On Measuring Excess Capacity in Neural Networks

Anonymous Author(s)
Affiliation
Address
email

Abstract

We study the *excess capacity* of deep networks in the context of supervised classification. That is, given a capacity measure of the underlying hypothesis class – in our case, empirical Rademacher complexity – by how much can we (a priori) constrain this class while retaining an empirical error on a par with the unconstrained regime? To assess excess capacity in modern architectures (such as *residual networks*), we extend and unify prior Rademacher complexity bounds to accommodate function *composition* and *addition*, as well as the structure of convolutions. The capacity-driving terms in our bounds are the Lipschitz constants of the layers and a (2,1) group norm distance to the initializations of the convolution weights. Experiments on benchmark datasets of varying task difficulty indicate that (1) there is a substantial amount of excess capacity per task, and (2) capacity can be kept at a surprisingly similar level across tasks. Overall, this suggests a notion of *compressibility* with respect to weight norms, orthogonal to classic compression via weight pruning.

1 Introduction

Understanding the generalization behavior of deep networks in supervised classification is still a largely open problem, despite a long history of theoretical advances. The observation that (over-parametrized) models can easily fit—i.e., reach zero training error—to randomly permuted training labels [39, 40] but, when trained on unpermuted labels, yield good generalization performance, has fueled much of the progress in this area. Recent works range from relating generalization to weight norms [3, 11, 22, 32, 33], measures of the distance to initializations [31], implicit regularization induced by the optimization algorithm [7, 35], or model compression [2, 6, 36]. Other works study connections to optimal transport [8], or generalization in the neural tangent kernel setting [1, 20].

When seeking to establish generalization guarantees within the classic uniform convergence regime, bounding a capacity measure, such as the Rademacher complexity [4], of the hypothesis class is the crucial step. While the resultant generalization bounds are typically vacuous and can exhibit concerning behavior [30], the *capacity bounds* themselves offer invaluable insights through the behavior of the bound-driving quantities, such as various types of weight norms or Lipschitz constants. Particularly relevant to our work is the observation that the bound-driving quantities tend to increase with *task difficulty*. The figure to the right illustrates this behavior for a ResNet18 (Section 4) in terms of Lipschitz constants per layer and the distance of each layer's weight to its initialization (measured via a group norm we develop in Section 3.1). This raises two immediate questions: First (**Q1**), can a network maintain empirical testing performance at substantially lower capacity? Second, (**Q2**) is any imposed capacity reduction inevitably tied to task difficulty, i.e., would a reduced-capacity model for an easy task fail on a difficult task? Both of these questions aim at the amount of "unneeded" capacity, which we refer to as *excess capacity* in the remainder of this work.

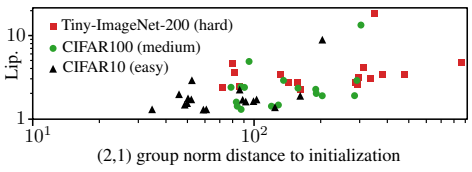

Submitted to 36th Conference on Neural Information Processing Systems (NeurIPS 2022). Do not distribute.

We will address questions (**Q1**) and (**Q2**) by means of controlling the empirical Rademacher complexity of a neural network. To this end, we consolidate and extend prior results from the literature on Rademacher complexity bounds to accommodate a broad range of network components in a unified way, including convolutions and skip connections, two ubiquitous elements in state-of-the-art models.

Our **contributions** can be summarized as follows:

1. We establish two bounds (in Theorem 3.5) for the empirical Rademacher complexity of networks that use *convolutions* and implement functions built from *composition* and *addition*. Specifically, we introduce two novel, convolution-specific, single-layer covering number bounds in Section 3.2 and contrast them to prior art, then modularize the single-layer to multi-layer covering approach of Bartlett et al. [3] in Section 3.3, and eventually present one incarnation of our framework for convolutional residual networks in Section 3.4.
2. We present an extensive set of experiments (in Section 4) with a ResNet18 model across benchmark datasets of varying task difficulty, demonstrating that model capacity, when measured via our weight norm based bound, (1) can be kept surprisingly small per task, and (2) can be kept at *roughly the same level* regardless of task difficulty. Both observations suggest *compressibility* of neural networks with respect to weight norms, orthogonal to the well-known compressibility with respect to the number of parameters [2, 36].

2 Related Work

Many prior works establish uniform-convergence type generalization bounds for neural networks through Rademacher complexity analysis. We review such approaches, highlighting challenges that arise with modern network architectures and the peculiarities of convolution layers.

One direct approach to bound the empirical Rademacher complexity is via a *layer-peeling* strategy [11, 33, 38], where the Rademacher complexity of L-layer networks is expressed by a factor times the Rademacher complexity of $(L-1)$-layer networks; in other words, the last layer is peeled off. This factor is typically a matrix $l_{p,q}$ norm, and thus the bounds usually scale with the product of these norms. Notably, the nonlinearities need to be *element-wise* operations, and some approaches only work for specific nonlinearities, such as ReLUs or asymmetric activations. A second strategy is to bound the empirical Rademacher complexity via a *covering numbers* approach [3, 22, 25, 41], typically achieved via Dudley's entropy integral [10]. This strategy is particularly flexible as it allows for arbitrary (but fixed) nonlinearities and for various paths to bounding the covering number of network parts, e.g., via Maurey's sparsification lemma, or via parameter counting. The corresponding whole-network bounds typically scale with the product of each layer's Lipschitz constant or local empirical estimates thereof [37]. Irrespective of the particular proof strategy, most formal arguments only hold for neural networks constructed from function composition, i.e., maps of the form

$$x \mapsto \sigma_L(A_L \sigma_{L-1}(A_{L-1} \dots \sigma_1(A_1 x) \dots)) \ , \tag{1}$$

where $\sigma_i : \mathbb{R}^{d_{i-1}} \to \mathbb{R}^{d_i}$ are nonlinearities and A_i are the weight matrices of the i-th linear map. However, modern architectures often rely on operations specifically tailored to the data, such as convolutions, and typically incorporate skip connections as in residual networks [17], rendering many results inapplicable or suboptimal for these models. In this work, we handle convolutions and skip connections, thus increasing the applicability and utility of such bounds.

For example, while residual networks have been studied extensively, theory mostly focuses on expressivity or optimization aspects [3, 15, 24, 38]. Yun et al. [38] provide a Rademacher complexity bound via layer-peeling for fully-connected layers and element-wise activations. He et al. [16] establish a generalization bound for residual networks via covering number arguments, resting upon earlier work by Bartlett et al. [3] for linear maps. However, when directly applied to convolutions, both bounds scale unfavorably w.r.t. the spatial input size (see Section 3.2). Other works provide generalization guarantees *specifically* tailored to convolutional networks, cf. [13, 22, 25, 27], and, although such bounds scale benignly with input size, they only apply for models as in Eq. (1).

