

Size-Independent Sample Complexity of Neural Networks

Noah Golowich
Harvard University

Alexander Rakhlin
University of Pennsylvania

Ohad Shamir
Weizmann Institute of Science
and Microsoft Research

Abstract

We study the sample complexity of learning neural networks, by providing new bounds on their Rademacher complexity assuming norm constraints on the parameter matrix of each layer. Compared to previous work, these complexity bounds have improved dependence on the network depth, and under some additional assumptions, are fully independent of the network size (both depth and width). These results are derived using some novel techniques, which may be of independent interest.

1 Introduction

One of the major challenges involving neural networks is explaining their ability to generalize well, even if they are very large and have the potential to overfit the training data [Neyshabur et al., 2014, Zhang et al., 2016]. Learning theory teaches us that this must be due to some inductive bias, which constrains one to learn networks of specific configurations (either explicitly, e.g., via regularization, or implicitly, via the algorithm used to train them). However, understanding the nature of this inductive bias is still largely an open problem.

A useful starting point is to consider the much more restricted class of linear predictors ($\mathbf{x} \mapsto \mathbf{w}^\top \mathbf{x}$). For this class, we have a very good understanding of how its generalization behavior is dictated by the *norm* of \mathbf{w} . In particular, assuming that $\|\mathbf{w}\| \leq M$ (where $\|\cdot\|$ signifies Euclidean norm), and the distribution is such that $\|\mathbf{x}\| \leq B$ almost surely, it is well-known that the generalization error (w.r.t. Lipschitz losses) given m training examples scales as $\mathcal{O}(MB/\sqrt{m})$, completely independent of the dimension of \mathbf{w} . Thus, it is very natural to ask whether in the more general case of neural networks, one can obtain similar “size-independent” results (independent of the networks’ depth and width), under appropriate norm constraints on the parameters. This is also a natural question, considering that the size of modern neural networks is often much larger than the number of training examples.

Classical results on the sample complexity of neural networks (see Anthony and Bartlett [2009]) do not satisfy this desideratum, and have a strong explicit dependence on the network size. In particular, they can be trivial once the number of parameters exceeds the number of training examples. More recently, there have been several works aiming at improving the sample complexity bounds, assuming various norm constraints on the parameter matrices. For example, Neyshabur et al. [2015] use Rademacher complexity tools to show that if the parameter matrices W_1, \dots, W_d in each of the d layers have Frobenius norms upper-bounded by $M_F(1), \dots, M_F(d)$ respectively, and under suitable assumptions on the activation functions, the generalization error scales (with high probability) as

$$\mathcal{O}\left(\frac{B2^d \prod_{j=1}^d M_F(j)}{\sqrt{m}}\right). \quad (1)$$

Although this bound has no explicit dependence on the network width (that is, the dimensions of W_1, \dots, W_d), it has a very strong, exponential dependence on the network depth d , even if $M_F(j) \leq 1$ for all j . Neyshabur et al. [2015] also showed that this dependence can sometimes be avoided for anti-symmetric activations, but unfortunately this is a non-trivial assumption, which is not satisfied for common activations such as the ReLU. Bartlett et al. [2017] use a covering numbers argument to show a bound scaling as

$$\tilde{\mathcal{O}} \left(\frac{B \left(\prod_{j=1}^d \|W_j\| \right) \left(\sum_{j=1}^d \left(\frac{\|W_j^T\|_{2,1}}{\|W_j\|} \right)^{2/3} \right)^{3/2}}{\sqrt{m}} \right), \quad (2)$$

where $\|W\|$ denotes the spectral norm of W , $\|W^T\|_{2,1} := \sum_l \sqrt{\sum_k W_j(l,k)^2}$ denotes the 1-norm of the 2-norms of the rows of W , and where we ignore factors logarithmic in m and the network width. Unlike Eq. (1), here there is no explicit exponential dependence on the depth. However, there is still a strong and unavoidable polynomial dependence: To see this, note that for any W_j , $\frac{\sum_l \sqrt{\sum_k W_j(l,k)^2}}{\|W_j\|} \geq \frac{\|W_j\|_F}{\|W_j\|} \geq 1$, so the bound above can never be smaller than

$$\tilde{\mathcal{O}} \left(B \left(\prod_{j=1}^d \|W_j\| \right) \sqrt{\frac{d^3}{m}} \right).$$

In particular, even if we assume that $B \left(\prod_{j=1}^d \|W_j\| \right)$ is a constant, the bound becomes trivial once $d \geq \Omega(m^{1/3})$. Finally, and using the same notation, Neyshabur et al. [2017] utilize a PAC-Bayesian analysis to prove a bound scaling as

$$\tilde{\mathcal{O}} \left(B \left(\prod_{j=1}^d \|W_j\| \right) \sqrt{\frac{d^2 h \sum_{i=1}^d \frac{\|W_j\|_F^2}{\|W_j\|^2}}{m}} \right), \quad (3)$$

where h denotes the network width.¹ Again, since $\frac{\|W_j\|_F}{\|W_j\|} \geq 1$ for any parameter matrix W_j , this bound can never be smaller than $\tilde{\mathcal{O}} \left(B \left(\prod_{j=1}^d \|W_j\| \right) \sqrt{\frac{d^3 h}{m}} \right)$, and becomes trivial once $d^3 h \geq \Omega(m^{1/3})$. To summarize, although some of the bounds above have logarithmic or no dependence on the network width, we are not aware of a bound in the literature which avoids a strong dependence on the depth, even if various norms are controlled.

Can this depth dependency be avoided, assuming the norms are sufficiently constrained? We argue that in some cases, it must be true. To see this, let us return to the well-understood case of linear predictors, and consider generalized linear predictors of the form

$$\left\{ \mathbf{x} \mapsto \sigma(\mathbf{w}^\top \mathbf{x}) : \|\mathbf{w}\| \leq M \right\},$$

where $\sigma(z) = \max\{0, z\}$ is the ReLU function. Like plain-vanilla linear predictors, the generalization error of this class is well-known to be $\mathcal{O}(MB/\sqrt{m})$, assuming the inputs satisfy $\|\mathbf{x}\| \leq B$ almost surely.

¹Bartlett et al. [2017] note that this PAC-Bayesian bound is never better than the bound in Eq. (2) derived from a covering numbers argument.

However, it is not difficult to show that this class can be equivalently written as a class of “ultra-thin” ReLU networks of the form

$$\left\{ \mathbf{x} \mapsto \sigma(w_d \cdot \sigma(\dots w_2 \cdot \sigma(\mathbf{w}_1^\top \mathbf{x}))) : \|\mathbf{w}_1\| \cdot \prod_{j=2}^d |w_j| \leq M \right\}, \quad (4)$$

(where \mathbf{w}_1 is a vector and w_2, \dots, w_d are scalars), where the depth d is arbitrary. Therefore, the sample complexity of this class must also scale as $\mathcal{O}(MB/\sqrt{m})$: This depends on the norm product M , but is completely independent of the network depth d as well as the dimension of \mathbf{w}_1 . We argue that a “satisfactory” sample complexity analysis should have similar independence properties when applied on this class.

