Provably Strict Generalisation Benefit for Equivariant Models

Bryn Elesedy [1] **Sheheryar Zaidi** [2]

Abstract

It is widely believed that engineering a model to be invariant/equivariant improves generalisation. Despite the growing popularity of this approach, a precise characterisation of the generalisation benefit is lacking. By considering the simplest case of linear models, this paper provides the first *provably non-zero* improvement in generalisation for invariant/equivariant models when the target distribution is invariant/equivariant with respect to a compact group. Moreover, our work reveals an interesting relationship between generalisation, the number of training examples and properties of the group action. Our results rest on an observation of the structure of function spaces under averaging operators which, along with its consequences for feature averaging, may be of independent interest.

1. Introduction

There is significant and growing interest in models, especially neural networks, that are invariant or equivariant to the action of a group on their inputs. It is widely believed that these models enjoy improved generalisation when the group is correctly specified. The intuition being that if the salient aspects of a task are unchanged by some transformation, then more flexible models would need to learn (as opposed to being hard-coded) to ignore these transformations, requiring more samples to generalise. Work in this area has progressed quickly (Cohen & Welling, 2016; Cohen et al., 2018; 2019) and has found application in domains where the symmetry is known a priori, for instance in particle physics (Pfau et al., 2020).

In contrast with practical successes, our theoretical understanding of invariant/equivariant models is limited. Many previous works that have attempted to address the generalisation of invariant/equivariant models, such

as Sokolic et al. (2017); Sannai & Imaizumi (2019), cover only the worst case performance of algorithms. These works use complexity measures to find tighter upper bounds on the test risk for invariant/equivariant models, but a strict benefit is not demonstrated. The VC dimension governs distribution independent generalisation, so these complexity measures can show no more than a separation between the VC dimensions of a model and its invariant/equivariant version. This would imply the existence of a distribution on which the model with smaller VC dimension will generalise better, but would not rule out that on many common distributions training an invariant/equivariant model using standard procedures provides no benefit. For instance, there could be many invariant/equivariant distributions on which SGD automatically favours parameters that result in (possibly approximately) invariant/equivariant predictors, regardless of architecture.

The results of this paper move to address this issue, by quantifying exactly the generalisation benefit of invariant/equivariant linear models. We do this in the case of the minimum L_2-norm least-squares solution, which reflects the implicit bias of gradient descent in overparameterised linear models. While the linear model provides a tractable first-step towards understanding more complex models such as neural networks, the underlying ideas of this paper are equally applicable to non-linear predictors. We emphasise this by providing new perspectives on feature averaging and suggestions for how to apply the ideas of this paper to find new methods for training invariant/equivariant neural networks.

1.1. Our Contributions

The main result of this paper is Theorem 13, which quantifies the generalisation benefit of equivariance in a linear model. We define the *generalisation gap* $\Delta(f, f')$ between predictors f and f' to be the difference in their test errors on a given task. A positive generalisation gap $\Delta(f, f') > 0$ means that f' has strictly smaller test error than f. Theorem 13 concerns $\Delta(f, f')$ in the case that $f : \mathbb{R}^d \to \mathbb{R}^k$ is the minimum-norm least-squares predictor and f' is its equivariant version. Let a compact group $\mathcal{G}$ act via orthogonal representations ϕ and ψ on inputs $X \sim \mathcal{N}(0, I_d)$ and outputs $Y = h(X) + \xi \in \mathbb{R}^k$ respectively, where $h : \mathbb{R}^d \to \mathbb{R}^k$ is an equivariant linear

---

[1]Department of Computer Science, University of Oxford, Oxford, United Kingdom [2]Department of Statistics, University of Oxford, Oxford, United Kingdom. Correspondence to: Bryn Elesedy <bryn@robots.ox.ac.uk>.

map. Let $(\chi_\psi | \chi_\phi) = \int_{\mathcal{G}} \mathrm{Tr}(\psi(g)) \, \mathrm{Tr}(\phi(g)) \, \mathrm{d}\lambda(g)$ denote the scalar product of the characters of the representations. The generalisation benefit of enforcing equivariance in a linear model is given by

$$\mathbb{E}[\Delta(f, f')] = \mathrm{Var}\,[\xi]\, r(n, d)(dk - (\chi_\psi | \chi_\phi)) + \mathcal{E}_{\mathcal{G}}(n, d)$$

where

$$r(n, d) = \begin{cases} \frac{n}{d(d-n-1)} & n < d - 1 \\ (n - d - 1)^{-1} & n > d + 1 \\ \infty & \text{otherwise} \end{cases}$$

and $\mathcal{E}_{\mathcal{G}}(n, d) \geq 0$ is the generalisation gap of the corresponding noiseless problem, that vanishes when $n \geq d$. The divergence at the interpolation threshold $n \in [d-1, d+1]$ is consistent with the double descent literature (Hastie et al., 2019).

The quantity $dk - (\chi_\psi | \chi_\phi)$ represents the significance of the group symmetry to the task. The dimension of the space of linear maps $\mathbb{R}^d \to \mathbb{R}^k$ is dk, while $(\chi_\psi | \chi_\phi)$ is the dimension of the space of equivariant linear maps. We will see later that the quantity $dk - (\chi_\psi | \chi_\phi)$ represents the dimension of the space of linear maps that vanish when averaged over $\mathcal{G}$; it is through the dimension of this space that the symmetry in the task controls the generalisation gap. Although invariance is a special case of equivariance, we find it instructive to discuss it separately. In Theorem 7 we provide a result that is analogous to Theorem 13 for invariant predictors, along with a separate proof.

In order to arrive at Theorems 7 and 13 we make use of general results about the structure of function spaces under averaging operators. In Section 4 we show how averaging operators can be used to decompose function spaces into orthogonal subspaces of symmetric (invariant/equivariant) and anti-symmetric (vanish when averaged) functions. In Section 5 we use these insights to provide new perspectives on feature averaging. Our main results are in Section 6. Finally, in Section 7 we apply our insights to derive principled methods for training invariant/equivariant neural networks and provide open questions for future work.

2. Related Work

Implementations   While there has been a recent surge in interest, symmetry is not a new concept in machine learning. Recent literature is dominated by neural networks, but other methods do exist: e.g. kernels (Haasdonk et al., 2005), support vector machines (Schölkopf et al., 1996) or feature spaces such as polynomials (Schulz-Mirbach, 1994; 1992). The engineering of invariant neural networks dates back at least to Wood & Shawe-Taylor (1996), in which ideas from representation theory are applied to find weight tying schemes that result in group invariant architectures;

similar themes are present in Ravanbakhsh et al. (2017). Recent work follows in this vein, borrowing ideas from fundamental physics to construct invariant/equivariant convolutional architectures (Cohen & Welling, 2016; Cohen et al., 2018). Correspondingly, a sophisticated theory of invariant/equivariant networks has arisen (Kondor & Trivedi, 2018; Cohen et al., 2019) including universal approximation results (Maron et al., 2019; Yarotsky, 2018).

Learning and Generalisation   The intuition that invariant or equivariant models are more sample efficient or generalisable is widespread in the literature, but arguments are often heuristic and, to the best of our knowledge, a provably strict (non-zero) generalisation benefit has not appeared before this paper. It was noted (Abu-Mostafa, 1993) that constraining a model to be invariant cannot increase its VC dimension. An intuitive argument for reduced sample complexity is made in Mroueh et al. (2015) in the case that the input space has finite cardinality. The sample complexity of linear classifiers with invariant representations trained on a simplified image task is discussed briefly in Anselmi et al. (2014), the authors conjecture that a general result may be obtained using wavelet transforms. The framework of robustness (Xu & Mannor, 2012) is used in Sokolic et al. (2017) to obtain a generalisation bound for interpolating large-margin classifiers that are invariant to a finite set of transformations; note that the results contain an implicit margin constraint on the training data. The generalisation of models invariant or equivariant to finite permutation groups is considered in Sannai & Imaizumi (2019). Both of Lyle et al. (2019; 2020) cover the PAC Bayes approach to generalisation of invariant models, the latter also considers the relative benefits of feature averaging and data augmentation.