An initial numerical comparison. Bounds on the empirical Rademacher complexity differ in their dependence on various quantities, such as matrix $l_{p,q}$ group norms, Lipschitz constants, or the number of parameters. This makes a precise formal comparison challenging as, depending on the setting, different bounds might be preferable. To provide some intuition about magnitude differences, we evaluated several existing bounds (including ours from Section 3.4) on two convolutional (ReLU) networks with 6 and 11 layers, see Fig. 1 and Appendix A.2 for details.

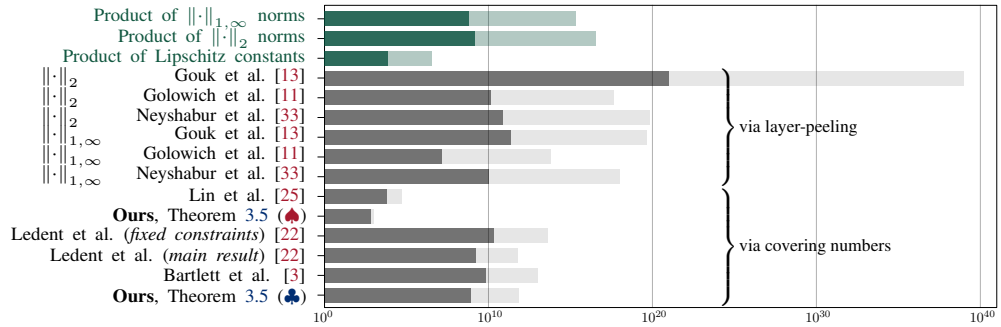

Figure 1: Empirical Rademacher complexity bounds (grouped by proof strategy; lower is better), for a 6- (■■) and an 11-layer (■■) convolutional network, trained on CIFAR10. Bounds are listed in Appendix A.2 and quantities that typically appear in these bounds are highlighted in **green** (top part of figure) for reference.

3 Rademacher Complexity Analysis

To derive bounds on the empirical Rademacher complexity, we follow the (margin-based) multiclass learning formalization and the flexible proof strategy of Bartlett et al. [3]. Section 3.2 introduces novel single-layer covering number bounds for convolutions; Section 3.3 modularizes and extends the single- to multi-layer covering step to account for architectures such as residual networks (Section 3.4).

3.1 Preliminaries

In a k-class classification task, we are given n instance/label pairs $S = ((x_1, y_1), \ldots, (x_n, y_n))$, drawn iid from a probability measure $\mathcal{D}$ on $\mathbb{R}^d \times \{1, \ldots, k\}$. For a neural network f in a hypothesis class $\mathcal{F} \subset \{f \colon \mathbb{R}^d \to \mathbb{R}^k\}$, a class label for input x is obtained by taking the argmax over the components of $f(x) \in \mathbb{R}^k$. The *margin operator* $\mathcal{M} \colon \mathbb{R}^k \times \{1, \ldots, k\} \to \mathbb{R}, (v, y) \mapsto v_y - \max_{i \neq y} v_i$ leads to the *ramp loss* ℓ_γ and the *empirical ramp risk* $\hat{R}_\gamma$, at margin $\gamma > 0$, defined as

$$\ell_\gamma : \mathbb{R} \to \mathbb{R}^+, r \mapsto (1 + {}^r\!/_\gamma)\mathbb{1}_{r \in [-\gamma, 0]} + \mathbb{1}_{r > 0} \quad \text{and} \quad \hat{R}_\gamma(f) = \frac{1}{n} \sum_{i=1}^{n} \ell_\gamma(-\mathcal{M}(f(x_i), y_i)) \ . \tag{2}$$

To derive generalization bounds via classical Rademacher complexity analysis [29], without having to resort to vector-contraction inequalities [28], we consider the hypothesis class

$$\mathcal{F}_\gamma = \{(x, y) \mapsto \ell_\gamma(-\mathcal{M}(f(x), y)) : f \in \mathcal{F}\} \ . \tag{3}$$

Then, defining the *empirical Rademacher complexity* of any class $\mathcal{H}$ of real-valued functions as

$$\hat{\mathfrak{R}}_S(\mathcal{H}) = \mathop{\mathbb{E}}_{\sigma} \left[\sup_{h \in \mathcal{H}} \frac{1}{n} \sum_{i=1}^{n} \sigma_i h(x_i, y_i) \right] \ , \tag{4}$$

with iid Rademacher variables $\sigma = (\sigma_1, \ldots, \sigma_n)$ from a uniform distribution on $\{\pm 1\}$, facilitates to study $\mathcal{F}_\gamma$ via Eq. (4). The following lemma [3, Lemma 3.1] establishes the link to a margin-based multiclass generalization bound for any $f \in \mathcal{F}$.

Lemma 3.1. *Given a hypothesis class $\mathcal{F}$ of functions $f \colon \mathbb{R}^d \to \mathbb{R}^k$ and classification margin $\gamma > 0$, then, with probability of at least $1 - \delta$ over the choice of $S \sim \mathcal{D}^n$, for any $f \in \mathcal{F}$, it holds that*

$$\mathbb{P}[\arg\max_{i \in \{1, \ldots, k\}} f(x)_i \neq y] \leq \hat{R}_\gamma(f) + 2\hat{\mathfrak{R}}_S(\mathcal{F}_\gamma) + 3\sqrt{\frac{\log\left(\frac{2}{\delta}\right)}{2n}} \ . \tag{5}$$

To obtain a computable expression for the right-hand side of Lemma 3.1, we seek a bound on $\hat{\mathfrak{R}}_S(\mathcal{F}_\gamma)$ tied to some measurable quantities of the network realizing $\mathcal{F}$. For our purposes, the relationship of $\hat{\mathfrak{R}}_S(\mathcal{F}_\gamma)$ and the (log) covering number of $\mathcal{F}_\gamma$ turns out to be a flexible approach. In general, given a normed space $(\mathcal{G}, \|\cdot\|_{\mathcal{G}})$, the *covering number* $\mathcal{N}(\mathcal{G}, \epsilon, \|\cdot\|_{\mathcal{G}})$ is the cardinality of the smallest ϵ-cover of $\mathcal{G}$, i.e., of the smallest subset $\mathcal{U} \subset \mathcal{G}$ such that, for any $g \in \mathcal{G}$, there exists $u \in \mathcal{U}$ with

$\|g-u\|_{\mathcal{G}} \leq \epsilon$. In our setting, $\mathcal{G}$ is a class of functions $g : (\mathcal{X}, \|\cdot\|_{\mathcal{X}}) \to (\mathcal{Y}, \|\cdot\|_{\mathcal{Y}})$ between normed spaces $\mathcal{X}$ and $\mathcal{Y}$. Given data $X = (x_1, \ldots, x_n) \in \mathcal{X}^n$, we define a *data-dependent norm* on $\mathcal{G}|_X$ as

$$\|g\|_X = \sqrt{\sum_i \|g(x_i)\|_{\mathcal{Y}}^2} \ , \tag{6}$$

i.e., as the l_2 norm on the restriction of $\mathcal{G}$ to X. Specifically, we seek to bound $\log \mathcal{N}(\mathcal{F}_\gamma, \epsilon, \|\cdot\|_S)$, as this facilitates, via Dudley's entropy integral, to upper bound the empirical Rademacher complexity of $\mathcal{F}_\gamma$; see Eq. (4). Typically, such covering numbers depend on the norm of the data itself, i.e., $\|X\| = \sqrt{\|x_1\|_X^2 + \cdots \|x_n\|_X^2}$.