In more general neural networks, the vector \mathbf{w}_1 and scalars w_2, \dots, w_d become matrices W_1, \dots, W_d , and the simple observation above no longer applies. However, using the same intuition, it is natural to try and derive generalization bounds by controlling $\prod_{j=1}^d \|W_j\|$, where $\|\cdot\|$ is a suitable matrix norm. Perhaps the simplest choice is the spectral norm (and indeed, a product of spectral norms was utilized in some of the previous results mentioned earlier). However, as we formally show in Sec. 5, the spectral norm alone is too weak to get size-independent bounds, even if the network depth is small. Instead, we show that controlling other suitable norms can indeed lead to better depth dependence, or even fully size-independent bounds, improving on earlier works. Specifically, we make the following contributions:

- In Sec. 3, we show that the exponential depth dependence in Rademacher complexity-based analysis (e.g. Neyshabur et al. [2015]) can be avoided by applying contraction to a slightly different object than what has become standard since the work of Bartlett and Mendelson [2002]. For example, for networks with parameter matrices of Frobenius norm at most $M_F(1), \dots, M_F(d)$, the bound in Eq. (1) can be improved to

$$\mathcal{O}\left(\frac{B\sqrt{d} \prod_{j=1}^d M_F(j)}{\sqrt{m}}\right). \quad (5)$$

The technique can also be applied to other types of norm constraints. For example, if we consider an ℓ_1/ℓ_∞ setup, corresponding to the class of depth- d networks, where the 1-norm of each row of W_j is at most $M(j)$, we attain a bound of

$$\mathcal{O}\left(\frac{B\sqrt{d + \log(n)} \cdot \prod_{j=1}^d M(j)}{\sqrt{m}}\right),$$

where n is the input dimension. Again, the dependence on d is polynomial and quite mild. In contrast, Neyshabur et al. [2015] studied a similar setup and only managed to obtain an exponential dependence on d .

- In Sec. 4, we develop a generic technique to convert depth-dependent bounds to depth-independent bounds, assuming some control over any Schatten norm of the parameter matrices (which includes, for instance, the Frobenius norm and the trace norm as special cases). The key observation we utilize is that the prediction function computed by such networks can be approximated by the composition of a shallow network and univariate Lipschitz functions. For example, again assuming that the Frobenius norms of the layers are bounded by $M_F(1), \dots, M_F(d)$, we can further improve Eq. (5) to

$$\tilde{\mathcal{O}}\left(B \prod_{j=1}^d M_F(j) \cdot \min\left\{\frac{1}{m^{1/4}}, \sqrt{\frac{d}{m}}\right\}\right). \quad (6)$$

ignoring logarithmic factors. Note that this can be upper bounded by $\tilde{O}(B \prod_{j=1}^d M_F(j)/m^{1/4})$, which to the best of our knowledge, is the first explicit bound for standard neural networks which is fully size-independent, assuming only suitable norm constraints. Moreover, it captures the depth-independent sample complexity behavior of the network class in Eq. (4) discussed earlier. We also apply this technique to get a depth-independent version of the bound in [Bartlett et al., 2017]: Specifically, if we assume that $\|W_j\| \leq M(j)$ for all j , and $\max_j \frac{\|W_j^T\|_{2,1}}{\|W_j\|} \leq L$, then the bound in Eq. (2) provided by Bartlett et al. [2017] becomes

$$\tilde{O} \left(BL \prod_{j=1}^d M(j) \cdot \sqrt{\frac{d^3}{m}} \right).$$

In contrast, we show the following bound for any $p \geq 1$ (ignoring some lower-order logarithmic factors):

$$\tilde{O} \left(BL \prod_{j=1}^d M(j) \cdot \min \left\{ \frac{\log \left(\frac{B}{\Gamma} \prod_{j=1}^d M_p(j) \right)^{\frac{1}{\frac{2}{3}+p}}}{m^{\frac{1}{2+3p}}}, \sqrt{\frac{d^3}{m}} \right\} \right),$$

where $M_p(j)$ is a bound on the Schatten p -norm of W_j , and Γ is the *largest* possible Euclidean norm of the network's output, over all input vectors of norm at most B . Again, by upper bounding the min by its first argument, we get a bound independent of the depth d , assuming the norms are suitably constrained.

- In Sec. 5, we provide a lower bound, showing that for any p , the class of depth- d , width- h neural networks, where each parameter matrix W_j has Schatten p -norm at most $M_p(j)$, can have Rademacher complexity of at least

$$\Omega \left(\frac{B \prod_{j=1}^d M_p(j) \cdot h^{\max\{0, \frac{1}{2} - \frac{1}{p}\}}}{\sqrt{m}} \right).$$

This somewhat improves on Bartlett et al. [2017, Theorem 3.6], which only showed such a result for $p = \infty$ (i.e. with spectral norm control), and without the h term. For $p = 2$, it matches the upper bound in Eq. (6) in terms of the norm dependencies and B . Moreover, it establishes that controlling the spectral norm alone (and indeed, any Schatten p -norm control with $p > 2$) cannot lead to bounds independent of the size of the network. Finally, the bound shows (similar to Bartlett et al. [2017]) that a dependence on products of norms across layers is generally inevitable.

Besides the above, we provide some additional remarks and observations in Sec. 6. Most of our technical proofs are presented in Sec. 7.

2 Preliminaries

Notation. We use bold-faced letters to denote vectors, and capital letters to denote matrices or fixed parameters (which should be clear from context). Given a vector $\mathbf{w} \in \mathbb{R}^h$, $\|\mathbf{w}\|$ will refer to the Euclidean norm, and for $p \geq 1$, $\|\mathbf{w}\|_p = \left(\sum_{i=1}^h |\mathbf{w}_i|^p \right)^{1/p}$ will refer to the ℓ_p norm. For a matrix W , we use $\|W\|_p$, where $p \in [1, \infty]$, to denote the Schatten p -norm (that is, the p -norm of the spectrum of W , written as a vector).

For example, $p = \infty$ refers to the spectral norm, $p = 2$ refers to the Frobenius norm, and $p = 1$ refers to the trace norm. For the case of the spectral norm, we will drop the ∞ subscript, and use just $\|W\|$. Also, in order to follow standard convention, we use $\|W\|_F$ to denote the Frobenius norm. Finally, given a matrix W and reals $p, q \geq 1$, we let $\|W\|_{p,q} := \left(\sum_k \left(\sum_j |W_{j,k}|^p \right)^{q/p} \right)^{1/q}$ denote the q -norm of the p -norms of the columns of W .

Neural Networks. Given the domain $\mathcal{X} = \{\mathbf{x} : \|\mathbf{x}\| \leq B\}$ in Euclidean space, we consider (scalar or vector-valued) standard neural networks, of the form

$$\mathbf{x} \mapsto W_d \sigma_{d-1}(W_{d-1} \sigma_{d-2}(\dots \sigma_1(W_1 \mathbf{x}))),$$

where each W_j is a parameter matrix (with some fixed dimensions), and each σ_j is some fixed Lipschitz continuous function between Euclidean spaces, satisfying $\sigma_j(\mathbf{0}) = \mathbf{0}$. In the above, we denote d as the depth of the network, and its width h is defined as the maximal row or column dimensions of W_1, \dots, W_d . Without loss of generality, we will assume that σ_j has a Lipschitz constant of at most 1 (otherwise, the Lipschitz constant can be absorbed into the norm constraint of the neighboring parameter matrix). We say that σ is element-wise if it can be written as an application of the same univariate function over each coordinate of its input (in which case, somewhat abusing notation, we will also use σ to denote that univariate function). We say that σ is positive-homogeneous if it is element-wise and satisfies $\sigma(\alpha z) = \alpha \sigma(z)$ for all $\alpha \geq 0$ and $z \in \mathbb{R}$. An important example of the above are ReLU networks, where every σ_j corresponds to applying the (positive-homogeneous) ReLU function $\sigma(z) = \max\{0, z\}$ on each element. To simplify notation, we let W_b^r be shorthand for the matrix tuple $\{W_b, W_{b+1}, \dots, W_r\}$, and $N_{W_b^r}$ denote the function computed by the sub-network composed of layers b through r , that is

$$\mathbf{x} \mapsto W_r \sigma_{r-1}(W_{r-1} \sigma_{r-2}(\dots \sigma_b(W_b \mathbf{x}))).$$