3. Preliminaries

We assume familiarity with the basic notions of group theory, as well as the definition of a group action and a linear representation. The reader may consult Wadsley (2012); Serre (1977, Chapters 1-4) for background. We define some key concepts and notation here and introduce more as necessary throughout the paper.

Notation and Technicalities   We write $\mathcal{X}$ and $\mathcal{Y}$ for input and output spaces respectively. We assume for simplicity that $\mathcal{Y} = (\mathbb{R}^k, +)$ is a k-dimensional real vector space (with k finite) but we expect our results to apply in other settings too. Throughout the paper, $\mathcal{G}$ will represent an arbitrary compact group that has a measurable action ϕ on $\mathcal{X}$ and a representation ψ on $\mathcal{Y}$. By this we mean that $\phi : \mathcal{G} \times \mathcal{X} \to \mathcal{X}$ is a measurable map and the same for ψ. We sometimes write gx as a shorthand for $\phi(g)x$ and similarly for actions on $\mathcal{Y}$. Some notation for

specific groups: C_m and S_m are, respectively, the cyclic and symmetric groups on m elements; while $O(m)$ and $SO(m)$ are the m-dimensional orthogonal and special orthogonal groups respectively. For any matrix A we write A^+ for the Moore-Penrose pseudo-inverse and $\|A\|_F = \sqrt{\text{Tr}(A^\top A)}$ for the Frobenius/Hilbert-Schmidt norm. We write $\mathbb{G}_n(\mathbb{R}^d)$ for the Grassmannian manifold of subspaces of dimension n in $\mathbb{R}^d$. The results of this paper require some mild care with technical considerations such as topology/measurability. We do not stress these in the main paper but they do appear in the proofs, all of which are deferred to the supplementary material.

Invariance, Equivariance and Symmetry A function $f : \mathcal{X} \to \mathcal{Y}$ is $\mathcal{G}$-invariant if $f(\phi(g)x) = f(x) \; \forall x \in \mathcal{X}$ $\forall g \in \mathcal{G}$ and is $\mathcal{G}$-equivariant if $f(\phi(g)x) = \psi(g)f(x)$ $\forall x \in \mathcal{X} \; \forall g \in \mathcal{G}$. A measure μ on $\mathcal{X}$ is $\mathcal{G}$-invariant if $\forall g \in \mathcal{G}$ and any μ-measurable $B \subset \mathcal{X}$ the pushforward of μ by the action ϕ equals μ, i.e. $(\phi(g)_*\mu)(B) = \mu(B)$. This means that if $X \sim \mu$ then $\phi(g)X \sim \mu \; \forall g \in \mathcal{G}$. We will sometimes use the catch-all term symmetric to describe an object that is invariant or equivariant.

4. Symmetric and Anti-Symmetric Functions

Averaging the inputs of a function over the orbit of a group is a well known method to enforce invariance, for instance see Schulz-Mirbach (1994). Approaching this from another perspective, averaging can also be used to identify invariance. That is, a function is $\mathcal{G}$-invariant if and only if it is preserved by orbit averaging with respect to $\mathcal{G}$. The same can be said for equivariant functions, using a modified average. After introducing the relevant concepts, we will use this observation and other properties of the averaging operators to decompose function spaces into mutually orthogonal symmetric (invariant/equivariant) and anti-symmetric (vanish when averaged) subspaces. This observation provides the foundation for many results later in the paper.

4.1. Setup

Haar Measure Let $\mathcal{G}$ be a compact group. The Haar measure is the unique invariant measure on $\mathcal{G}$ and we denote it by λ. By invariance we mean that for any measurable subset $A \subset \mathcal{G}$ and for any $g \in \mathcal{G}$, $\lambda(gA) = \lambda(Ag) = \lambda(A)$. We assume normalisation such that $\lambda(\mathcal{G}) = 1$, which is always possible when $\mathcal{G}$ is compact. The (normalised) Haar measure can be interpreted as the uniform distribution on $\mathcal{G}$. See Kallenberg (2006) for more details.

Orbit Averaging For any feature map $f : \mathcal{X} \to \mathcal{Y}$, we can construct a $\mathcal{G}$-invariant feature map by averaging with

respect to λ. We represent this by the operator $\mathcal{O}$, where

$$(\mathcal{O}f)(x) = \int_{\mathcal{G}} f(gx) \, \mathrm{d}\lambda(g).$$

Similarly, if ψ is a representation of $\mathcal{G}$ on $\mathcal{Y}$, we can transform f into an equivariant feature map by applying $\mathcal{Q}$, where

$$(\mathcal{Q}f)(x) = \int_{\mathcal{G}} \psi(g^{-1})f(gx) \, \mathrm{d}\lambda(g).$$

Notice that $\mathcal{O}$ is a special case of $\mathcal{Q}$ corresponding to ψ being the trivial representation. The operator $\mathcal{O}$ can be thought of as performing feature averaging with respect to $\mathcal{G}$. This interpretation is widely adopted, for instance appearing in Lyle et al. (2020).

Function Spaces We now show how to construct the relevant spaces of functions. We present this in an abstract way, but these functions can be interpreted as predictors, feature maps, feature extractors and so on. Let μ be a $\mathcal{G}$-invariant measure on $\mathcal{X}$ and let $\langle a, b \rangle$ for $a, b \in \mathcal{Y}$ be an inner product on $\mathcal{Y} = \mathbb{R}^k$ that is preserved by ψ.[1] By preserved we mean that $\langle \psi(g)a, \psi(g)b \rangle = \langle a, b \rangle$, $\forall g \in \mathcal{G}$, $\forall a, b \in \mathcal{Y}$ and any inner product can be transformed to satisfy this property using the Weyl trick $\langle a, b \rangle \mapsto \int_{\mathcal{G}} \langle \psi(g)a, \psi(g)b \rangle \, \mathrm{d}\lambda(g)$. Given two functions $f, h : \mathcal{X} \to \mathcal{Y}$, we define their inner product by

$$\langle f, h \rangle_\mu = \int_{\mathcal{X}} \langle f(x), h(x) \rangle \, \mathrm{d}\mu(x).$$

This inner product can be thought of as comparing the similarity between functions and can used to define a notion of distance with the norm $\|f\|_\mu = \sqrt{\langle f, f \rangle_\mu}$. We then define V as the space of all (measurable) functions $f : \mathcal{X} \to \mathcal{Y}$ such that $\|f\|_\mu < \infty$.[2] Formally, V is a Bochner space.

4.2. Averaging and the Structure of Function Spaces

We have seen how to define orbit averaging operators to produce invariant and equivariant functions as well as how to construct spaces of functions on which these operators can act. The reason for all of this is the following result, which shows that the averaging operators allow us to decompose any function in V into orthogonal $\mathcal{G}$-symmetric and $\mathcal{G}$-anti-symmetric parts. Recall that since $\mathcal{O}$ is just a special case of $\mathcal{Q}$, Lemma 1 applies to both operators.

---

[1] It is permissible for inner product itself to depend on the point x at which the feature maps are evaluated. The only requirement is that evaluating the inner product between two fixed vectors is a $\mathcal{G}$-invariant function $\langle a, b \rangle(x) = \langle a, b \rangle(gx), \forall a, b \in \mathbb{R}^k, g \in \mathcal{G}$ and $x \in \mathcal{X}$. We believe that this allows our results to extend to the case of the features defined as maps from a manifold to its tangent bundle.

[2] Equality is defined μ-almost-everywhere.