3.2 Covering number bounds for convolutions

We consider 2D convolutions, acting on images with c_{in} channels of width w and height h, i.e., $x \in \mathbb{R}^{c_{in} \times h \times w}$. For readability only, we discuss convolutions of stride 1 and input-size preserving padding; this is not an assumption required for Theorem 3.2. Formally, a convolutional layer is a linear (as we omit bias parameters) map $\phi_K \colon \mathbb{R}^{c_{in} \times h \times w} \to \mathbb{R}^{c_{out} \times h \times w}$, where c_{in} and c_{out} denote the number of input/output channels. The map is parametrized by a tensor $K \in \mathbb{R}^{c_{out} \times c_{in} \times k_h \times k_w}$ of spatial extension/kernel size (k_h, k_w). Since convolutions are linear maps, they correspond to the action of matrices on the (reshaped) inputs and one could invoke existing covering number bounds. However, this is suboptimal, as any structure specific to convolutions is lost. In particular, norm-based generalization bounds agnostic to this structure incur unfavorable scaling behavior w.r.t. the dimensionality of the input. To be more specific, the weight tensor K of a convolutional layer does not directly specify the corresponding matrix; instead, it parametrizes c_{out} filters, i.e., local linear maps, which are applied to the $(c_{in} \times k_h \times k_w)$-sized pixel neighborhoods of the input. Hence, the matrix corresponding to the *global* linear map, denoted by M_K, consists of many copies of the elements of this tensor, one for each of the hw patches the convolution is applied to. Thus, the p-norm of the matrix M_K is $\|M_K\|_p = (hw)^{1/p} \|K\|_p$ (see Appendix A.1). We mitigate this scaling issue by tying the covering number of a convolutional layer to a variant of the (2,1) group norm on the *tensor* K. We define this norm as the sum over the l_2 norms taken along the *input channels* of K, i.e.,

$$\|K\|_{2,1} = \sum_{ikl} \|K_{i \cdot kl}\|_2 = \sum_{ikl} \sqrt{\sum_j K_{ijkl}^2} \ . \tag{7}$$

For the special case of inputs of size $(h, w) = (1, 1)$ and kernel size $(k_h, k_w) = (1, 1)$, convolution is just matrix multiplication along the channels. In this case, $M_K = K_{\cdot\cdot 11}$ and our norm from Eq. (7) agrees with the standard $(2, 1)$ group norm on $M_K^\top$, i.e., $\|K\|_{2,1} = \left\|M_K^\top\right\|_{2,1}$. Theorem 3.2 establishes two covering number bounds for convolutions.

Theorem 3.2. *Let $b > 0$ and $\mathcal{F} = \{\phi_K |\ K \in \mathbb{R}^{c_{out} \times c_{in} \times k_h \times k_w}, \|K\|_{2,1} \leq b\}$ denote the class of 2D convolutions with c_{in} input channels, c_{out} output channels, kernel size $k_h \times k_w$, parametrized by tensors K with W parameters. Then, for any $X \in \mathbb{R}^{n \times c_{in} \times h \times w}$ and covering radius $\epsilon > 0$,*

$$\log \mathcal{N}(\mathcal{F}, \epsilon, \|\cdot\|_X) \leq \begin{cases} \left\lceil \dfrac{\|X\|^2 b^2}{\epsilon^2} \right\rceil \log(2W) & \text{(8a)} \\ 2W \log\left(1 + \left\lceil \dfrac{\|X\|^2 b^2}{\epsilon^2} \right\rceil\right) \ . & \text{(8b)} \end{cases}$$

Eq. (8a) is analogous to the single-layer bound of Bartlett et al. [3, Lemma 3.2] for fully-connected layers, but replaces the $(2, 1)$ group norm constraint on matrices M with a constraint on tensors K. This is tighter than invoking [3, Lemma 3.2] directly on $M_K \in \mathbb{R}^{c_{in} hw \times c_{out} hw}$, as K has only $c_{out} c_{in} k_h k_w$ parameters and $\left\|M_K^\top\right\|_{2,1} \geq {}^{hw}/_{\sqrt{k_h k_w}} \|K\|_{2,1}$, see Appendix A.1. A meaningful comparison between the two bounds in Theorem 3.2 is more nuanced, though, as preferring one over the other depends on the ratio between the number of parameters W and $\|X\|^2 b^2/\epsilon^2$. The latter, in turn, requires to consider all covering radii ϵ. Hence, we defer this discussion to Section 3.4, where differences more clearly manifest in the overall empirical Rademacher complexity bounds.

Proof sketch. The statement of Theorem 3.2 follows from an application of Maurey's sparsification lemma, which guarantees the existence of an ϵ-cover of $\mathcal{F}$ (of known cardinality) if there is a finite subset $\{V_1, \ldots, V_d\} \subset \mathcal{F}$ s.t. every $f \in \mathcal{F}$ is a convex combination of the V_i. We show that one can find such a finite subset of cardinality $d = 2c_{in}c_{out}k_h k_w = 2W$. The cardinality of the cover is then determined by a *combinatorial quantity* which additionally depends on $\|X\|$ and the norm constraint b. Bounding this quantity, i.e., a binomial coefficient, in two different ways, establishes the bounds.

Relation to prior work. Closely related is recent work by Ledent et al. [22] who derive l_∞ covering number bounds for convolutional layers based on a classic result by Zhang [41]. Similar to Eq. (8a), their bound depends on the square of a weight norm directly on the tensor K, the square of a data norm, as well as a logarithmic term. The data norm is the maximal l_2 norm of a single patch. Compared to our result, this implicitly removes a factor of the spatial dimension hw, which however reenters via the Lipschitz constant when transitioning to multi-layer bounds (as the spatial dimension of the output after subsequent pooling). Overall, the quadratic terms across both results scale similarly (with our data norm being less sensitive to outliers), but we improve on Ledent et al. [22] in the logarithmic term. By contrast, the use of l_∞ covers in [22] yields whole-network bounds with improved dependency on the number of classes; see Appendix A.5 for an in-depth comparison. In other related work, Lin et al. [25] derive an l_2 covering number bound for convolutional layers similar to Eq. (8b), which depends linearly on the number of parameters and logarithmically on norms. Their proof handles the local linear maps, determined by the weight tensor, via standard covering numbers for balls in $\mathbb{R}^d$. This incurs an additional factor inside the logarithm that corresponds to the number of these applications, i.e., the spatial dimension of the output. Importantly, non-convolution specific approaches can equally mitigate undesirable scaling issues, e.g., by utilizing $(1, \infty)$ group norms on the matrices representing the linear map [11, 13, 33]; as differences primarily manifest in the resulting bounds on the Rademacher complexity, we refer to our discussion in Section 3.4.