Rademacher Complexity. The results in this paper focus on Rademacher complexity, which is a standard tool to control the uniform convergence (and hence the sample complexity) of given classes of predictors (see Bartlett and Mendelson [2002], Shalev-Shwartz and Ben-David [2014] for more details). Formally, given a real-valued function class \mathcal{H} and some set of data points $\mathbf{x}_1, \dots, \mathbf{x}_m \in \mathcal{X}$, we define the (empirical) Rademacher complexity $\hat{\mathcal{R}}_m(\mathcal{H})$ as

$$\hat{\mathcal{R}}_m(\mathcal{H}) = \mathbb{E}_\varepsilon \left[\sup_{h \in \mathcal{H}} \frac{1}{m} \sum_{i=1}^m \varepsilon_i h(\mathbf{x}_i) \right], \quad (7)$$

where $\varepsilon = (\varepsilon_1, \dots, \varepsilon_m)$ is a vector uniformly distributed in $\{-1, +1\}^m$. Our main results provide bounds on the Rademacher complexity (sometimes independent of $\mathbf{x}_1, \dots, \mathbf{x}_m$, as long as they are assumed to have norm at most B), with respect to classes of neural networks with various norm constraints. Using standard arguments, such bounds can be converted to bounds on the generalization error, assuming access to a sample of m i.i.d. training examples.

3 From Exponential to Polynomial Depth Dependence

To get bounds on the Rademacher complexity of deep neural networks, a reasonable approach (employed in Neyshabur et al. [2015]) is to use a ‘‘peeling’’ argument, where the complexity bound for depth d networks is reduced to a complexity bound for depth $d - 1$ networks, and then applying the reduction d times. For

example, consider the class \mathcal{H}_d of depth- d ReLU real-valued neural networks, with each layer’s parameter matrix with Frobenius norm at most $M_F(j)$. Using some straightforward manipulations, it is possible to show that $\hat{\mathcal{R}}_m(\mathcal{H}_d)$, which by definition equals

$$\mathbb{E}_\epsilon \sup_{h \in \mathcal{H}_d} \frac{1}{m} \sum_{i=1}^m \epsilon_i h(\mathbf{x}_i) = \mathbb{E}_\epsilon \sup_{h \in \mathcal{H}_{d-1}} \sup_{W_d: \|W_d\|_F \leq M_F(d)} \frac{1}{m} \sum_{i=1}^m \epsilon_i W_d \sigma(h(\mathbf{x}_i)),$$

can be upper bounded by

$$M_F(d) \cdot \mathbb{E}_\epsilon \sup_{h \in \mathcal{H}_{d-1}} \left\| \frac{1}{m} \sum_{i=1}^m \epsilon_i \sigma(h(\mathbf{x}_i)) \right\| \leq 2M_F(d) \cdot \mathbb{E}_\epsilon \sup_{h \in \mathcal{H}_{d-1}} \left\| \frac{1}{m} \sum_{i=1}^m \epsilon_i h(\mathbf{x}_i) \right\|. \quad (8)$$

Iterating this inequality d times, one ends up with a bound scaling as $2^d \prod_{j=1}^d M_F(j)$ (as in Neyshabur et al. [2015], see also Eq. (1)). The exponential 2^d factor follows from the 2 factor in Eq. (8), which in turn follows from applying the Rademacher contraction principle to get rid of the σ function. Unfortunately, this 2 factor is generally unavoidable (see the discussion in Ledoux and Talagrand [1991] following Theorem 4.12).

In this section, we point out a simple trick, which can be used to reduce such exponential depth dependencies to polynomial ones. In a nutshell, using Jensen’s inequality, we can rewrite the (scaled) Rademacher complexity $m \cdot \hat{\mathcal{R}}_m(\mathcal{H}) = \mathbb{E}_\epsilon \sup_{h \in \mathcal{H}} \sum_{i=1}^m \epsilon_i h(\mathbf{x}_i)$ as

$$\frac{1}{\lambda} \log \exp \left(\lambda \cdot \mathbb{E}_\epsilon \sup_{h \in \mathcal{H}} \sum_{i=1}^m \epsilon_i h(\mathbf{x}_i) \right) \leq \frac{1}{\lambda} \log \left(\mathbb{E}_\epsilon \sup_{h \in \mathcal{H}} \exp \left(\lambda \sum_{i=1}^m \epsilon_i h(\mathbf{x}_i) \right) \right),$$

where $\lambda > 0$ is an arbitrary parameter. We then perform a “peeling” argument similar to before, resulting in a multiplicative 2 factor after every peeling step. Crucially, these factors accumulate *inside* the log factor, so that the end result contains only a $\log(2^d) = d$ factor, which by appropriate tuning of λ , can be further reduced to \sqrt{d} .

The formalization of this argument depends on the matrix norm we are using, and we will begin with the case of the Frobenius norm. A key technical condition for the argument to work is that we can perform the “peeling” inside the exp function. This is captured by the following lemma:

Lemma 1. *Let σ be a 1-Lipschitz, positive-homogeneous activation function which is applied element-wise (such as the ReLU). Then for any vector-valued class \mathcal{F} , and any convex and monotonically increasing function $g : \mathbb{R} \rightarrow [0, \infty)$,*

$$\mathbb{E}_\epsilon \sup_{f \in \mathcal{F}, W: \|W\|_F \leq R} g \left(\left\| \sum_{i=1}^m \epsilon_i \sigma(W f(\mathbf{x}_i)) \right\| \right) \leq 2 \cdot \mathbb{E}_\epsilon \sup_{f \in \mathcal{F}} g \left(R \cdot \left\| \sum_{i=1}^m \epsilon_i f(\mathbf{x}_i) \right\| \right).$$

Proof. Letting $\mathbf{w}_1, \mathbf{w}_2, \dots, \mathbf{w}_h$ be the rows of the matrix W , we have

$$\left\| \sum_{i=1}^m \epsilon_i \sigma(W f(\mathbf{x}_i)) \right\|^2 = \sum_{j=1}^h \|\mathbf{w}_j\|^2 \left(\sum_{i=1}^m \epsilon_i \sigma \left(\frac{\mathbf{w}_j^\top}{\|\mathbf{w}_j\|} f(\mathbf{x}_i) \right) \right)^2.$$

The supremum of this over all $\mathbf{w}_1, \dots, \mathbf{w}_h$ such that $\|W\|_F^2 = \sum_{j=1}^h \|\mathbf{w}_j\|^2 \leq R^2$ must be attained when $\|\mathbf{w}_j\| = R$ for some j , and $\|\mathbf{w}_i\| = 0$ for all $i \neq j$. Therefore,

$$\mathbb{E}_\epsilon \sup_{f \in \mathcal{F}, W: \|W\|_F \leq R} g \left(\left\| \sum_{i=1}^m \epsilon_i \sigma(W f(\mathbf{x}_i)) \right\| \right) = \mathbb{E}_\epsilon \sup_{f \in \mathcal{F}, \mathbf{w}: \|\mathbf{w}\|=R} g \left(\left\| \sum_{i=1}^m \epsilon_i \sigma(\mathbf{w}^\top f(\mathbf{x}_i)) \right\| \right).$$