Lemma 1. Let U be any subspace of V that is closed under $\mathcal{Q}$. Define the subspaces S and A of, respectively, the $\mathcal{G}$-symmetric and $\mathcal{G}$-anti-symmetric functions in U: $S = \{f \in U : f \text{ is } \mathcal{G}\text{-equivariant}\}$ and $A = \{f \in U : \mathcal{Q}f = 0\}$. Then U admits admits an orthogonal decomposition into symmetric and anti-symmetric parts

$$U = S \oplus A.$$

The proof consists of establishing the following properties: (A) for any $f \in V$, $\mathcal{Q}f \in V$; (B) any $f \in V$ is $\mathcal{G}$-equivariant if and only if $\mathcal{Q}f = f$; (C) $\mathcal{Q}$ has only two eigenvalues, 1 and 0; and, (D) $\mathcal{Q}$ is self-adjoint with respect to $\langle \cdot, \cdot \rangle_\mu$. The last of these is critical and depends on the $\mathcal{G}$-invariance of μ. There are many tasks for which $\mathcal{G}$-invariance of the input distribution is a natural assumption, for instance in medical imaging (Winkels & Cohen, 2018).

Lemma 1 says that any function $u \in U$ can be written $u = s + a$, where s is $\mathcal{G}$-equivariant, $\mathcal{Q}a = 0$ and $\langle s, a \rangle_\mu = 0$. We refer to s and a as the symmetric and anti-symmetric parts of u. In general this does not imply that a is odd, that it outputs an anti-symmetric matrix or that it is negated by swapping two inputs. These are, however, special cases. If $\mathcal{G} = C_2$ acts by $x \mapsto -x$ then odd functions $f : \mathbb{R} \to \mathbb{R}$ will be anti-symmetric in the sense of this paper. If $\mathcal{G} = C_2$ acts on matrices by $M \mapsto M^\top$ then $f : M \mapsto \frac{1}{2}(M - M^\top)$ is also anti-symmetric. Finally, if $\mathcal{G} = S_n$ and $f : \mathbb{R}^n \to \mathbb{R}$ with $f(x_1, \ldots, x_j, x_{j+1}, \ldots, x_n) = -f(x_1, \ldots, x_{j+1}, x_j, \ldots, x_n)$, then f is anti-symmetric in the sense of this paper.

Although it is straightforward to demonstrate and has surely been observed before, we will see in the rest of the paper that the perspective provided by Lemma 1 is highly fruitful. Before that, we conclude this section with an example for intuition.

Example 2. Let V consist of all functions $f : \mathbb{R}^2 \to \mathbb{R}$ such that $\mathbb{E}[f(X)^2] < \infty$ where $X \sim \mathcal{N}(0, I_2)$. Let $\mathcal{G} = \mathrm{SO}(2)$ act by rotation about the origin, with respect to which the normal distribution is invariant. Using Lemma 1 we may write $V = S \oplus A$. Alternatively, consider polar coordinates (r, θ), then for any feature map $f(r, \theta)$ we have $\mathcal{O}f(r, \theta) = \frac{1}{2\pi} \int_0^{2\pi} f(r, \theta) \, d\theta$. So any $\mathcal{G}$-invariant feature map (i.e. anything in S) depends only on the radial coordinate. Similarly, any h for which $\mathcal{O}h = 0$ must have $\mathcal{O}h(r, \theta) = \frac{1}{2\pi} \int_0^{2\pi} h(r, \theta) \, d\theta = 0$ for any r, and A consists entirely of such functions. For example, $r^3 \cos \theta \in A$. We then recover $\langle s, h \rangle_\mu = \frac{1}{2\pi} \int_{\mathcal{X}} s(r)h(r, \theta)e^{-r^2/2}r \, dr \, d\theta = 0$ for any $s \in S$ by integrating h over θ. Intuitively, one can think of the functions in S as varying perpendicular to the flow of $\mathcal{G}$ on $\mathcal{X} = \mathbb{R}^2$ and so are preserved by it, while the functions in A average to 0 along this flow, see Fig. 1.
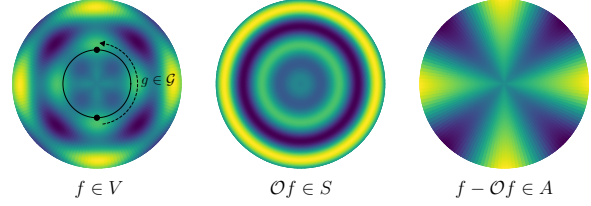

Figure 1. Example of a function decomposition. The figure shows $f(r, \theta) = r \cos(r - 2\theta) \cos(r + 2\theta)$ decomposed into its symmetric and anti-symmetric parts in $V = S \oplus A$ under the natural action of $\mathcal{G} = \mathrm{SO}(2)$ on $\mathbb{R}^2$. See Example 2. Best viewed in colour.

5. Feature Averaging

We remarked earlier that $\mathcal{O}$ can be thought of as performing feature averaging. Before our study of the generalisation of symmetric models, we use this interpretation to derive our first consequence of Lemma 1. We show that feature averaging can be viewed as solving a least squares problem in the space of features extractors V. That is, feature averaging sends f to $\bar{f}$, where $\bar{f}$ is the *closest $\mathcal{G}$-invariant feature extractor to f*.

Proposition 3 (Feature Averaging as a Least-Squares Problem). Let V be the space of all normalisable feature extractors as defined above. Define S and A as in Lemma 1. For any $f \in V$, feature averaging with $\mathcal{O}$ maps $f \mapsto \bar{f}$ where $\bar{f}$ is the (μ-a.e.) unique solution to the least-squares problem

$$\bar{f} = \operatorname*{argmin}_{s \in S} \|f - s\|_\mu^2.$$

Example 4. Consider again the setting of Example 2. For simplicity, let $f(r, \theta) = f_{\mathrm{rad}}(r) f_{\mathrm{ang}}(\theta)$ be separable in polar coordinates. Notice that $\mathcal{O}f = c_f f_{\mathrm{rad}}$ where $c_f = \frac{1}{2\pi} \int_0^{2\pi} f_{\mathrm{ang}}(\theta) \, d\theta$. Then for any $s \in S$ can calculate:

$$
\begin{aligned}
\|f - s\|_\mu^2 &= \frac{1}{2\pi} \int_{\mathcal{X}} (f(r, \theta) - s(r))^2 e^{-r^2/2} r \, dr \, d\theta \\
&= \frac{1}{2\pi} \int_{\mathcal{X}} (f(r, \theta) - c_f f_{\mathrm{rad}}(r))^2 e^{-r^2/2} r \, dr \, d\theta \\
&\quad + \frac{1}{2\pi} \int_{\mathcal{X}} (c_f f_{\mathrm{rad}}(r) - s)^2 e^{-r^2/2} r \, dr \, d\theta
\end{aligned}
$$

which is minimised by $s = c_f f_{\mathrm{rad}}$, as predicted.

5.1. Feature Averaging and Generalisation

We end our discussion of feature averaging with an analysis of its impact on generalisation. To do this we consider the reduction in the Rademacher complexity of a feature averaged class.

Rademacher Complexity Let $T = \{t_1, \ldots, t_n\}$ be a collection of points of some space $\mathcal{T}$. The empirical

Rademacher complexity of a set $\mathcal{F}$ of functions $f : \mathcal{T} \to \mathbb{R}$ evaluated on T is defined by

$$\widehat{\mathfrak{R}}_T(\mathcal{F}) = \mathbb{E}_{\boldsymbol{\sigma}} \left[\sup_{f \in \mathcal{F}} \left| \frac{1}{n} \sum_{i=1}^n \sigma_i f(t_i) \right| \right]$$

where $\sigma_i \sim \text{Unif}\{-1, 1\}$ for $i = 1, \ldots, n$ are Rademacher random variables over which the expectation $\mathbb{E}_{\boldsymbol{\sigma}}$ is taken. Let ν be a distribution on $\mathcal{T}$ and take $T \sim \nu^n$, in which case the empirical Rademacher complexity $\widehat{\mathfrak{R}}_T(\mathcal{F})$ is a random quantity. The Rademacher complexity $\mathfrak{R}_n(\mathcal{F})$ is defined by taking an expectation over T: $\mathfrak{R}_n(\mathcal{F}) = \mathbb{E}[\widehat{\mathfrak{R}}_T(\mathcal{F})]$. The Rademacher complexity appears in the study of generalisation in statistical learning, for instance see Wainwright (2019, Theorem 4.10 and Proposition 4.12).