3.3 Covering number bounds for composition & addition

As many neural networks are built from composition and summation of layers, we also study covering numbers under these operations. The key building blocks are the following, easy to verify, inequalities.

Lemma 3.3. *Let $\mathcal{F}_1, \mathcal{F}_2$ be classes of functions on normed spaces $(\mathcal{X}, \|\cdot\|_{\mathcal{X}}) \to (\mathcal{Y}, \|\cdot\|_{\mathcal{Y}})$ and let $\mathcal{G}$ be a class of c-Lipschitz functions $(\mathcal{Y}, \|\cdot\|_{\mathcal{Y}}) \to (\mathcal{Z}, \|\cdot\|_{\mathcal{Z}})$. Then, for any $X \in \mathcal{X}^n$ and any $\epsilon_{\mathcal{F}_1}$, $\epsilon_{\mathcal{F}_2}, \epsilon_{\mathcal{G}} > 0$, it holds that*

$$\mathcal{N}(\{f_1 + f_2 \mid f_1 \in \mathcal{F}_1, f_2 \in \mathcal{F}_2\}, \epsilon_{\mathcal{F}_1} + \epsilon_{\mathcal{F}_2}, \|\cdot\|_X) \leq \mathcal{N}(\mathcal{F}_1, \epsilon_{\mathcal{F}_1}, \|\cdot\|_X)\mathcal{N}(\mathcal{F}_2, \epsilon_{\mathcal{F}_2}, \|\cdot\|_X) \quad (9)$$

and

$$\mathcal{N}(\{g \circ f \mid g \in \mathcal{G}, f \in \mathcal{F}_2\}, \epsilon_{\mathcal{G}} + c\epsilon_{\mathcal{F}_2}, \|\cdot\|_X) \leq \mathcal{N}(\mathcal{F}_2, \epsilon_{\mathcal{F}_2}, \|\cdot\|_X) \sup_{f \in \mathcal{F}_2} \mathcal{N}(\mathcal{G}, \epsilon_{\mathcal{G}}, \|\cdot\|_{f(X)}) \ . \quad (10)$$

To establish these inequalities, one chooses minimal covers of the original function spaces and links their elements via the considered operation, i.e., addition or composition. The resulting functions correspond to tuples of elements of the original covers. Hence, the right-hand side of the inequalities is a product of covering numbers. The crucial step is to determine a, preferably small, radius ϵ such that these functions form an ϵ-cover. In Lemma 3.3, this is achieved via standard properties of norms. Notably, iterative application of Lemma 3.3 allows bounding the covering numbers of any function class built from compositions and additions of simpler classes.

In Appendix C.3, we apply Lemma 3.3 on two examples, i.e., (1) $f \in \mathcal{F} = \{f_L \circ \cdots \circ f_1\}$ and (2) $h \in \mathcal{H} = \{g + h_L \circ \cdots \circ h_1\}$. Instantiating the first example for $f_i = \sigma_i \circ \phi_i$, with σ_i fixed and ϕ_i from a family of linear maps, yields covering number bounds for networks as in Eq. (1). As the second example corresponds to residual blocks (with g possibly the identity map), the combination of (1) and (2) yields covering number bounds for residual networks; see Example C.3.

Overall, this strategy not only allows to derive covering number bounds for a broad range of architectures, but also facilitates integrating linkings between function spaces in a modular way. For instance, Lemma C.14 provides a variant of Lemma 3.3 for *concatenation*, used in DenseNets [19].

Relation to prior work. He et al. [16] investigate covering number bounds for function spaces as considered above. They present covering number bounds for residual networks and show that the covering number $\mathcal{N}(\mathcal{F}, \epsilon, \|\cdot\|_X)$ of such models with layers $\mathcal{F}_\alpha$ is bounded by the product

$\prod_\alpha \sup_{\phi\in\mathcal{G}_\alpha} \mathcal{N}(\mathcal{F}_\alpha, \epsilon_{\mathcal{F}_\alpha}, \|\cdot\|_{\phi(X)})$ for appropriately defined function spaces $\mathcal{G}_\alpha$. Yet, the dependency of the whole network covering radius ϵ on the single-layer covering radii $\epsilon_{\mathcal{F}_\alpha}$ is *only* derived for a very specific residual network. Our addition to the theory is a more modular and structured way of approaching the problem, which we believe to be valuable on its own.

3.4 Covering number bounds for residual networks

We next state our whole-network covering number bounds for residual networks and the corresponding bounds on the empirical Rademacher complexity. Accompanying generalization guarantees (obtained via Lemma 3.1) are given in Appendix C.5. The results of this section hold for a hypothesis class $\mathcal{F}$ of networks implementing functions of the form

$$f = \sigma_L \circ f_L \circ \cdots \circ \sigma_1 \circ f_1 \quad \text{with} \quad f_i(x) = g_i(x) + (\sigma_{iL_i} \circ h_{iL_i} \circ \cdots \circ \sigma_{i1} \circ h_{i1})(x) \ , \tag{11}$$

i.e., a composition of L residual blocks. Here, the nonlinearities σ_i and σ_{ij} are fixed and ρ_i-, resp., ρ_{ij}-Lipschitz continuous with $\sigma_i(0) = 0$ and $\sigma_{ij}(0) = 0$. We further fixate the shortcuts to maps with $g_i(0) = 0$. The map h_{ij} identifies the j-th layer in the i-th residual block with Lipschitz constraints s_{ij} and distance constraints b_{ij} (w.r.t. reference weights M_{ij}). In particular, if h_{ij} is *fully-connected*,

$$h_{ij} \in \left\{ \phi : x \mapsto A_{ij}x \,\middle|\, \mathrm{Lip}(\phi) \leq s_{ij},\ \left\|A_{ij}^\top - M_{ij}^\top\right\|_{2,1} \leq b_{ij} \right\} \ , \tag{12}$$

and, in case h_{ij} is *convolutional*,

$$h_{ij} \in \left\{ \phi_{K_{ij}} \,\middle|\, \mathrm{Lip}(\phi_{K_{ij}}) \leq s_{ij},\ \|K_{ij} - M_{ij}\|_{2,1} \leq b_{ij} \right\} . \tag{13}$$

In terms of notation, $s_i = \mathrm{Lip}(g_i) + \prod_{j=1}^{L_i} \rho_{ij} s_{ij}$ further denotes the upper bound on the Lipschitz constant of the i-th residual block f_i. The Lipschitz constants are w.r.t. Euclidean norms; for a fully-connected layer this coincides with the spectral norm of the weight matrix.