Since $g(|z|) \leq g(z) + g(-z)$, this can be upper bounded by

$$\mathbb{E}_\epsilon \sup g \left(\sum_{i=1}^m \epsilon_i \sigma(\mathbf{w}^\top f(\mathbf{x}_i)) \right) + \mathbb{E}_\epsilon \sup g \left(- \sum_{i=1}^m \epsilon_i \sigma(\mathbf{w}^\top f(\mathbf{x}_i)) \right) = 2 \cdot \mathbb{E}_\epsilon \sup g \left(\sum_{i=1}^m \epsilon_i \sigma(\mathbf{w}^\top f(\mathbf{x}_i)) \right),$$

where the equality follows from the symmetry in the distribution of the ϵ_i random variables. The right hand side in turn can be upper bounded by

$$\begin{aligned} 2 \cdot \mathbb{E}_\epsilon \sup_{f \in \mathcal{F}, \mathbf{w}: \|\mathbf{w}\|=R} g \left(\sum_{i=1}^m \epsilon_i \mathbf{w}^\top f(\mathbf{x}_i) \right) &\leq 2 \cdot \mathbb{E}_\epsilon \sup_{f \in \mathcal{F}, \mathbf{w}: \|\mathbf{w}\|=R} g \left(\|\mathbf{w}\| \left\| \sum_{i=1}^m \epsilon_i f(\mathbf{x}_i) \right\| \right) \\ &= 2 \cdot \mathbb{E}_\epsilon \sup_{f \in \mathcal{F}} g \left(R \cdot \left\| \sum_{i=1}^m \epsilon_i f(\mathbf{x}_i) \right\| \right). \end{aligned}$$

(see equation 4.20 in Ledoux and Talagrand [1991]). \square

With this lemma in hand, we can provide a bound on the Rademacher complexity of Frobenius-norm-bounded neural networks, which is as clean as Eq. (1), but where the 2^d factor is replaced by \sqrt{d} :

Theorem 1. *Let \mathcal{H}_d be the class of real-valued networks of depth d over the domain \mathcal{X} , where each parameter matrix W_j has Frobenius norm at most $M_F(j)$, and with activation functions satisfying Lemma 1. Then*

$$\hat{\mathcal{R}}_m(\mathcal{H}_d) \leq \frac{1}{m} \prod_{j=1}^d M_F(j) \cdot \left(\sqrt{2 \log(2)d} + 1 \right) \sqrt{\sum_{i=1}^m \|\mathbf{x}_i\|^2} \leq \frac{B \left(\sqrt{2 \log(2)d} + 1 \right) \prod_{j=1}^d M_F(j)}{\sqrt{m}}.$$

Proof. Fix $\lambda > 0$, to be chosen later. The Rademacher complexity can be upper bounded as

$$\begin{aligned} m \hat{\mathcal{R}}_m(\mathcal{H}_d) &= \mathbb{E}_\epsilon \sup_{N_{W_1^{d-1}}, W_d} \sum_{i=1}^m \epsilon_i W_d \sigma_{d-1}(N_{W_1^{d-1}}(\mathbf{x}_i)) \\ &\leq \frac{1}{\lambda} \log \mathbb{E}_\epsilon \sup \exp \left(\lambda \sum_{i=1}^m \epsilon_i W_d \sigma_{d-1}(N_{W_1^{d-1}}(\mathbf{x}_i)) \right) \\ &\leq \frac{1}{\lambda} \log \mathbb{E}_\epsilon \sup \exp \left(M_F(d) \cdot \left\| \lambda \sum_{i=1}^m \epsilon_i \sigma_{d-1}(N_{W_1^{d-1}}(\mathbf{x}_i)) \right\| \right) \end{aligned}$$

We write this last expression as

$$\begin{aligned} &\frac{1}{\lambda} \log \mathbb{E}_\epsilon \sup_{f, \|W_{d-1}\|_F \leq M_F(d-1)} \exp \left(M_F(d) \cdot \lambda \left\| \sum_{i=1}^m \epsilon_i \sigma_{d-1}(W_{d-1} f(\mathbf{x}_i)) \right\| \right) \\ &\leq \frac{1}{\lambda} \log \left(2 \cdot \mathbb{E}_\epsilon \sup_f \exp \left(M_F(d) \cdot M_F(d-1) \cdot \lambda \left\| \sum_{i=1}^m \epsilon_i f(\mathbf{x}_i) \right\| \right) \right) \end{aligned}$$

where f ranges over all possible functions $\sigma_{d-2} \circ N_{W_1^{d-2}}(\mathbf{x})$. Here we applied Lemma 1 with $g(\alpha) = \exp\{M_F(d)\lambda \cdot \alpha\}$. Repeating the process, we arrive at

$$m\hat{\mathcal{R}}_m(\mathcal{H}_d) \leq \frac{1}{\lambda} \log \left(2^d \cdot \mathbb{E}_\epsilon \exp \left(M\lambda \left\| \sum_{i=1}^m \epsilon_i \mathbf{x}_i \right\| \right) \right) \quad (9)$$

where $M = \prod_{j=1}^d M_F(j)$. Define a random variable

$$Z = M \cdot \left\| \sum_{i=1}^m \epsilon_i \mathbf{x}_i \right\|,$$

(random as a function of the random variables $\epsilon_1, \dots, \epsilon_m$). Then

$$\frac{1}{\lambda} \log \left\{ 2^d \cdot \mathbb{E} \exp \lambda Z \right\} = \frac{d \log(2)}{\lambda} + \frac{1}{\lambda} \log \left\{ \mathbb{E} \exp \lambda (Z - \mathbb{E}Z) \right\} + \mathbb{E}Z. \quad (10)$$

By Jensen's inequality, $\mathbb{E}[Z]$ can be upper bounded by

$$M \sqrt{\mathbb{E}_\epsilon \left[\left\| \sum_{i=1}^m \epsilon_i \mathbf{x}_i \right\|^2 \right]} = M \sqrt{\mathbb{E}_\epsilon \left[\sum_{i,i'=1}^m \epsilon_i \epsilon_{i'} \mathbf{x}_i^\top \mathbf{x}_{i'} \right]} = M \sqrt{\sum_{i=1}^m \|\mathbf{x}_i\|^2}.$$

To handle the $\log \left\{ \mathbb{E} \exp \lambda (Z - \mathbb{E}Z) \right\}$ term in Eq. (10), note that Z is a deterministic function of the i.i.d. random variables $\epsilon_1, \dots, \epsilon_m$, and satisfies

$$Z(\epsilon_1, \dots, \epsilon_i, \dots, \epsilon_m) - Z(\epsilon_1, \dots, -\epsilon_i, \dots, \epsilon_m) \leq 2M\|\mathbf{x}_i\|.$$

This means that Z satisfies a bounded-difference condition, which by the proof of Theorem 6.2 in [Boucheron et al., 2013], implies that Z is sub-Gaussian, with variance factor

$$v = \frac{1}{4} \sum_{i=1}^m (2M\|\mathbf{x}_i\|)^2 = M^2 \sum_{i=1}^m \|\mathbf{x}_i\|^2,$$

and satisfies

$$\frac{1}{\lambda} \log \left\{ \mathbb{E} \exp \lambda (Z - \mathbb{E}Z) \right\} \leq \frac{1}{\lambda} \frac{\lambda^2 M^2 \sum_{i=1}^m \|\mathbf{x}_i\|^2}{2} = \frac{\lambda M^2 \sum_{i=1}^m \|\mathbf{x}_i\|^2}{2}.$$

Choosing $\lambda = \frac{\sqrt{2 \log(2)d}}{M \sqrt{\sum_{i=1}^m \|\mathbf{x}_i\|^2}}$ and using the above, we get that Eq. (9) can be upper bounded as follows:

$$\frac{1}{\lambda} \log \left\{ 2^d \cdot \mathbb{E} \exp \lambda Z \right\} \leq \mathbb{E}Z + \sqrt{2 \log(2)d} \cdot M \sqrt{\sum_{i=1}^m \|\mathbf{x}_i\|^2} \leq M \left(\sqrt{2 \log(2)d} + 1 \right) \sqrt{\sum_{i=1}^m \|\mathbf{x}_i\|^2},$$

from which the result follows. \square

Remark 1. We note that for simplicity, the bound in Thm. 1 is stated for real-valued networks, but the argument easily carries over to vector-valued networks, composed with some real-valued Lipschitz loss function. In that case, one uses a variant of Lemma 1 to peel off the losses, and then proceed in the same manner as in the proof of Thm. 1. We omit the precise details for brevity.