Proposition 5. Let $\mathcal{G}$ be a compact group acting measurably on a set $\mathcal{T}$. Let $\mathcal{F}$ be a class of functions $f : \mathcal{T} \to \mathbb{R}$ and define the symmetric and anti-symmetric classes $\overline{\mathcal{F}} = \{\mathcal{O}f : f \in \mathcal{F}\}$ and $\mathcal{F}^\perp = \{f - \mathcal{O}f : f \in \mathcal{F}\}$. Let ν be a distribution over $\mathcal{T}$ that is $\mathcal{G}$-invariant. Then the Rademacher complexity of the feature averaged class satisfies

$$0 \le \mathfrak{R}_n(\mathcal{F}) - \mathfrak{R}_n(\overline{\mathcal{F}}) \le \mathfrak{R}_n(\mathcal{F}^\perp)$$

where the expectations in the definition of $\mathfrak{R}_n$ are taken over $t_i \sim \nu$ i.i.d..

Proposition 5 says that the Rademacher complexity is reduced by feature averaging, but not by more than the complexity of the anti-symmetric component of the class. This can be thought of as quantifying the benefit of enforcing invariance by averaging in terms of the extent to which the inductive bias is already present in the function class. Although the form of this result is suggestive, it does not imply a strict benefit. We provide stronger results in the following section.

6. Generalisation Benefit from Symmetric Models

In this section we apply Lemma 1 to derive a strict (i.e. non-zero) generalisation gap between models that have and have not been specified to have the invariance/equivariance that is present in the task. We start with the following general result, which equates the generalisation gap between any predictor and its closest equivariant function to the norm of the anti-symmetric component of the predictor.

Lemma 6. Let $X \sim \mu$ where μ is a $\mathcal{G}$-invariant distribution on $\mathcal{X}$. Let $Y = f^*(X) + \xi \in \mathbb{R}^k$, where ξ is a random element of $\mathbb{R}^k$ that is independent of X with zero mean and finite variance and $f^* : \mathcal{X} \to \mathbb{R}^k$ is $\mathcal{G}$-equivariant. Then,

for any $f \in V$, the generalisation gap satisfies

$$\Delta(f, \mathcal{Q}f) := \mathbb{E}[\|Y - f(X)\|_2^2] - \mathbb{E}[\|Y - \mathcal{Q}f(X)\|_2^2]$$
$$= \|f^\perp\|_\mu^2.$$

Lemma 6 demonstrates the existence of barrier in the generalisation of any predictor on a problem that has a symmetry. Notice that the barrier turns out to be the measure of how well the predictor encodes the symmetry. Clearly, the only way of overcoming this is to set $f^\perp = 0$ (μ-a.e.), which from Lemma 1 equivalent to enforcing $\mathcal{G}$-equivariance in f (μ-a.e.). Lemma 6 therefore provides a *strict generalisation benefit for equivariant predictors*.

In a sense, $\mathcal{Q}f$ is the archetypal equivariant predictor to which f should be compared. A trivial extension to Proposition 3 shows that $\mathcal{Q}f$ is the closest equivariant predictor to f and, more importantly, if h is a $\mathcal{G}$-equivariant predictor with smaller test risk than $\mathcal{Q}f$ then $\Delta(f, h) = \Delta(f, \mathcal{Q}f) + \Delta(\mathcal{Q}f, h) \ge \Delta(f, \mathcal{Q}f)$ which cannot weaken our result.

Later in this section we will use Lemma 6 to explicitly calculate the generalisation gap for invariant/equivariant linear models. We will see that it displays a natural relationship between the number of training examples and the dimension of the space of anti-symmetric models A, which is a property of the group action. Intuitively, the model needs enough examples to learn to be orthogonal to A.

This result also has a useful theoretical implication for test-time data augmentation, which is commonly used to increase test accuracy (Simonyan & Zisserman, 2015; Szegedy et al., 2015; He et al., 2016). Test-time augmentation consists of averaging the output of a learned function f over random transformations of the same input when making predictions at test-time. When the transformations belong to a group $\mathcal{G}$ and are sampled from its Haar measure, test-time averaging can be viewed as a Monte Carlo estimate of $\mathcal{O}f$. Lemma 6 then shows that test-time averaging is beneficial for generalisation when the target function is itself $\mathcal{G}$-invariant, regardless of the learned function f.

6.1. Regression with Invariant Target

Let $\mathcal{X} = \mathbb{R}^d$ with the Euclidean inner product and $\mathcal{Y} = \mathbb{R}$. Consider linear regression with the squared-error loss $\ell(y, y') = (y - y')^2$. Let $\mathcal{G}$ be a compact group that acts on $\mathcal{X}$ via an orthogonal representation $\phi : \mathcal{G} \to O(d)$ and let $X \sim \mu$ where μ is now an arbitrary $\mathcal{G}$-invariant probability distribution on $\mathcal{X}$ with $\Sigma := \mathbb{E}[XX^\top]$ finite and positive definite.[3] We consider linear predictors $h_w : \mathcal{X} \to \mathcal{Y}$ with

---

[3]If Σ is only positive semi-definite then the developments are similar. We assume $\Sigma > 0$ for simplicity.

$h_w(x) = w^\top x$ where $w \in \mathcal{X}$. Define the space of all linear predictors $V_{\text{lin}} = \{h_w : w \in \mathcal{X}\}$ which is a subspace of V. Notice that V_{lin} is closed under $\mathcal{O}$: for any $x \in \mathcal{X}$

$$
\begin{aligned}
\mathcal{O}h_w(x) &= \int_{\mathcal{G}} h_w(gx) \, \mathrm{d}\lambda(g) \\
&= \int_{\mathcal{G}} w^\top \phi(g) x \, \mathrm{d}\lambda(g) \\
&= \left(\int_{\mathcal{G}} \phi(g^{-1}) w \, \mathrm{d}\lambda(g) \right)^\top x \\
&= h_{\Phi_{\mathcal{G}}(w)}(x)
\end{aligned}
$$

where in the last line we substituted $g \mapsto g^{-1}$ and defined the linear map $\Phi_{\mathcal{G}} : \mathbb{R}^d \to \mathbb{R}^d$ by $\Phi_{\mathcal{G}}(w) = \int_{\mathcal{G}} \phi(g) w \, \mathrm{d}\lambda(g)$.[4] We also have

$$
\langle h_a, h_b \rangle_\mu = \int_{\mathcal{X}} a^\top x x^\top b \, \mathrm{d}\mu(x) = a^\top \Sigma b
$$

and we denote the induced inner product on $\mathcal{X}$ by $\langle a, b \rangle_\Sigma := a^\top \Sigma b$ and the corresponding norm by $\|\cdot\|_\Sigma$. Since V_{lin} is closed under $\mathcal{O}$ we can apply Lemma 1 to decompose $V_{\text{lin}} = S \oplus A$ where the orthogonality is with respect to $\langle \cdot, \cdot \rangle_\mu$. It follows that we can write any $h_w \in V_{\text{lin}}$ as

$$
h_w = \overline{h_w} + h_w^\perp
$$

where we have shown that there must exist $\bar{w}, w^\perp \in \mathcal{X}$ with $\langle \bar{w}, w^\perp \rangle_\Sigma = 0$ such that $\overline{h_w} = h_{\bar{w}}$ and $h_w^\perp = h_{w^\perp}$. By choosing a basis for $\mathcal{X}$, there is a natural bijection $\mathcal{X} \to V_{\text{lin}}$ where $w \mapsto h_w$. Using this identification, we abuse notation slightly and write $\mathcal{X} = S \oplus A$ to represent the induced structure on $\mathcal{X}$.