The following covering number bounds in Theorems 3.4 and 3.5 depend on three types of quantities: (1) the total number of layers $\bar{L} = \sum_i L_i$, (2) the numbers W_{ij} of parameters of the j-th layer in the i-th residual block, their maximum $W = \max_{ij} W_{ij}$, and (3) terms C_{ij} that quantify the part of a layer's capacity attributed to weight and data norms. With respect to the latter, we define

$$C_{ij}(X) = 2\,\frac{\|X\|}{\sqrt{n}} \left(\prod_{l=1}^{L} s_l \rho_l\right) \frac{\prod_{k=1}^{L_i} \rho_{ik} s_{ik}}{s_i} \frac{b_{ij}}{s_{ij}} \tag{14}$$

and write $C_{ij} = C_{ij}(X)$ for brevity. Importantly, $\|X\| \leq \sqrt{n} \max_i \|x_i\|$ and so the C_{ij} can be bounded independently of the sample size. With this notation, the residual network covering number bounds are as follows.

Theorem 3.4. *The covering number of the class of residual networks $\mathcal{F}$ as specified above, satisfies*

$$\log \mathcal{N}(\mathcal{F}, \epsilon, \|\cdot\|_X) \leq \begin{cases} \left(\log(2W) \sum_{i=1}^{L} \sum_{j=1}^{L_i} \left\lceil C_{ij}^{2/3} \right\rceil\right)^3 \frac{n}{\epsilon^2} & \text{(15a)} \\ \sum_{i=1}^{L} \sum_{j=1}^{L_i} 2W_{ij} \log\left(1 + \left\lceil \bar{L}^2 C_{ij}^2 \right\rceil \frac{n}{\epsilon^2}\right) \ . & \text{(15b)} \end{cases}$$

In combination with Dudley's entropy integral, Theorem 3.4 implies the empirical Rademacher complexity bounds in Theorem 3.5 below. These bounds equally hold for non-residual networks as in Eq. (1), i.e., the special case of setting the shortcuts g_i to the zero map (with $L = 1$ block).

Theorem 3.5. *Let $\gamma > 0$ and define $\tilde{C}_{ij} = 2C_{ij}/\gamma$. Further, let $H_{n-1} = \sum_{m=1}^{n-1} 1/m = \mathcal{O}(\log(n))$ denote the $(n-1)$-th harmonic number. Then, the empirical Rademacher complexity of $\mathcal{F}_\gamma$ satisfies*

$$\hat{\mathfrak{R}}_S(\mathcal{F}_\gamma) \leq \frac{4}{n} + \frac{12 H_{n-1}}{\sqrt{n}} \sqrt{\log(2W)} \left(\sum_{i=1}^{L} \sum_{i=j}^{L_i} \left\lceil \tilde{C}_{ij}^{2/3} \right\rceil\right)^{3/2} \tag{♣}$$

and

$$\hat{\mathfrak{R}}_S(\mathcal{F}_\gamma) \leq \frac{12}{\sqrt{n}} \sqrt{\sum_{i=1}^{L} \sum_{i=j}^{L_i} 2W_{ij} \log\left(1 + \left\lceil \bar{L}^2 \tilde{C}_{ij}^2 \right\rceil\right) + \psi\left(\left\lceil \bar{L} \tilde{C}_{ij} \right\rceil\right)} \ , \qquad (\spadesuit)$$

where ψ is a monotonically increasing function, satisfying $\psi(0) = 0$ and $\forall x : \psi(x) < 2.7$.

Interpretation. To facilitate a clean comparison between the bounds in Theorem 3.5, we disregard the ceiling function and apply Jensen's inequality to the first bound ($\clubsuit$), yielding

$$\hat{\mathfrak{R}}_S(\mathcal{F}_\gamma) \leq \frac{4}{n} + 12 \sqrt{\log(2W) \sum\nolimits_{ij} \bar{L}^2 \tilde{C}_{ij}^2} \ \frac{H_{n-1}}{\sqrt{n}} \ . \qquad (16)$$

Denoting $\tilde{C} = \max_{ij} \tilde{C}_{ij}$, Eq. (16) reveals that the bounds essentially differ only in that ($\clubsuit$) depends on $(\log(2W)\bar{L}^2\tilde{C}^2)^{1/2}$ and ($\spadesuit$) depends on $(2W \log(1 + \bar{L}^2\tilde{C}^2))^{1/2}$. Thus, the question of *which one is tighter*, hinges on the ratio of $2W$ and $\bar{L}^2\tilde{C}^2$, i.e., a tradeoff between the number of parameters and the weight norms. As $\tilde{C}$ depends exponentially on the network depth via the product of Lipschitz constants, this can also be interpreted as a tradeoff between layer width and network depth. Notably, Rademacher complexity bounds not for $\mathcal{F}_\gamma$, but for networks composed with a Lipschitz augmented loss function [37], as well as inferred generalization bounds, do not suffer from the exponential depth dependency via the product of Lipschitz constants. Instead, such bounds depend on data dependent empirical estimates thereof, which are typically much smaller. Yet, we believe that in our setting bounds on $\hat{\mathfrak{R}}_S(\mathcal{F}_\gamma)$ are preferable, as the reduced data dependency allows for cleaner comparison of excess capacity across datasets. As we see in Fig. 1, for simple, unconstrained networks, our second bound ($\spadesuit$) is much tighter. However, due to the logarithmic dependency on $\tilde{C}$, it is less affected by constraining the distance to initialization b and the Lipschitz constants s. In Section 4, we show that this effect causes ($\clubsuit$) to be a more faithful measure of excess capacity.

Relation to prior work. Prior works [11, 13, 33] that tie generalization to $(1, \infty)$ group norms of matrices of fully-connected layers are equally applicable to convolutional networks without unfavorable scaling w.r.t. input size. In particular, for $(1, \infty)$ group norms of M_K, we have $\|M_K\|_{1,\infty} = \max_o \|K_{o\cdots}\|_1$, i.e., the maximum 1-norm over each (input channel, width, height) slice of K. Yet, due to the layer-peeling strategy common to these works, the bounds scale with the product of matrix group norms *vs.* the product of Lipschitz constants (as in the $\tilde{C}_{ij}$ in Theorem 3.5) for covering number based strategies. While one can construct settings where the product of $(1, \infty)$ group norms is smaller than the product of Lipschitz constants, this is typically not observed empirically, cf. Fig. 1. Alternatively, Long & Sedghi [27] derive a generalization bound which does not depend on p-norms or group norms, but only on the distance to initialization with respect to the spectral norm. Notably, an intermediate result in this reference yields a generalization bound of similar form as ($\spadesuit$), scaling with the logarithm of the product of Lipschitz constants and with the square root of the number of parameters, see Appendix A.4. The distance to initialization then enters the main result [27, Theorem 3.1] at the cost of a Lipschitz constraint on the initialization. We argue that ($\spadesuit$) incorporates the distance to initialization more naturally, as it comes without constraints on the initialization itself. Further, it holds for any sample size n and comes without unknown constants.