A result similar to the above can also be derived for other entry-based matrix norms. For example, given a matrix W , let $\|W\|_{1,\infty}$ denote the maximal 1-norm of its rows, and consider the class \mathcal{H}_d of depth- d networks, where each parameter matrix W_j satisfies $\|W_j\|_{1,\infty} \leq M(j)$ for all j (this corresponds to a setting, also studied in Neyshabur et al. [2015], where the 1-norm of the weights of each neuron in the network is bounded). In this case, we can derive a variant of Lemma 1, which in fact does not require positive-homogeneity of the activation function:

Lemma 2. *Let σ be a 1-Lipschitz activation function with $\sigma(0) = 0$, applied element-wise. Then for any vector-valued class \mathcal{F} , and any convex and monotonically increasing function $g : \mathbb{R} \rightarrow [0, \infty)$,*

$$\mathbb{E}_\epsilon \sup_{f \in \mathcal{F}, W: \|W\|_{1,\infty} \leq R} g \left(\left\| \sum_{i=1}^m \epsilon_i \sigma(W f(\mathbf{x}_i)) \right\|_\infty \right) \leq 2 \cdot \mathbb{E}_\epsilon \sup_{f \in \mathcal{F}} g \left(R \cdot \left\| \sum_{i=1}^m \epsilon_i f(\mathbf{x}_i) \right\|_\infty \right), \quad (11)$$

where $\|\cdot\|_\infty$ denotes the vector infinity norm.

Using the same technique as before, we can use this lemma to get a bound on the Rademacher complexity for \mathcal{H}_d :

Theorem 2. *Let \mathcal{H}_d be the class of real-valued networks of depth d over the domain \mathcal{X} , where $\|W_j\|_{1,\infty} \leq M(j)$ for all $j \in \{1, \dots, d\}$, and with activation functions satisfying the condition of Lemma 2. Then*

$$\hat{\mathcal{R}}_m(\mathcal{H}_d) \leq \frac{2}{m} \prod_{j=1}^d M(j) \cdot \sqrt{d+1+\log(n)} \cdot \sqrt{\max_{j \in \{1, \dots, n\}} \sum_{i=1}^m x_{i,j}^2} \leq \frac{2B\sqrt{d+1+\log(n)} \cdot \prod_{j=1}^d M(j)}{\sqrt{m}},$$

where $x_{i,j}$ is the j -th coordinate of the vector \mathbf{x}_i .

The proofs of the theorem, as well as Lemma 2, appear in Sec. 7.

The constructions used in the results of this section use the function $g(z) = \exp(\lambda z)$ together with its inverse $g^{-1}(z) = \frac{1}{\lambda} \log(z)$, to get depth dependencies scaling as $g^{-1}(2^d)$. Thus, it might be tempting to try and further improve the depth dependence, by using other functions g for which g^{-1} increases sub-logarithmically. Unfortunately, the argument still requires us to control $\mathbb{E}_\epsilon g(\|\sum_{i=1}^m \epsilon_i \mathbf{x}_i\|)$, which is difficult if g increases more than exponentially. In the next section, we introduce a different idea, which under suitable assumptions, allows us to get rid of the depth dependence altogether.

4 From Depth Dependence to Depth Independence

In this section, we develop a general result, which allows one to convert any depth-dependent bound on the Rademacher complexity of neural networks, to a depth-independent one, assuming that the Schatten p -norms of the parameter matrices (for any $p \in [1, \infty)$) is controlled. We develop and formalize the main result in Subsection 4.1, and provide applications in Subsection 4.2. The proofs the results in this section appear in Sec. 7.

4.1 A General Result

To motivate our approach, let us consider a special case of depth- d networks, where

- Each parameter matrix W_1, \dots, W_{d-1} is constrained to be diagonal and of size $h \times h$.

- The Frobenius norm of every W_1, \dots, W_d is at most 1.
- All activation functions are the identity (so the network computes a linear function).

Letting \mathbf{w}_i be the diagonal of W_i , such networks are equivalent to

$$\mathbf{x} \mapsto (\mathbf{w}_d \circ \mathbf{w}_{d-1} \circ \dots \circ \mathbf{w}_1)^\top \mathbf{x},$$

where \circ denotes element-wise product. Therefore, if we would like the network to compute a non-trivial function, we clearly need that $\mathbf{w}_d \circ \dots \circ \mathbf{w}_1$ be bounded away from zero (e.g., not exponentially small in d), while still satisfying the constraint $\|\mathbf{w}_j\| \leq 1$ for all j . In fact, the only way to satisfy both requirements simultaneously is if $\mathbf{w}_1, \dots, \mathbf{w}_d$ are all close to some 1-sparse unit vector, which implies that the matrices W_1, \dots, W_d must be close to being rank-1.

It turns out that this intuition holds much more generally, even if we do not restrict ourselves to identity activations and diagonal parameter matrices as above. Essentially, what we can show is that if some network computes a non-trivial function, and the product of its Schatten p -norms (for any $p < \infty$) is bounded, then there must be at least one parameter matrix, which is not far from being rank-1. Therefore, if we replace that parameter matrix by an appropriate rank-1 matrix, the function computed by the network does not change much. This is captured by the following theorem:

Theorem 3. *For any $p \in [1, \infty)$, any network $N_{W_1^d}$ such that $\sup_{\mathbf{x} \in \mathcal{X}} \|N_{W_1^d}(\mathbf{x})\| \geq \Gamma$ and $\prod_{j=1}^d \|W_j\|_p \leq M$, and for any $r \in \{1, \dots, d\}$, there exists another network $N_{\tilde{W}_1^d}$ (of the same depth and layer dimensions) with the following properties:*

- $\tilde{W}_1^d = \{\tilde{W}_1, \dots, \tilde{W}_d\}$ is identical to W_1^d , except for the parameter matrix \tilde{W}_r in the r -th layer, which is of rank at most 1, and equals $s\mathbf{u}\mathbf{v}^\top$ where $s, \mathbf{u}, \mathbf{v}$ are some leading singular value and singular vectors pairs of W_r .
- $\sup_{\mathbf{x} \in \mathcal{X}} \|N_{W_1^d}(\mathbf{x}) - N_{\tilde{W}_1^d}(\mathbf{x})\| \leq B \left(\prod_{j=1}^d \|W_j\| \right) \left(\frac{2p \log(MB/\Gamma)}{r} \right)^{1/p}$.

We now make the following crucial observation: A real-valued network with a rank-1 parameter matrix $W_r = s\mathbf{u}\mathbf{v}^\top$ computes the function

$$\mathbf{x} \mapsto W_d \sigma_{d-1}(\dots \sigma_r(s\mathbf{u}\mathbf{v}^\top \sigma_{r-1}(\dots \sigma_1(W_1 \mathbf{x}) \dots))) .$$

This can be seen as the composition of the *depth- r* network

$$\mathbf{x} \mapsto \mathbf{v}^\top \sigma_{r-1}(\dots \sigma_1(W_1 \mathbf{x}) \dots),$$

and the *univariate* function

$$x \mapsto W_d \sigma_{d-1}(\dots \sigma_r(s\mathbf{u}x)) .$$

Moreover, the norm constraints imply that the latter function is Lipschitz. Therefore, the class of networks we are considering is a subset of the class of depth- r networks composed with univariate Lipschitz functions. Fortunately, given *any* class with bounded complexity, one can effectively bound the Rademacher complexity of its composition with univariate Lipschitz functions, as formalized in the following theorem.