Suppose examples are labelled by a target function $h_\theta \in V_{\text{lin}}$ that is $\mathcal{G}$-invariant. Let $X \sim \mu$ and $Y = \theta^\top X + \xi$ where ξ is independent of X, has mean 0 and finite variance. Recall the definition of the *generalisation gap* between predictors as the difference in their test errors. We study the generalisation gap $\Delta(h_w, h_{\bar{w}})$ between predictors h_w and $h_{\bar{w}}$ defined above. Lemma 6 gives $\Delta(h_w, h_{\bar{w}}) = \|h_{w^\perp}\|_\mu^2 = \|w^\perp\|_\Sigma^2$. In Theorem 7 we calculate this quantity where w is the minimum-norm least-squares estimator and $\bar{w} = \Phi_{\mathcal{G}}(w)$. To the best of our knowledge, this is the first result to specify exactly the generalisation benefit for invariant models.

Theorem 7. Let $\mathcal{X} = \mathbb{R}^d$, $\mathcal{Y} = \mathbb{R}$ and let $\mathcal{G}$ be a compact group with an orthogonal representation ϕ on $\mathcal{X}$. Let $X \sim \mathcal{N}(0, \sigma_X^2 I)$ and $Y = h_\theta(X) + \xi$ where $h_\theta(x) = \theta^\top x$ is $\mathcal{G}$-invariant with $\theta \in \mathbb{R}^d$ and ξ has mean 0, variance $\sigma_\xi^2 < \infty$ and is independent of X. Let w be the least-squares estimate

---
[4]Since $\mathcal{G}$ is compact it is unimodular and this change of variables is valid, e.g. see Folland (2016, Corollary 2.28) and adjacent results.

of θ from i.i.d. examples $\{(X_i, Y_i) : i = 1, \ldots, n\}$ and let A be the orthogonal complement of the subspace of $\mathcal{G}$-invariant linear predictors (c.f. Lemma 1).

- If $n > d + 1$ then the generalisation gap is
$$
\mathbb{E}[\Delta(h_w, h_{\bar{w}})] = \sigma_\xi^2 \frac{\dim A}{n - d - 1}.
$$

- At the interpolation threshold $n \in [d-1, d+1]$, if h_w is not $\mathcal{G}$-invariant then the generalisation gap diverges to ∞.

- If $n < d - 1$ the generalisation gap is
$$
\begin{aligned}
&\mathbb{E}[\Delta(h_w, h_{\bar{w}})] \\
&= \dim A \left(\frac{\sigma_X^2 \|\theta\|_2^2 \, n(d-n)}{d(d-1)(d+2)} + \frac{\sigma_\xi^2 \, n}{d(d-n-1)} \right).
\end{aligned}
$$

In each case, the generalisation gap has a term of the form $\sigma_\xi^2 r(n, d) \dim A$ that arises due to the noise in the target distribution. In the overparameterised setting $d > n + 1$ there is an additional term (the first) that represents the generalisation gap in the noiseless setting $\xi \overset{\text{a.s.}}{=} 0$. This term is the error in the least-squares estimate of θ in the noiseless problem, which of course vanishes in the fully determined case $n > d + 1$. In addition, the divergence at the so called interpolation threshold $n \approx d$ is consistent with the literature on double descent (Hastie et al., 2019).

Notice the central role of $\dim A$ in Theorem 7. This quantity is a property of the group action as it describes the codimension of the set of invariant models. The generalisation gap is then dictated by how 'significant' the symmetry is to the problem. We give two examples that represent extremal cases of this 'significance'.

Example 8 (Permutations, $\dim A = d - 1$)**.** Let S_d act on $\mathcal{X} = \mathbb{R}^d$ by permutation of the coordinates, so $(\phi(\rho)w)_i = w_{\rho(i)}$ for $\rho \in S_d$. Observe that, since the Haar measure λ is uniform on S_d, for any $i = 1, \ldots, d$

$$
\Phi_{S_d}(w)_i = \frac{1}{d!} \sum_{\rho \in S_d} w_{\rho(i)} = \frac{1}{d} \sum_j w_j
$$

so S is the one dimensional subspace $\{t(1, \ldots, 1)^\top : t \in \mathbb{R}\}$. Since $\mathcal{X} = S \oplus A$ we get $\dim A = d - 1$.

Example 9 (Reflection, $\dim A = 1$)**.** Let C_2 be the cyclic group of order 2 and let it act on $\mathcal{X} = \mathbb{R}^d$ by reflection in the first coordinate. A is then the subspace consisting of w such that for any $j = 2, \ldots, d$

$$
\Phi_{C_2}(w)_j = \frac{1}{|C_2|} \sum_{g \in C_2} (\phi(g)w)_j = w_j = 0
$$

since the action fixes all coordinates apart from the first. Hence $A = \{t(1, 0, \ldots, 0)^\top : t \in \mathbb{R}\}$.

6.2. Regression with Equivariant Target

One can apply the same construction to equivariant models. Assume the same setup, but now let $\mathcal{Y} = \mathbb{R}^k$ with the Euclidean inner product and let the space of predictors be $W_{\text{lin}} = \{f_W : \mathbb{R}^d \to \mathbb{R}^k, \ f_W(x) = W^\top x : W \in \mathbb{R}^{d \times k}\}$. We consider linear regression with the squared-error loss $\ell(y, y') = \|y - y'\|_2^2$. Let ψ be an orthogonal representation of $\mathcal{G}$ on $\mathcal{Y}$. We define the linear map, which we call the intertwining average, $\Psi_{\mathcal{G}} : \mathbb{R}^{d \times k} \to \mathbb{R}^{d \times k}$ by[5]

$$\Psi_{\mathcal{G}}(W) = \int_{\mathcal{G}} \phi(g) W \psi(g^{-1}) \, d\lambda(g).$$

Similarly, define the intertwining complement as $\Psi_{\mathcal{G}^\perp} : \mathbb{R}^{d \times k} \to \mathbb{R}^{d \times k}$ by $\Psi_{\mathcal{G}^\perp}(W) = W - \Psi_{\mathcal{G}}(W)$. We establish the following results, which are in fact generalisations of the invariant case. In the proofs we will leverage the expression of $\Psi_{\mathcal{G}}$ as a 4-tensor with components $\Psi_{\mathcal{G}\,abce} = \int_{\mathcal{G}} \phi(g)_{ac} \psi(g)_{be} \, d\lambda(g)$ where $a, c = 1, \ldots d$ and $b, e = 1, \ldots, k$.[6]

Proposition 10. For any $f_W \in W_{\text{lin}}$, $\mathcal{Q} f_W = f_{\Psi_{\mathcal{G}}(W)}$ and hence W_{lin} is closed under $\mathcal{Q}$.

Proposition 11. The inner product on W_{lin} satisfies, for any $f_{W_1}, f_{W_2} \in W_{\text{lin}}$,

$$\langle f_{W_1}, f_{W_2} \rangle_\mu = \text{Tr}(W_1^\top \Sigma W_2)$$

where $\Sigma = \mathbb{E}[X X^\top]$ and $X \sim \mu$.