4 Empirical Evaluation

To assess excess capacity, for each network trained via a standard protocol on some dataset, we seek a hypothesis class that contains a network of the same architecture with comparable testing error, but *smaller* capacity. Controlling capacity via the bounds in Theorem 3.5 requires *simultaneously* constraining (1) the Lipschitz constants per layer and (2) the (2,1) group norm distance of each layer's weight to its initialization. We discuss how to enforce the constraints next, then fix a convolutional residual network and train on datasets of increasing difficulty while varying the constraint strengths.

Capacity reduction. Controlling hypothesis class capacity necessitates ensuring that optimization yields a network parametrization that satisfies the desired constraints. To this end, we implement a variant of *projected* stochastic gradient descent (SGD) where, after a certain number of update steps, we project onto the intersection of the corresponding constraint sets $\mathcal{C}_1$ and $\mathcal{C}_2$. For convolutional layers, parametrized by tensors K, these are the convex sets $\mathcal{C}_1 = \{K : \|K - K^0\|_{2,1} \leq b\}$ and $\mathcal{C}_2 = \{K : \mathrm{Lip}(\phi_K) \leq s\}$. Hence, jointly satisfying the constraints is a convex feasibility problem of finding a point in $\mathcal{C} = \mathcal{C}_1 \cap \mathcal{C}_2$. To ensure $\mathcal{C} \neq \emptyset$, we initially (prior to optimization) scale each layer's weight K so that $\mathrm{Lip}(\phi_{K^0}) = s$. This starting point (per layer) resides in $\mathcal{C}$ by construction.

To project onto $\mathcal{C}$, we rely on alternating *orthogonal projections* which map K to a tensor in $\mathcal{C}_1$, resp. $\mathcal{C}_2$ with minimal distance to K. Repeated application of these projections converges to a point in the intersection $\mathcal{C}$ [5]. To implement the orthogonal projection onto $\mathcal{C}_1$, we rely on work by Liu et al. [26]; recent work by Sedghi et al. [34] allows to compute the orthogonal projection onto $\mathcal{C}_2$ under certain architectural prerequisites. As a consequence, we need to use convolutions of stride 1 (though our bounds equally hold for strides > 1) and reduce spatial dimensionality only via max-pooling. Further, we use circular padding and kernel sizes not larger than the input dimensionalities. For details on the projection algorithm and a comparison to alternating *radial* projections, see Appendix B.2.

Architecture. We use a slightly modified (pre-activation) ResNet18 [18]. Modifications include: (1) the *removal of batch normalization* and biases; (2) skip connections for residual blocks where the number of channels doubles and spatial dimensionality is halved are implemented via a fixed map. Each half of the resulting channels is obtained via 2×2 spatial max-pooling (shifted by one pixel). This map has Lipschitz constant $\sqrt{2}$ and is similar to the shortcut variant (A) of [17]; finally, (3) we use a *fixed* linear classifier with weights set to the vertices of a $k-1$ unit simplex in the output space of the network. Fixing the classifier is motivated by [15] and the simplex configuration is inspired by [14, 42] who show that this configuration corresponds to the geometric weight arrangement one would obtain at minimal cross-entropy loss. By construction, this classifier has Lipschitz constant $\sqrt{k/k-1}$. Modification (1) is required as the running mean and variance of batch normalization layers prohibit to treat them as either fixed activation functions or linear layers. Thus, including batch normalization would leave the setting of Eq. (11) in Section 3, which specifies the function class for which (♣) and (♠) are faithful capacity measures. Notably, modifications (2) and (3) *do not* harm performance, with empirical testing errors on a par with a standard ResNet18 w/o batch normalization.

Datasets & Training. We test on three benchmark datasets: CIFAR10/100 [21], and Tiny-ImageNet-200, listed in order of increasing task difficulty. We minimize the cross-entropy loss using SGD with momentum (0.9) for 200 epochs with batch size 256 and follow a CIFAR-typical stepwise learning rate schedule, decaying the initial learning rate (of 3e-3) by a factor of $5\times$ at epochs 60, 120 & 160 [9]. *No data augmentation is used*. When projecting onto the constraint sets, we found one alternating projection step every 15th SGD update to be sufficient in order to remain close to $\mathcal{C}$. At the final SGD update, we increase the alternating projection steps to 15 which, in combination with a small learning rate, is empirically enough for convergence of the iterated projection. This ensures that a trained model is within the capacity-constrained class.

4.1 Results

First, we assess the capacity-driving quantities in our bounds for models trained *without* constraints. Table 1 (top) lists a comparison across datasets, along with the capacity measures (♣, ♠), the training/testing error, and the empirical generalization gap. The only explicit regularization is due to small (1e-4) weight decay. In accordance with our motivating figure in Section 1, we observe an overall *increase* in both capacity-driving quantities as a function of task difficulty.

To assess excess capacity in the context of questions (**Q1**) and (**Q2**), we first identify, per dataset, the *most restrictive* constraint combination where the testing error[1] is *as close as possible* to the unconstrained regime. We refer to this setting as the *operating point* for the constraints, characterizing the function class $\mathcal{F}$ that serves as a reference to measure excess capacity. The operating points per dataset, as well as the corresponding results are listed in Table 1 (bottom).

With respect to (**Q1**) we find that networks can indeed maintain, or even improve, testing error at substantially lowered capacity (see performance comparison in Fig. 5 relative to ◆). Further, as capacity of the constrained models surprisingly remains in the *same order* of magnitude across tasks of varying difficulty, this equally supports an affirmative answer to question (**Q2**). In fact, task difficulty primarily manifests as *excess* capacity, as can be seen from comparing the top *vs.* bottom part of Table 1 with respect to column (♣). Another manifestation of task difficulty is evident from the more detailed analysis in Fig. 5 (bottom), where we see that tightening both constraints beyond the identified operating point leads to a more rapid deterioration of the testing error as task difficulty increases. Interestingly, at the operating point, the constrained models do not only share similar capacity across datasets, but also similar *generalization gaps*, primarily due to leaving the ubiquitous zero-training-error regime. The latter is particularly noteworthy, as strong regularization

[1] We are primarily interested in *what is feasible* in terms of tolerable capacity reduction. Hence, leveraging the testing split of each dataset for this purpose is legitimate from this exploratory perspective.