Theorem 4. Let \mathcal{H} be a class of functions from Euclidean space to $[-R, R]$. Let $\mathcal{F}_{L,a}$ be the class of L -Lipschitz functions from $[-R, R]$ to \mathbb{R} , such that $f(0) = a$ for some fixed a . Letting $\mathcal{F}_{L,a} \circ \mathcal{H} := \{f(h(\cdot)) : f \in \mathcal{F}_{L,a}, h \in \mathcal{H}\}$, its Rademacher complexity satisfies

$$\hat{\mathcal{R}}_m(\mathcal{F}_{L,a} \circ \mathcal{H}) \leq cL \left(\frac{R}{\sqrt{m}} + \log^{3/2}(m) \cdot \hat{\mathcal{R}}_m(\mathcal{H}) \right),$$

where $c > 0$ is a universal constant.

Remark 2. The $\log^{3/2}(m) \cdot \hat{\mathcal{R}}_m(\mathcal{H})$ can be replaced by $\log(m) \cdot \hat{\mathcal{G}}_m(\mathcal{H})$, where $\hat{\mathcal{G}}_m(\mathcal{H})$ is the empirical Gaussian complexity of \mathcal{H} – see the proof in Sec. 7 for details.

Combining these ideas, our plan of attack is the following: Given some class of depth- d networks, and an arbitrary parameter $r \in \{1, \dots, d\}$, we use Thm. 3 to relate their Rademacher complexity to the complexity of similar networks, but where one of the first r parameter matrices is of rank-1. We then use Thm. 4 to bound that complexity in turn using the Rademacher complexity of depth- r networks. Crucially, the resulting bound has no explicit dependence on the original depth d , only on the new parameter r . Formally, we have the following theorem, which is the main result of this section:

Theorem 5. Consider the following hypothesis class of networks on $\mathcal{X} = \{\mathbf{x} : \|\mathbf{x}\| \leq B\}$:

$$\mathcal{H} = \left\{ N_{W_1^d} : \forall j \in \{1 \dots d\}, W_j \in \mathcal{W}_j, \max \left\{ \frac{\|W_j\|}{M(j)}, \frac{\|W_j\|_p}{M_p(j)} \right\} \leq 1 \right\},$$

for some parameters $p, \Gamma \geq 1, \{M(j), M_p(j), \mathcal{W}_j\}_{j=1}^d$. Also, for any $r \in \{1, \dots, d\}$, define

$$\mathcal{H}_r = \left\{ N_{W_1^r} : \begin{array}{l} N_{W_1^r} \text{ maps to } \mathbb{R} \\ \forall j \in \{1 \dots r-1\}, W_j \in \mathcal{W}_j \\ \forall j \in \{1 \dots r\}, \max \left\{ \frac{\|W_j\|}{M(j)}, \frac{\|W_j\|_p}{M_p(j)} \right\} \leq 1 \end{array} \right\}.$$

Finally, for $m > 1$, let $\ell \circ \mathcal{H} = \{(\ell_1(h(\mathbf{x}_1)), \dots, \ell_m(h(\mathbf{x}_m))) : h \in \mathcal{H}\}$, where ℓ_1, \dots, ℓ_m are real-valued loss functions which are $\frac{1}{\gamma}$ -Lipschitz and satisfy $\ell_1(\mathbf{0}) = \ell_2(\mathbf{0}) = \dots = \ell_m(\mathbf{0})$.

Then the Rademacher complexity $\hat{\mathcal{R}}_m(\ell \circ \mathcal{H})$ is upper bounded by

$$\frac{cB \prod_{j=1}^d M(j)}{\gamma} \left(\min_{r \in \{1, \dots, d\}} \left\{ \frac{\log^{3/2}(m) \cdot \hat{\mathcal{R}}_m(\mathcal{H}_r)}{B \prod_{j=1}^r M(j)} + \left(\frac{\log \left(\frac{B}{\Gamma} \prod_{j=1}^d M_p(j) \right)}{r} \right)^{1/p} \right\} + \frac{1}{\sqrt{m}} \right),$$

where $c > 0$ is a universal constant.

In particular, one can upper bound this result by any choice of r in $\{1, \dots, d\}$. By tuning r appropriately, we can get bounds independent of the depth d . In the next subsection, we provide some concrete applications for specific choices of \mathcal{H} .

Remark 3. The parameters Γ and γ , which divide the norm terms in Thm. 5, are both closely related to the notion of a margin. Indeed, if we consider binary or multi-class classification, then bounds as above w.r.t. $\frac{1}{\gamma}$ -Lipschitz losses can be converted to a bound on the misclassification error rate in terms of the average γ -margin error on the training data (see Bartlett et al. [2017, Section 3.1] for a discussion). Also, Γ can be viewed as the “maximal” margin attainable over the input domain.

4.2 Applications of Thm. 5

In this section we exemplify how Thm. 5 can be used to obtain depth-independent bounds on the sample complexity of various classes of neural networks. The general technique is as follows: First, we prove a bound on $\hat{R}_m(\mathcal{H}_r)$, which generally depends on the depth r , and scales as r^α/\sqrt{m} for some $\alpha > 0$. Then, we plug it into Thm. 5, and utilize the following lemma to tune r appropriately:

Lemma 3. *For any $\alpha > 0$, $\beta \in (0, 1]$ and $b, c, n \geq 1$, it holds that*

$$\min \left\{ \min_{r \in \{1, \dots, d\}} \frac{cr^\alpha}{n} + \frac{b}{r^\beta}, \frac{d^\alpha}{n} \right\} \leq \min \left\{ 3 \cdot \frac{b^{\frac{\alpha}{\alpha+\beta}}}{(n/c)^{\frac{\beta}{\alpha+\beta}}}, \frac{d^\alpha}{n} \right\}.$$

We begin with proving a depth-independent version of Thm. 1. That theorem implies that for the class \mathcal{H}_r of depth- r neural networks with Frobenius norm bounds $M_F(1), \dots, M_F(r)$ (up to and including $r = d$),

$$\hat{R}_m(\mathcal{H}_r) \leq \mathcal{O} \left(B \prod_{j=1}^d M_F(j) \sqrt{\frac{r}{m}} \right) \quad (12)$$

Plugging this into Thm. 5, and using Lemma 3, it is straightforward to derive the following corollary (see Sec. 7 for a formal derivation):

Corollary 1. *Let \mathcal{H} be the class of depth- d neural networks, where each parameter matrix W_j satisfies $\|W_j\|_F \leq M_F(j)$, and with 1-Lipschitz, positive-homogeneous, element-wise activation functions. Assuming the loss function ℓ and \mathcal{H} satisfy the conditions of Thm. 5 (with the sets \mathcal{W}_j being unconstrained), it holds that*

$$\hat{\mathcal{R}}_m(\ell \circ \mathcal{H}) \leq \mathcal{O} \left(\frac{B \prod_{j=1}^d M_F(j)}{\gamma} \cdot \min \left\{ \frac{\log^{3/4}(m) \sqrt{\log \left(\frac{B}{F} \prod_{j=1}^d M_F(j) \right)}}{m^{1/4}}, \sqrt{\frac{d}{m}} \right\} \right).$$

Ignoring logarithmic factors and replacing the min by its first argument, the bound in the corollary is at most

$$\tilde{\mathcal{O}} \left(\frac{B \prod_{j=1}^d M_F(j)}{\gamma m^{1/4}} \right).$$

The bound is completely independent of either the width or the depth of the networks. In other words, it is possible to make this bound smaller than any fixed ϵ , with a sample size m independent of the network's size, as long as $\frac{B}{F} \prod_{j=1}^d M_F(j)$ is bounded. On the other hand, the bound in Corollary 1 is also bounded by

$$\mathcal{O} \left(\frac{B \prod_{j=1}^d M_F(j) \sqrt{d}}{\gamma \sqrt{m}} \right),$$

which is the bound one would get from an immediate application of Thm. 1, and implies that the asymptotic rate (as a function of m) is still maintained.

