(n)} = \bigoplus_j W_j^{(n)}: V^{(n)} \rightarrow U^{(n)}$.

Now suppose \mathcal{F} is an algebraically free \mathcal{V} -module as above with a surjection $\mathcal{F} \rightarrow \mathcal{V}$ whose kernel $\mathcal{K} = \{K^{(n)}\}$ is generated in degree d_p . Define the composition

$$\widetilde{W}^{(n_0)}: F^{(n_0)} \rightarrow V^{(n_0)} \xrightarrow{W^{(n_0)}} U^{(n_0)},$$

which satisfies $\widetilde{W}^{(n_0)}(F^{(j)}) \subseteq U^{(j)}$ for all $j \leq d_g$ by assumption and $\widetilde{W}^{(n_0)}(K^{(n_0)}) = 0$ by its definition. By the previous paragraph, it extends to a morphism $\{\widetilde{W}^{(n)}: F^{(n)} \rightarrow U^{(n)}\}$. Because \mathcal{K} is generated in degree d_p and $n_0 \geq d_p$, we have $K^{(n)} = \mathbb{R}[G^{(n)}]K^{(n_0)}$. Because $\widetilde{W}^{(n)}$ is equivariant, we have $\widetilde{W}^{(n)}(K^{(n)}) = 0$. Therefore, $\widetilde{W}^{(n)}$ can be factored as $F^{(n)} \rightarrow F^{(n)}/K^{(n)} = V^{(n)} \xrightarrow{W^{(n)}} U^{(n)}$, where the maps $W^{(n)}$ in this factorization give the desired extension of $W^{(n_0)}$ to a morphism $\mathcal{V} \rightarrow \mathcal{U}$. \square

D CORRECTNESS OF THE COMPUTATIONAL RECIPE

In this section, we prove the correctness of Algorithm 1. That is, if n_0 is set to be large enough, then the algorithm will generate the unique free (or compatible) extension of the network learned at level n_0 . It is useful to use the notation $d_p(\mathcal{V})$ to denote the presentation degree of the sequence \mathcal{V} .

Theorem D.1. *Consider a neural network architecture with consistent sequences $\mathcal{V}_1, \mathcal{U}_1, \dots, \mathcal{V}_L, \mathcal{U}_L$, and \mathcal{V}_{L+1}*

Experiment	Architecture		
	Input	Hidden	Output
Trace	$\mathcal{V}^{\otimes 2}$	$2\mathcal{V} + 2\mathcal{V}^{\otimes 2}$	\mathcal{S}
Diagonal extraction	$\mathcal{V}^{\otimes 2}$	$4\mathcal{V} + 4\mathcal{V}^{\otimes 2}$	$\mathcal{V}^{\otimes 2}$
Symmetric projection	$\mathcal{V}^{\otimes 2}$	$4\mathcal{V} + 4\mathcal{V}^{\otimes 2}$	$\mathcal{V}^{\otimes 2}$
Top singular vector	$\mathcal{V}^{\otimes 2}$	$25\mathcal{S} + 10\mathcal{V} + 2\mathcal{V}^{\otimes 2} + \mathcal{V}^{\otimes 3}$	\mathcal{V}
Orthogonal invariance	$2\mathcal{V}$	$25\mathcal{S} + 10\mathcal{V} + 2\mathcal{V}^{\otimes 2} + \mathcal{V}^{\otimes 3}$	\mathcal{S}

Table 2: Architectures used for the numerical examples. Recall that $\mathcal{S} = \{\mathbb{R}\}$ is the scalar sequence endowed with the trivial action $g \cdot v = v$.

and consistent activation functions $\sigma_1, \dots, \sigma_L$ as described Figure 1. Assume we run Algorithm 1 to learn a neural network at level n_0 . Then, the following two hold.

(Free) If $n_0 \geq \max_i d_p(\mathcal{V}_i \otimes \mathcal{U}_i)$ and $n_0 \geq \max_i d_p(\mathcal{U}_i)$, then, Algorithm 1 extends the trained network at level n_0 to a unique free neural network.

(Compatible) If $n_0 \geq \max_i d_p(\mathcal{V}_i)$ and $n_0 \geq \max_i d_p(\mathcal{U}_i)$, then, Algorithm 1 extends the trained network at level n_0 to a unique compatible neural network.

Proof. We start by proving the second statement. This is a simple corollary of Theorem 3.2. Note that once n_0 is bigger than the presentation degrees of \mathcal{V}_i and \mathcal{U}_i , then, by Theorem 3.2, there is a unique extension of the weights and biases satisfying CompNN, which corresponds to the unique solution of (ExtSyst). The first statement follows by an analogous argument by substituting Theorem 3.2 with Proposition B.8. This completes the proof. \square

E DETAILED DESCRIPTION OF THE NUMERICAL EXPERIMENTS

In this section, we elaborate on the implementation details of the numerical experiments. The code and instructions to reproduce the experiments can be found in <https://github.com/mateodd25/free-nets>. All experiments were run on a 2021 Macbook Air M1 with 16GB of RAM.

Initialization. We initialize the weights and biases at random with small i.i.d Gaussian entries $N(0, 1/100)$.

Data generation. The datasets $\{(X_i, f^{(n)}(X_i))\}$ for both training and testing (in every dimension) are generated as i.i.d. standard Gaussian samples with the appropriate size.

Architecture. For all the experiments, we use three layers, i.e., two hidden layers. All the sequences we take are sums and tensors products of the base sequence $\mathcal{V} = \{\mathbb{R}^d\}$. The activation functions are a composition of the form

$$\mathbf{h}(\mathbf{bilinear}(x) + x)$$

where **bilinear** (Finzi et al., 2021) is described in Section 3 and **h** changes depending on the group; for the permutation group, **h** applies entry-wise ReLU for the permutation group, while for the orthogonal group, it applies a gated nonlinearity (Weiler et al., 2018) —also described in Section 3.

We chose the smallest hidden layer size that yielded competitive results in order to control the size of the linear systems we employ for extending the networks. The top singular vector and orthogonal invariant examples required a larger architecture, so we only extended them to smaller dimensions. The architecture for each example are summarized in Table 2.