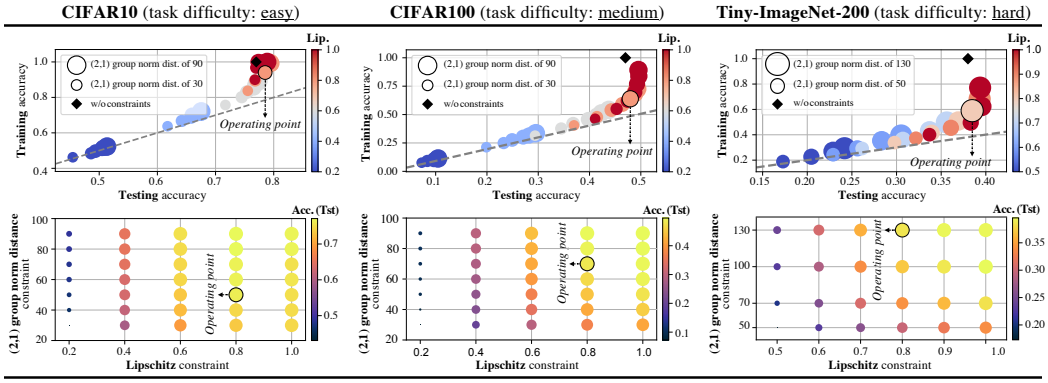

Figure 2: Fine-grained analysis of training/testing accuracy in relation to the Lipschitz constraint and the (2,1) group norm distance to initialization constraint. We see that testing accuracy can be retained (relative to ◆) for a range of fairly restrictive constraints (*top row*), compared to the unconstrained regime (cf. **Lip./Dist.** columns in the top part of Table 1). However, this range noticeably narrows with increasing task difficulty (*bottom row*) .

Table 1: Assessment of the capacity-driving quantities. We list the *median* over the Lipschitz constants (**Lip.**) and the (2,1) group norm distances (**Dist.**) across all layers. **Err.** denotes the training/testing error, **Capacity** denotes the measures (♣, ♠) from Theorem 3.5 and **Gap** the empirical generalization gap. The top part lists results in the unconstrained regime (see ◆ in Fig. 5), the bottom part lists results at the operating point of the most restrictive constraint combination where the testing error is on a par with the unconstrained case.

	Lip.	**Dist.**	**Err. (Tst)**	**Err. (Trn)**	**Capacity (♣, ♠)**	**Gap**
CIFAR10	1.63	60.3	0.24	0.0	$1.0{\cdot}10^{10}$ / $8.8{\cdot}10^{2}$	0.24
CIFAR100	2.17	129.1	0.54	0.0	$1.7{\cdot}10^{11}$ / $9.3{\cdot}10^{2}$	0.54
Tiny-ImageNet-200	3.05	287.3	0.62	0.0	$4.5{\cdot}10^{13}$ / $1.1{\cdot}10^{3}$	0.62
CIFAR10	0.80	50.0	0.21	0.06	$1.8{\cdot}10^{8}$ / $7.8{\cdot}10^{2}$	0.15
CIFAR100	0.80	70.0	0.52	0.36	$2.6{\cdot}10^{8}$ / $7.9{\cdot}10^{2}$	0.16
Tiny-ImageNet-200	0.80	130.0	0.62	0.41	$8.9{\cdot}10^{8}$ / $8.9{\cdot}10^{2}$	0.21

(e.g., via weight decay) can equally enforce this behavior, but typically at the cost of a large increase in testing error (which we do not observe). Finally, the parameter-counting variant of our bound (see Table 1, column ♠) is, by construction, much less affected by the constraints and apparently fails to capture the observations above. This highlights the relevance of tying capacity to weight norms and underscores their utility in our context.

5 Discussion

Studying the capacity of neural networks hinges crucially on the measure that is used to quantify it. In our case, capacity rests upon two bounds on the empirical Rademacher complexity, both depending on weight norms and the number of parameters, but to different extents. Hence, exerting control over the *weight norms* manifests in different ways: in case of the more weight norm dependent capacity measure, our results show substantial task-dependent *excess capacity*, while, when relying more on parameter counting, this effect is less pronounced. Although the latter measure yields tighter bounds, its utility in terms of explaining the observed empirical behavior is limited: in fact, capacity tied to weight norms not only better correlates with observed generalization gaps (both with and without constraints), but the amount of tolerable capacity reduction also reflects the smaller generalization gaps in the constrained regime. Note that our results rest upon carefully implementing constraint enforcement during optimization. Hence, numerical schemes to better account for this setting might potentially reveal an even more pronounced excess capacity effect.

In summary, our experiments, guided by the theoretical bounds, strongly suggest a notion of *compressibility* of networks with respect to weight norms, and this compressibility only moderately reduces with task difficulty. We believe these observations to be particularly relevant, and foresee them sparking future work along this direction.

References

[1] Sanjeev Arora, Simon Du, Wei Hu, Zhiyuan Li, and Ruosong Wang. Fine-grained analysis of optimization and generalization for overparameterized two-layer neural networks. In *ICML*, 2019.

[2] Sanjeev Arora, Rong Ge, Behnam Neyshabur, and Yi Zhang. Stronger generalization bounds for deep nets via a compression approach. In *ICML*, 2018.

[3] Peter Bartlett, Dylan J Foster, and Matus Telgarsky. Spectrally-normalized margin bounds for neural networks. In *NeurIPS*, 2017.

[4] Peter L Bartlett and Shahar Mendelson. Rademacher and Gaussian complexities: Risk bounds and structural results. *J. Mach. Learn. Res. (JMLR)*, 3(Nov):463–482, 2002.

[5] Heinz H. Bauschke and Jonathan M. Borwein. On projection algorithms for solving convex feasibility problems. *SIAM Review*, 36(3):367–426, 1996.

[6] Cenk Baykal, Lucas Liebenwein, Igor Gilitschenski, Dan Feldman, and Daniela Rus. Data-dependent coresets for compressing neural networks with application to generalization bounds. In *ICLR*, 2019.

[7] Yuan Cao and Quanquan Gu. Generalization bounds of stochastic gradient descent for wide and deep neural networks. In *NeurIPS*, 2019.

[8] Ching-Yao Chuang, Youssef Mroueh, Kristjan Greenwald, Antonio Torralba, and Stefanie Jegelka. Measuring generalization with optimal transport. In *NeurIPS*, 2021.

[9] Terrance DeVries and Graham W. Taylor. Improved regularization of convolutional neural networks with cutout. *arXiv preprint* `arxiv.org/abs/1708.04552`, 2017.

[10] Richard M. Dudley. The sizes of compact subsets of Hilbert space and continuity of gaussian processes. *J. Funct. Anal.*, 1(3):290–330, 1967.

[11] Noah Golowich, Alexander Rakhlin, and Ohad Shamir. Size-independent sample complexity of neural networks. In *COLT*, 2018.

[12] Henry Gouk, Eibe Frank, Bernhard Pfahringer, and Michael J. Cree. Regularisation of neural networks by enforcing lipschitz continuity. *Machine Learning*, 110(2):393–416, Feb 2021.