Next, we apply Thm. 5 to the results in Bartlett et al. [2017], which as discussed in the introduction, provide a depth-dependent bound using a different set of norms. Specifically, they obtain the following intermediary result in deriving their generalization bound:

Theorem 6 (Bartlett et al. [2017]). *Let \mathcal{H} be the hypothesis class of depth- d , width- h real-valued networks on $\mathcal{X} = \{\mathbf{x} : \|\mathbf{x}\| \leq B\}$, using 1-Lipschitz activation functions, given by*

$$\mathcal{H} = \left\{ N_{W_1^d} : \forall j \in \{1, \dots, d\}, \|W_j\| \leq M(j), \frac{\|W_j^T\|_{2,1}}{\|W_j\|} \leq L(j) \right\},$$

for some fixed parameters $\{L(j), M(j)\}_{j=1}^d$. Then the Rademacher complexity $\hat{\mathcal{R}}_m(\mathcal{H})$ is upper-bounded by

$$\hat{\mathcal{R}}_m(\mathcal{H}) \leq \mathcal{O} \left(\frac{B \log(h) \log(m) \prod_{j=1}^d M(j)}{\sqrt{m}} \cdot \left(\sum_{j=1}^d L(j)^{2/3} \right)^{3/2} \right).$$

As discussed in the introduction, $L(j)$ can never be smaller than 1, hence the bound scales at least as $\sqrt{d^3/m}$. However, using the bound above together with Thm. 5 and Lemma 3, we can get the following corollary (where for simplicity, we assume that $L(j)$ for all j are uniformly bounded by some L):

Corollary 2. *Let \mathcal{H} be the class of depth- d , width- h networks with 1-Lipschitz, positive-homogeneous, element-wise activation functions. Assuming the loss function ℓ and \mathcal{H} satisfy the conditions of Thm. 5 (with $W_j = \left\{ W_j : \frac{\|W_j^T\|_{2,1}}{\|W_j\|} \leq L \right\}$ for all j), it holds that the Rademacher complexity $\hat{\mathcal{R}}_m(\ell \circ \mathcal{H})$ is at most*

$$\mathcal{O} \left(\frac{BL \log(h) \log(m) \prod_{j=1}^d M(j)}{\gamma} \cdot \min \left\{ \frac{\log \left(\frac{B}{\Gamma} \prod_{j=1}^d M_p(j) \right)^{\frac{1}{\frac{3}{2}+p}} \left(\log^{3/2}(m) \right)^{1+\frac{1}{\frac{3}{2}p}}}{m^{\frac{1}{2+3p}}}, \frac{d^{3/2}}{\sqrt{m}} \right\} \right).$$

As before, by replacing the min by its first argument, we get a bound which is fully independent of the network size, assuming the norms are suitably bounded. To give a concrete example, if we take $p = 2$ (so that the $M_p(j)$ constraints correspond to the Frobenius norm), and ignore lower-order logarithmic factors, we get a bound scaling as

$$\hat{\mathcal{R}}_m(\ell \circ \mathcal{H}) \leq \tilde{\mathcal{O}} \left(\frac{BL \prod_{j=1}^d M(j)}{\gamma} \cdot \min \left\{ \sqrt[4]{\frac{\log^{3/2} \left(\frac{B}{\Gamma} \prod_{j=1}^d M_p(j) \right)}{m^{1/2}}}, \sqrt{\frac{d^3}{m}} \right\} \right).$$

In contrast, a direct application of Thm. 6 in the same setting leads to a bound of

$$\hat{\mathcal{R}}_m(\ell \circ \mathcal{H}) \leq \tilde{\mathcal{O}} \left(\frac{BL \prod_{j=1}^d M(j)}{\gamma} \cdot \sqrt{\frac{d^3}{m}} \right).$$

Finally, we note that since the Eq. (3), based on a PAC-Bayes analysis, is always weaker than Eq. (2) (as noted in Bartlett et al. [2017]), Corollary 2 also gives a size-independent version of Eq. (3).

5 A Lower Bound for Schatten Norms

In this section, we present a lower bound on the Rademacher complexity, for the class of neural networks with parameter matrices of bounded Schatten norms. The formal result is the following:

Theorem 7. Let \mathcal{H} be the class of depth- d , width- h neural networks, where each parameter matrix W_j satisfies $\|W_j\|_p \leq M_p(j)$ for some Schatten p -norm $\|\cdot\|_p$ (and where we use the convention that $p = \infty$ refers to the spectral norm). Then there exists a choice of $\frac{1}{\gamma}$ -Lipschitz loss ℓ and data points $\mathbf{x}_1, \dots, \mathbf{x}_m \in \mathcal{X}$, with respect to which

$$\hat{\mathcal{R}}_m(\ell \circ \mathcal{H}) \geq \Omega \left(\frac{B \prod_{j=1}^d M_p(j) \cdot h^{\max\{0, \frac{1}{2} - \frac{1}{p}\}}}{\gamma \sqrt{m}} \right).$$

This theorem strengthens Theorem 3.6 in Bartlett et al. [2017], which only considered the case $p = \infty$, and did not have a dependence on h (on the other hand, they consider bounds which hold for any choice of $\mathbf{x}_1, \dots, \mathbf{x}_m$, while we consider bounds uniform over $\mathbf{x}_1, \dots, \mathbf{x}_m$ for simplicity). Furthermore, it has the following implications:

- Like Bartlett et al. [2017], the theorem implies that by controlling just the norms of each parameter matrix, a dependence on the product of the norms is generally inevitable.
- For $p = \infty$, we see that there is an inevitable $h^{1/2}$ factor in the bound, which implies that controlling the spectral norm is insufficient to get size-independent bounds (at least, independent of the width h). More generally, any Schatten p -norm control with $p > 2$ will be insufficient to get such size independence.
- For $p = 2$ (i.e. Frobenius norm bounds $M_F(1), \dots, M_F(d)$), the lower bound becomes size-independent, and on the order of $\frac{B \prod_{j=1}^d M_F(j)}{\sqrt{m}}$. This matches our upper bound in Corollary 1 up to logarithmic factors, except for the worse polynomial dependence on m .

6 Additional Remarks

6.1 Post-hoc Guarantees

So far we have proved upper bounds on the empirical Rademacher complexity of a fixed class of neural networks, of the form

$$\mathcal{H}_L = \{h : \text{compl}(h) \leq L\},$$

for some complexity measure compl and a parameter L . These imply high-probability learning guarantees for algorithms which return predictors in \mathcal{H}_L . However, in the context of norm-based constraints, practical algorithms for neural networks usually perform unconstrained optimization, and therefore are not guaranteed a-priori to return a predictor in some fixed \mathcal{H}_L . Fortunately, it is straightforward to convert these bounds into probabilistic guarantees for *any* neural network h , with the bound scaling appropriately with the complexity $\text{compl}(h)$ of the particular network. We note that such post-hoc guarantees have also been stated in the context of some previous sample complexity bounds for neural networks (e.g., Bartlett et al. [2017], Neyshabur et al. [2017]). It is achieved, for instance, by a union bound over, say, a doubling scale of the complexity. We refer to the proof of the margin bound of Koltchinskii and Panchenko [2002, Theorem 2] for an example of this technique.