Proposition 10 allows us to apply Lemma 1 to write $W_{\text{lin}} = S \oplus A$, so for any $f_W \in W_{\text{lin}}$ there exists $\overline{f_W} \in S$ and $f_W^\perp \in A$ with $\langle \overline{f_W}, f_W^\perp \rangle_\mu = 0$. The corresponding parameters $\overline{W} = \Psi_{\mathcal{G}}(W)$ and $W^\perp = \Psi_{\mathcal{G}^\perp}(W)$ must therefore satisfy $\text{Tr}(\overline{W}^\top \Sigma W^\perp) = 0$, with Σ defined as in Proposition 11. Repeating our abuse of notation, we identify $\mathbb{R}^{d \times k} = S \oplus A$ with $S = \Psi_{\mathcal{G}}(\mathbb{R}^{d \times k})$ and A its orthogonal complement with respect to the induced inner product.

Proposition 12. Let $X \sim \mu$ and let ξ a random element of $\mathbb{R}^k$ that is independent of X with $\mathbb{E}[\xi] = 0$ and finite variance. Set $Y = f_{\Theta}(X) + \xi$ where f_{Θ} is $\mathcal{G}$-equivariant. For any $f_W \in W_{\text{lin}}$, the generalisation gap satisfies

$$\Delta(f_W, f_{\overline{W}}) := \mathbb{E}[\|Y - f_W(X)\|_2^2] - \mathbb{E}[\|Y - f_{\overline{W}}(X)\|_2^2]$$

$$= \|\Sigma^{1/2} W^\perp\|_F^2$$

where $\overline{W} = \Psi_{\mathcal{G}}(W)$, $W^\perp = \Psi_{\mathcal{G}^\perp}(W)$ and $\Sigma = \mathbb{E}[X X^\top]$.

---

[5]The reader may have noticed that we define $\Psi_{\mathcal{G}}$ backwards, in the sense that its image contains maps that are equivariant in the direction $\psi \to \phi$. This is because of the transpose in the linear model, which is there for consistency with the $k = 1$ invariance case. This choice is arbitrary and gives no loss in generality.

[6]If necessary, the reader can see the supplementary material for a derivation.

Having followed the same path as the previous section, we provide a characterisation of the generalisation benefit of equivariance. In the same fashion, we compare the least-squares estimate W with its equivariant version $\overline{W} = \Psi_{\mathcal{G}}(W)$. As we explained at the beginning of the section, the choice of $\overline{W} = \Psi_{\mathcal{G}}(W)$ is natural and costs us nothing.

Theorem 13. Let $\mathcal{X} = \mathbb{R}^d$, $\mathcal{Y} = \mathbb{R}^k$ and let $\mathcal{G}$ be a compact group with orthogonal representations ϕ on $\mathcal{X}$ and ψ on $\mathcal{Y}$. Let $X \sim \mathcal{N}(0, \sigma_X^2 I_d)$ and $Y = h_{\Theta}(X) + \xi$ where $h_{\Theta}(x) = \Theta^\top x$ is $\mathcal{G}$-equivariant and $\Theta \in \mathbb{R}^{d \times k}$. Assume ξ is a random element of $\mathbb{R}^k$, independent of X, with mean 0 and $\mathbb{E}[\xi \xi^\top] = \sigma_\xi^2 I_k < \infty$. Let W be the least-squares estimate of Θ from n i.i.d. examples $\{(X_i, Y_i) : i = 1, \ldots, n\}$ and let $(\chi_\psi | \chi_\phi) = \int_{\mathcal{G}} \chi_\psi(g) \chi_\phi(g) \, d\lambda(g)$ denote the scalar product of the characters of the representations of $\mathcal{G}$.

- If $n > d + 1$ the generalisation gap is

$$\mathbb{E}[\Delta(f_W, f_{\overline{W}})] = \sigma_\xi^2 \frac{dk - (\chi_\psi | \chi_\phi)}{n - d - 1}.$$

- At the interpolation threshold $n \in [d - 1, d + 1]$, if f_W is not $\mathcal{G}$-equivariant then the generalisation gap diverges to ∞.

- If $n < d - 1$ then the generalisation gap is

$$\mathbb{E}[\Delta(f_W, f_{\overline{W}})] =$$
$$\sigma_X^2 \frac{n(d-n)}{d(d-1)(d+2)} \left((d+1)\|\Theta\|_F^2 - \text{Tr}(J_{\mathcal{G}} \Theta^\top \Theta) \right)$$
$$+ \sigma_\xi^2 \frac{n(dk - (\chi_\psi | \chi_\phi))}{d(d - n - 1)}$$

where each term is non-negative and $J_{\mathcal{G}} \in \mathbb{R}^{k \times k}$ is given by

$$J_{\mathcal{G}} = \int_{\mathcal{G}} (\chi_\phi(g) \psi(g) + \psi(g^2)) \, d\lambda(g).$$

Theorem 13 is a direct generalisation of Theorem 7. As we remarked in the introduction, $dk - (\chi_\psi | \chi_\phi)$ plays the role of $\dim A$ in Theorem 7 and is a measure of the significance of the symmetry to the problem. The dimension of W_{lin} is dk, while $(\chi_\psi | \chi_\phi)$ is the dimension of the space of equivariant maps. In our notation $(\chi_\psi | \chi_\phi) = \dim S$.

Just as with Theorem 7, there is an additional term (the first) in the overparameterised case $d > n + 1$ that represents the estimation in a noiseless setting $\xi \overset{\text{a.s.}}{=} 0$. Notice that if $k = 1$ and ψ is trivial we find

$$J_{\mathcal{G}} = \int_{\mathcal{G}} \chi_\phi(g) \, d\lambda(g) + 1 = (\chi_\phi | 1) + 1 = \dim S + 1$$

which confirms that Theorem 13 reduces exactly to Theorem 7.

Interestingly, the first term in the $d > n+1$ case can be made independent of ψ, since the equivariance of h_Θ implies

$$\mathrm{Tr}(J_\mathcal{G}\Theta^\top\Theta) = \mathrm{Tr}(\Theta^\top J_\phi\Theta)$$

where

$$J_\phi = \int_\mathcal{G} (\chi_\phi(g)\phi(g) + \phi(g^2))\,\mathrm{d}\lambda(g).$$

Finally, we remark that Theorem 13 is possible for more general probability distributions on X. For instance, it sufficient that the distribution is absolutely continuous with respect to the Lebesgue measure, has finite variance and is $O(d)$ invariant. The final condition implies the existence of a scalar r_n such that $\mathbb{E}[(\boldsymbol{X}^\top\boldsymbol{X})^+] = r_n I_d$ where $\boldsymbol{X} \in \mathbb{R}^{n\times d}$ are the row-stacked training inputs as defined in the proof.

7. Neural Networks

In this section we discuss how the insights of this paper apply to neural networks and raise some open questions for future work. Let $F : \mathbb{R}^d \to \mathbb{R}^k$ be a feedforward neural network with L layers, layer widths κ_i $i = 1, \dots, L$ and weights $W^i \in \mathbb{R}^{\kappa_i \times \kappa_{i+1}}$ for $i = 1, \dots, L$ where $\kappa_1 = d$ and $\kappa_{L+1} = k$. We will assume F has the form

$$F(x) = W^L\sigma(W^{L-1}\sigma(\dots\sigma(W^1 x)\dots)) \qquad (1)$$

where σ is an element-wise non-linearity.

7.1. Invariant and Equivariant Networks

The standard method for engineering neural networks to be invariant/equivariant to the action of a finite group on the inputs is weight tying. This method has been around for a while (Wood & Shawe-Taylor, 1996) but has come to recent attention via Ravanbakhsh et al. (2017). We will briefly describe this approach, its connections to Theorems 7 and 13 and how the ideas of this paper can be used to find new algorithms for both enforced and learned invariance/equivariance. We leave analyses of these suggested approaches to future work.