[13] Henry Gouk, Timothy M. Hospedales, and Massimiliano Pontil. Distance-based regularization of deep networks for fine-tuning. In *ICLR*, 2021.

[14] Florian Graf, Christoph Hofer, Marc Niethammer, and Roland Kwitt. Dissecting supervised contrastive learning. In *ICML*, 2021.

[15] Moritz Hardt and Tengyu Ma. Identity matters in deep learning. In *ICLR*, 2017.

[16] Fengxiang He, Tongliang Liu, and Dacheng Tao. Why resnet works? residuals generalize. *IEEE Trans Neural Netw. Learn. Syst.*, 31(12):5349–5362, 2020.

[17] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016.

[18] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Identity mappings in deep residual networks. In *ECCV*, 2016.

[19] Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q. Weinberger. Densely connected convolutional networks. In *CVPR*, 2017.

[20] Arthur Jacot, Franck Gabriel, and Clement Hongler. Neural tangent kernel: Convergence and generalization in neural networks. In *NeurIPS*, 2018.

[21] Alex Krizhevsky. Learning multiple layers of features from tiny images. Technical report, University of Toronto, 2009.

[22] Antoine Ledent, Waleed Mustafa, Yunwen Lei, and Marius Kloft. Norm-based generalisation bounds for multi-class convolutional neural networks. In *AAAI*, 2021.

[23] Qiyang Li, Saminul Haque, Cem Anil, James Lucas, Roger Grosse, and Jörn-Henrik Jacobsen. Preventing gradient attenuation in lipschitz constrained convolutional networks. *NeurIPS*, 2019.

[24] Hongzhou Lin and Stefanie Jegelka. Resnet with one-neuron hidden layers is a universal approximator. In *NeurIPS*, 2018.

[25] Shan Lin and Jingwei Zhang. Generalization bounds for convolutional neural networks. *arXiv preprint* arxiv.org/abs/1910.01487, 2019.

[26] Jun Liu, Shuiwang Ji, and Jieping Ye. Multi-task feature learning via efficient $l_{2,1}$-norm minimization. In *UAI*, 2009.

[27] Philip Long and Hanie Sedghi. Generalization bounds for deep convolutional neural networks. In *ICLR*, 2020.

[28] Andreas Maurer. A vector-contraction inequality for Rademacher complexities. In *ALT*, 2016.

[29] Mehryar Mohri, Afshin Rostamizadeh, and Ameet Talwalkar. *Foundations of machine learning*. MIT press, 2018.

[30] Vaishnavh Nagarajan and J Zico Kolter. Uniform convergence may be unable to explain generalization in deep learning. In *NeurIPS*, 2019.

[31] Vaishnavh Nagarajan and Zico J. Kolter. Generalization in deep networks: The role of distance from initialization. In *NeurIPS workshop on Deep Learning: Bridging Theory and Practice*, 2017.

[32] Behnam Neyshabur, Srinadh Bhojanapalli, and Nathan Srebro. A pac-bayesian approach to spectrally-normalized margin bounds for neural networks. In *ICLR*, 2018.

[33] Behnam Neyshabur, Ryota Tomioka, and Nathan Srebro. Norm-based capacity control in neural networks. In *COLT*, 2015.

[34] Hanie Sedghi, Vineet Gupta, and Philip M Long. The singular values of convolutional layers. In *ICLR*, 2019.

[35] Daniel Soudry, Elad Hoffer, Mor Shpigel Nacson, and Nati Srebro. The implicit bias of gradient descent on separable data. In *ICLR*, 2018.

[36] Taiji Suzuki, Hiroshi Abe, and Tomoaki Nishimura. Compression based bound for non-compressed network: unified generalization error analysis of large compressible deep neural network. In *ICLR*, 2020.

[37] Colin Wei and Tengyu Ma. Data-dependent sample complexity of deep neural networks via lipschitz augmentation. In *NeurIPS*, 2019.

[38] Chulhee Yun, Suvrit Sra, and Ali Jadbabaie. Are deep ResNets provably better than linear predictors? In *NeurIPS*, 2019.

[39] Chiyuan Zhang, Samy Bengio, Moritz Hardt, Benjamin Recht, and Oriol Vinyals. Understanding deep learning requires rethinking generalization. In *ICLR*, 2017.

[40] Chiyuan Zhang, Samy Bengio, Moritz Hardt, Benjamin Recht, and Oriol Vinyals. Understanding deep learning (still) requires rethinking generalization. *Communications of the ACM*, 64(3):107–115, 2021.

[41] Tong Zhang. Covering number bounds of certain regularized linear function classes. *J. Mach. Learn. Res. (JMLR)*, (2):527–550, March 2002.

[42] Zhihui Zhu, Tianyu Ding, Jinxin Zhou, Xiao Li, Chong You, Jeremias Sulam, and Qing Qu. A geometric analysis of neural collapse with unconstrained features. In *NeurIPS*, 2021.

Checklist

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
   (b) Did you describe the limitations of your work? [Yes] *For all presented theoretical results, we discuss connections, as well as limitations or advantages/disadvantages with respect to prior work. Furthermore, our supplementary material (Appendix) contains substantial extensions of these discussions. Regarding our experimental results, we discuss potential limitations in Section 4 and Section 5.*
   (c) Did you discuss any potential negative societal impacts of your work? [N/A] *All experiments throughout this work are conducted on well-established, publicly available benchmark datasets for image classification. We primarily address theoretical aspects of existing neural network models for which we do not see a potential negative societal impact.*
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [Yes] *All key theoretical results (theorems, lemmas, etc.), including assumptions and definitions are listed in the main part of the paper.*
   (b) Did you include complete proofs of all theoretical results? [Yes] *Proofs for all theoretical results listed in the main part of the paper (theorems, lemmas, etc.) can be found in the supplementary material (Appendix), including any omitted technical details.*
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes] *Code to re-run experiments and reproduce results is included in the supplementary as part of this submission.*
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes] *We include training/architectural details in Section 4 of the main part of the paper. Details for our initial bound comparison in Fig. 1 can be found in the supplementary material (Appendix). Throughout all experiments, we adhere to the common training/testing splits of the three datasets we used, i.e., CIFAR10/100 and Tiny-ImageNet-200.*
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [N/A] *We did run experiments multiple times (using the same network, dataset and training regime), but under varying constraint strengths (rather than different seeds). Fig. 5 visualizes all these runs.*
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes] *Appendix B.3 lists all hardware resources used in our experiments.*
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? [Yes]
   (b) Did you mention the license of the assets? [No] *The datasets we used are publicly available and well-established benchmark datasets commonly used in computer vision research.*
   (c) Did you include any new assets either in the supplemental material or as a URL? [No]
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A] *All datasets in our experiments are used as is.*
   (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
   (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]

(b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]

(c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]