The methods of Wood & Shawe-Taylor (1996); Ravanbakhsh et al. (2017) can be described as follows. Let $\mathcal{G}_\mathrm{fin}$ be a finite group. For each $i = 2, \dots, L+1$, the user chooses a matrix representation $\psi_i : \mathcal{G}_\mathrm{fin} \to \mathrm{GL}_{k_i}(\mathbb{R})$ of $\mathcal{G}_\mathrm{fin}$ that acts on the inputs for each layer $i = 2, \dots, L$ and on the outputs of the network when $i = L + 1$.[7] For $i = 2, \dots, L$, these representations must be chosen such that they commute with the activation function

$$\sigma(\psi_i(g)\cdot) = \psi_i(g)\sigma(\cdot) \qquad (2)$$

$\forall g \in \mathcal{G}_\mathrm{fin}$.[8] One then chooses weights for the network such that at each layer and $\forall g \in \mathcal{G}_\mathrm{fin}$

$$W^i\psi_i(g) = \psi_{i+1}(g)W^i. \qquad (3)$$

By induction on the layers, satisfying Eqs. (2) and (3) ensures that the network is $\mathcal{G}_\mathrm{fin}$-equivariant. Invariance occurs when ψ_{L+1} is the trivial representation.

The condition in Eq. (3) can be phrased as saying that that W^i belongs to the space of *intertwiners* of the representations ψ_i and ψ_{i+1}. By denoting the space of all weight matrices in layer i as $U = \mathbb{R}^{\kappa_i \times \kappa_{i+1}}$, the space of intertwiners is immediately recognisable as $S = \Psi_\mathcal{G}(U)$ from Lemma 1.

Typically, the practitioner will hand-engineer the structure of weight matrices to belong to the correct intertwiner space. In the following sections we will propose alternative procedures that build naturally on the ideas of this paper. Moreover, as a benefit of our framework, these new approaches extend weight-tying to any compact group that admits a finite-dimensional, real representation. We end this section with a bound on the sample complexity of invariant networks, which follows from Bartlett et al. (2019). Similar results are possible for different activation functions.

Lemma 14. Let $\mathcal{G}$ be a compact group with layer-wise representations as described. Let $F : \mathbb{R}^d \to \mathbb{R}$ be a $\mathcal{G}$-invariant neural network with ReLU activation and weights that intertwine the representations. Let $\mathcal{H}$ be the class of all functions realisable by this network. Then

$$\mathrm{VC}(\mathcal{H}) \leq L + \frac{1}{2}\alpha(F)L(L+1)\max_{1\leq i\leq L}(\chi_i|\chi_{i+1})$$

where $\alpha(F) = \log_2\left(4\mathrm{e}\log_2\left(\sum_{i=1}^L 2\mathrm{e}i\kappa_i\right)\sum_{i=1}^L i\kappa_i\right)$.

Example 15 (Permutation invariant networks)**.** Permutation invariant networks (and permutation invariant functions more generally) are studied in Wood & Shawe-Taylor (1996); Zaheer et al. (2017); Bloem-Reddy & Teh (2020) and many other works, see references therein. In particular, multiple authors have given the form of a permutation equivariant weight matrix as $W = \lambda I + \gamma\mathbf{1}\mathbf{1}^\top$ for scalars $\alpha, \beta \in \mathbb{R}$ and with $\mathbf{1} = (1, \dots, 1)^\top$. Consider an L-layer ReLU network with, for simplicity, widths $\kappa_i = d$ $\forall i$. Let $\mathcal{H}$ be the class of all functions realisable by this network, then $\mathrm{VC}(\mathcal{H}) = O(L^2\log(Ld\log(Ld)))$.

7.2. Projected Gradients

As we have seen, provided that the activation function satisfies Eq. (2), specifying the weight matrices to intertwine

---

[7] ψ_1 is the representation on the inputs, which we consider as an aspect of the task and not a design choice.

[8] This condition is somewhat restrictive, but note that a permutation representation will commute with any element-wise non-linearity.

between layer-wise representations is sufficient to ensure equivariance in a neural network. We have also seen from Section 6.2 that it is possible to project any weight matrix into an intertwiner space using $\Psi_{\mathcal{G}}$. For each layer l of the network we have a linear map $\Psi_{\mathcal{G}}{}^l$, which is a 4-tensor with components $\Psi_{\mathcal{G}}{}^l_{abce} = \int_{\mathcal{G}} \psi_{l+1}(g)_{ac} \psi_l(g)_{be} \, d\lambda(g)$.

The tensors $\{\Psi_{\mathcal{G}}{}^l : l = 1, \dots, L\}$ depend only on the representations and so can be computed before training. One can therefore obtain invariant/equivariant networks by a form of projected gradient descent. Explicitly, with loss ℓ and learning rate η, the update rule for the l^{th} layer is

$$\widetilde{W}^l(t+1) = W^l(t) - \eta \nabla_{W^l} \ell(W^1(t), \dots, W^L(t))$$
$$W^l(t+1) = \Psi_{\mathcal{G}}{}^l \left(\widetilde{W}^l(t+1) \right).$$

If Eq. (2) holds the network will be exactly invariant/equivariant after any iteration.

7.3. Regularisation for Equivariance

We have seen from Lemma 1 and Section 6.2 that any weight matrix can be written $W = \overline{W} + W^\perp$ where $\overline{W} = \Psi_{\mathcal{G}}(W)$ belongs to an intertwiner space (so is equivariant) and $W^\perp = \Psi_{\mathcal{G}}{}^\perp(W)$ belongs to an orthogonal space that parametrises the anti-symmetric linear maps. This suggests a method of learned invariance/equivariance by using a regularisation term of the form $\sum_{l=1}^{L} \|\Psi_{\mathcal{G}}{}^{l\perp}(W^l)\|_{\mathrm{F}}^2$. Where the 4-tensor $\Psi_{\mathcal{G}}{}^{l\perp}$ has components $\Psi_{\mathcal{G}}{}^{l\perp}_{abce} = \delta_{ac}\delta_{be} - \Psi_{\mathcal{G}}{}^l_{abce}$ and can be computed before training. If $\Psi_{\mathcal{G}}{}^{l\perp}(W^l) = 0$ for $l = 1, \dots, L$ and the activation function satisfies Eq. (2), then the resulting network will be exactly $\mathcal{G}$-invariant/equivariant. This method could also allow for learned/approximate invariance. Indeed, Proposition 16 suggests $\|\Psi_{\mathcal{G}}{}^{l\perp}(W)\|_{\mathrm{F}}^2$ as a measure of the layer-wise invariance/equivariance of the network.

Proposition 16. Let $\mathcal{G}$ be a compact group. Let $f_W : \mathbb{R}^d \to \mathbb{R}^k$ with $f_W(x) = \sigma(Wx)$ be a single neural network layer with C-Lipschitz, element-wise activation σ. Let $\phi : \mathcal{G} \to O(d)$ and $\psi : \mathcal{G} \to O(k)$ be orthogonal representations of $\mathcal{G}$ on the input and output spaces respectively and assume that ψ commutes with σ as in Eq. (2). Let $X \in \mathbb{R}^d$, $X \sim \mathcal{N}(0, I_d)$. We can consider the network as belonging to V from Section 4 with $\mu = \mathcal{N}(0, I_d)$. Write $V = S \oplus A$, where S contains the equivariant functions in V, then

$$\inf_{s \in S} \mathbb{E}[\|f_W(X) - s(X)\|_2^2] \leq 2C^2 \|W^\perp\|_{\mathrm{F}}^2.$$

Proposition 16 shows that the distance between the outputs of a single layer neural network and its closest equivariant function is bounded by the norm of the anti-symmetric component of the weights $W^\perp$. This quantity can be interpreted as a measure of the equivariance of the

layer and regularising $\|W^\perp\|_{\mathrm{F}}$ will encourage the network to become (approximately) equivariant. It is easy to generalise Proposition 16 so that X follows any $\mathcal{G}$-invariant distribution with finite second moment.

7.4. Open Questions

Equivariant Convolutions There has been much work on engineering convolutional layers to be group equivariant, for instance Cohen & Welling (2016); Cohen et al. (2018); Kondor & Trivedi (2018); Cohen et al. (2019). The convolution is a linear operator parameterised by the kernel. This suggests that it may be possible to analyse the generalisation properties of group equivariant convolutions in the framework of Lemma 1, similar to Section 6.

Invariant/Equivariant Networks We have discussed enforcing invariance/equivariance in a neural network $F_{(W^1,\dots,W^L)}$ (with the dependence on the weights now explicit) by restricting weight matrices to intertwine between representations at each layer. We ask: is this the best way to encode symmetry? Mathematically, let $X \sim \mu$ with $\mathcal{G}$-invariant μ and embed the functions realised by the network in $V = S \oplus A$. Given an invariant/equivariant target $s \in S$, must the best approximating neural network be layer-wise invariant/equivariant? That is, are there $s \in S$ such that the following holds

$$\inf_{\mathcal{W}} \mathbb{E}[\|F_{(W^1,\dots,W^L)}(X) - s(X)\|_2^2]$$
$$< \inf_{\mathcal{U}} \mathbb{E}[\|F_{(U^1,\dots,U^L)}(X) - s(X)\|_2^2],$$

where $\mathcal{W} = \{W^l \in \mathbb{R}^{\kappa_l \times \kappa_{l+1}} : l = 1, \dots, L\}$ is the set of all possible weight matrices and $\mathcal{U} = \{U^l \in \Psi_{\mathcal{G}}{}^l(\mathbb{R}^{\kappa_l \times \kappa_{l+1}}) : l = 1, \dots, L\}$ is the set of all weight matrices restricted to be intertwiners? A resolution to this might shed light on new ways of encoding symmetry in neural networks.

Acknowledgements

References

Abu-Mostafa, Y. S. Hints and the VC dimension. *Neural Computation*, 5(2):278–288, 1993.

Anselmi, F., Leibo, J. Z., Rosasco, L., Mutch, J., Tacchetti,

A., and Poggio, T. Unsupervised learning of invariant representations in hierarchical architectures, 2014.

Bartlett, P. L., Harvey, N., Liaw, C., and Mehrabian, A. Nearly-tight VC-dimension and pseudodimension bounds for piecewise linear neural networks. *J. Mach. Learn. Res.*, 20:63–1, 2019.

Bloem-Reddy, B. and Teh, Y. W. Probabilistic symmetries and invariant neural networks. *Journal of Machine Learning Research*, 21(90):1–61, 2020. URL http://jmlr.org/papers/v21/19-322.html.

Cohen, T. and Welling, M. Group equivariant convolutional networks. In *International conference on machine learning*, pp. 2990–2999, 2016.

Cohen, T. S., Geiger, M., Köhler, J., and Welling, M. Spherical cnns. *arXiv preprint arXiv:1801.10130*, 2018.

Cohen, T. S., Geiger, M., and Weiler, M. A general theory of equivariant cnns on homogeneous spaces. In *Advances in Neural Information Processing Systems*, pp. 9145–9156, 2019.

Folland, G. B. *A course in abstract harmonic analysis*, volume 29. CRC press, 2016.

Haasdonk, B., Vossen, A., and Burkhardt, H. Invariance in kernel methods by haar integration kernels. In *SCIA 2005, Scandinavian Conference on Image Analysis*, pp. 841–851. Springer-Verlag, 2005.

Hastie, T., Montanari, A., Rosset, S., and Tibshirani, R. J. Surprises in high-dimensional ridgeless least squares interpolation. *arXiv preprint arXiv:1903.08560*, 2019.

He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 770–778, 2016. doi: 10.1109/CVPR.2016.90.

Kallenberg, O. *Foundations of modern probability*. Springer Science & Business Media, 2006.

Kondor, R. and Trivedi, S. On the generalization of equivariance and convolution in neural networks to the action of compact groups. In *International Conference on Machine Learning*, pp. 2747–2755, 2018.

Lyle, C., Kwiatkowksa, M., and Gal, Y. An analysis of the effect of invariance on generalization in neural networks. In *International conference on machine learning Workshop on Understanding and Improving Generalization in Deep Learning*, 2019.

Lyle, C., van der Wilk, M., Kwiatkowska, M., Gal, Y., and Bloem-Reddy, B. On the benefits of invariance in neural networks, 2020.

Maron, H., Fetaya, E., Segol, N., and Lipman, Y. On the universality of invariant networks. In *International Conference on Machine Learning*, pp. 4363–4371, 2019.

Mroueh, Y., Voinea, S., and Poggio, T. A. Learning with group invariant features: A kernel perspective. In *Advances in Neural Information Processing Systems*, pp. 1558–1566, 2015.

Pfau, D., Spencer, J. S., Matthews, A. G., and Foulkes, W. M. C. Ab initio solution of the many-electron schrödinger equation with deep neural networks. *Physical Review Research*, 2(3):033429, 2020.

Ravanbakhsh, S., Schneider, J., and Poczos, B. Equivariance through parameter-sharing. In *International Conference on Machine Learning*, pp. 2892–2901. PMLR, 2017.

Sannai, A. and Imaizumi, M. Improved generalization bound of group invariant / equivariant deep networks via quotient feature space, 2019.

Schölkopf, B., Burges, C., and Vapnik, V. Incorporating invariances in support vector learning machines. pp. 47–52. Springer, 1996.

Schulz-Mirbach, H. On the existence of complete invariant feature spaces in pattern recognition. In *International Conference On Pattern Recognition*, pp. 178–178. Citeseer, 1992.

Schulz-Mirbach, H. Constructing invariant features by averaging techniques. In *Proceedings of the 12th IAPR International Conference on Pattern Recognition, Vol. 3 - Conference C: Signal Processing (Cat. No.94CH3440-5)*, volume 2, pp. 387–390 vol.2, 1994.

Serre, J.-P. *Linear representations of finite groups*. Graduate texts in mathematics ; 42. Springer-Verlag, New York, 1977. ISBN 9780387901909.

Simonyan, K. and Zisserman, A. Very deep convolutional networks for large-scale image recognition. In *International Conference on Learning Representations*, 2015.

Sokolic, J., Giryes, R., Sapiro, G., and Rodrigues, M. Generalization error of invariant classifiers. In *Artificial Intelligence and Statistics*, pp. 1094–1103, 2017.

Szegedy, C., Liu, W., Jia, Y., Sermanet, P., Reed, S., Anguelov, D., Erhan, D., Vanhoucke, V., and Rabinovich, A. Going deeper with convolutions. In *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 1–9, 2015. doi: 10.1109/CVPR.2015.7298594.

Wadsley, S. Lecture notes on representation theory, October 2012. URL https://www.dpmms.cam.ac.uk/~sjw47/RepThLecturesMich2012.pdf.

Wainwright, M. J. *High-dimensional statistics: A non-asymptotic viewpoint*, volume 48. Cambridge University Press, 2019.

Winkels, M. and Cohen, T. S. 3d g-cnns for pulmonary nodule detection. *arXiv preprint arXiv:1804.04656*, 2018.

Wood, J. and Shawe-Taylor, J. Representation theory and invariant neural networks. *Discrete applied mathematics*, 69(1-2):33–60, 1996.

Xu, H. and Mannor, S. Robustness and generalization. *Machine learning*, 86(3):391–423, 2012.

Yarotsky, D. Universal approximations of invariant maps by neural networks, 2018.

Zaheer, M., Kottur, S., Ravanbakhsh, S., Poczos, B., Salakhutdinov, R. R., and Smola, A. J. Deep sets. In *Advances in neural information processing systems*, pp. 3391–3401, 2017.