6.2 Complexity of Lipschitz Networks

In proving the results of Sec. 4, the key element has been the observation that under appropriate norm constraints, a neural network must have a layer with parameter matrix close to being rank-1, and therefore the network can be viewed as a composition of a shallower network and a univariate Lipschitz function. In fact, this can be generalized: Whenever we have a network with parameter matrix close to being rank- k , we can view it as a composition of a shallow network and a Lipschitz function on \mathbb{R}^k . Although we do not develop this idea further in this paper, this observation might be useful in analyzing other types of neural network classes.

Taking this to the extreme, we can also bound the complexity of neural networks computing Lipschitz functions, by studying the complexity of all Lipschitz functions over the domain \mathcal{X} . It is easily verified that in our setting, if we consider the class \mathcal{H} of depth- d networks, where each parameter matrix has spectral norm at most $M(j)$, then the network must be $\prod_{j=1}^d M(j)$ -Lipschitz. Using well-known estimates of the covering numbers of Lipschitz functions over \mathcal{X} , we get that

$$\hat{\mathcal{R}}_m(\ell \circ \mathcal{H}) \leq \mathcal{O} \left(\frac{B \prod_{j=1}^d M(j)}{\gamma m^{1/\dim(\mathcal{X})}} \right),$$

where the loss ℓ is assumed $\frac{1}{\gamma}$ -Lipschitz, and where $\dim(\mathcal{X})$ is the dimensionality of \mathcal{X} . Of course, the bound has a very bad dependence on the input dimension (or equivalently, the width of the first layer in the network), but on the other hand, has no dependence on the network depth, nor on any matrix norm other than the spectral norm. Again, as discussed in the previous subsection, it is also possible to use this bound to get post-hoc guarantees, without constraining the Lipschitz parameter of the learned network in advance.

7 Proofs

7.1 Proofs of Lemma 2 and Thm. 2

We first prove Lemma 2. Letting \mathbf{w}_j denote the j -th row of a matrix W , we have

$$\begin{aligned} \sup_{f \in \mathcal{F}, W: \|W\|_{1, \infty} \leq R} g \left(\left\| \sum_{i=1}^m \epsilon_i \sigma(W f(\mathbf{x}_i)) \right\|_{\infty} \right) &= \sup_{f \in \mathcal{F}, W: \|\mathbf{w}_j\|_1 \leq R} \max_k g \left(\left\| \sum_{i=1}^m \epsilon_i \sigma(\mathbf{w}_k^\top f(\mathbf{x}_i)) \right\| \right) \\ &= \sup_{f \in \mathcal{F}, \|\mathbf{w}\|_1 \leq R} g \left(\left\| \sum_{i=1}^m \epsilon_i \sigma(\mathbf{w}^\top f(\mathbf{x}_i)) \right\| \right). \end{aligned}$$

Since $g(|z|) \leq g(z) + g(-z)$, the left-hand side of Eq. (11) is upper bounded by

$$\mathbb{E}_\epsilon \sup_{f \in \mathcal{F}, \|\mathbf{w}\|_1 \leq R} g \left(\sum_{i=1}^m \epsilon_i \sigma(\mathbf{w}^\top f(\mathbf{x}_i)) \right) + \mathbb{E}_\epsilon \sup_{f \in \mathcal{F}, \|\mathbf{w}\|_1 \leq R} g \left(- \sum_{i=1}^m \epsilon_i \sigma(\mathbf{w}^\top f(\mathbf{x}_i)) \right)$$

and the proof is concluded exactly as in Lemma 1 by appealing to Ledoux and Talagrand [1991].

We now turn to Thm. 2, whose proof is rather similar to that of Thm. 1. Fixing $\lambda > 0$ to be chosen later,

the Rademacher complexity can be upper bounded as

$$\begin{aligned}
m\hat{\mathcal{R}}_m(\mathcal{H}_d) &= \mathbb{E}_\epsilon \sup_{N_{W_1^{d-1}}, W_d} \sum_{i=1}^m \epsilon_i W_d \sigma_{d-1}(N_{W_1^{d-1}}(\mathbf{x}_i)) \\
&\leq \frac{1}{\lambda} \log \mathbb{E}_\epsilon \sup \exp \lambda \sum_{i=1}^m \epsilon_i W_d \sigma_{d-1}(N_{W_1^{d-1}}(\mathbf{x}_i)) \\
&\leq \frac{1}{\lambda} \log \mathbb{E}_\epsilon \sup \exp \left\{ M_F(d) \cdot \left\| \lambda \sum_{i=1}^m \epsilon_i \sigma_{d-1}(N_{W_1^{d-1}}(\mathbf{x}_i)) \right\|_\infty \right\}
\end{aligned}$$

Applying the same argument as in the proof of Thm. 1, and using Lemma 2, we can upper bound the above by

$$m\hat{\mathcal{R}}_m(\mathcal{H}_d) \leq \frac{1}{\lambda} \log \left(2^d \cdot \mathbb{E}_\epsilon \exp \left(M\lambda \left\| \sum_{i=1}^m \epsilon_i \mathbf{x}_i \right\|_\infty \right) \right), \quad (13)$$

where $M = \prod_{j=1}^d M(j)$. Letting $x_{i,j}$ denote the j -th coordinate of \mathbf{x}_i , and using symmetry, the expectation inside the log can be re-written as

$$\begin{aligned}
\mathbb{E}_\epsilon \exp \left(M\lambda \cdot \max_j \left| \sum_{i=1}^m \epsilon_i x_{i,j} \right| \right) &\leq \sum_{j=1}^n \mathbb{E}_\epsilon \exp \left(M\lambda \cdot \left| \sum_{i=1}^m \epsilon_i x_{i,j} \right| \right) \\
&\leq \sum_{j=1}^n \mathbb{E}_\epsilon \left[\exp \left(M\lambda \sum_{i=1}^m \epsilon_i x_{i,j} \right) + \exp \left(-M\lambda \sum_{i=1}^m \epsilon_i x_{i,j} \right) \right] \\
&= 2 \sum_{j=1}^n \mathbb{E}_\epsilon \exp \left(M\lambda \sum_{i=1}^m \epsilon_i x_{i,j} \right) = 2 \sum_{j=1}^n \prod_{i=1}^m \mathbb{E}_\epsilon \exp (M\lambda \epsilon_i x_{i,j}) \\
&= 2 \sum_{j=1}^n \prod_{i=1}^m \frac{\exp(M\lambda x_{i,j}) + \exp(-M\lambda x_{i,j})}{2} \leq 2 \sum_{j=1}^n \exp \left(M^2 \lambda^2 \sum_{i=1}^m x_{i,j}^2 \right),
\end{aligned}$$

where in the last step we used the fact that $\frac{\exp(z) + \exp(-z)}{2} \leq \exp(z^2/2)$. Further upper bounding this by $2n \max_j \exp \left(M^2 \lambda^2 \sum_{i=1}^d x_{i,j}^2 \right)$ and plugging back to Eq. (13), we get

$$\frac{1}{\lambda} \log \left(2^{d+1} n \cdot \max_j \exp \left(M^2 \lambda^2 \sum_{i=1}^m x_{i,j}^2 \right) \right) = \frac{d+1 + \log(n)}{\lambda} + M^2 \lambda \max_j \sum_{i=1}^m x_{i,j}^2.$$

Choosing $\lambda = \sqrt{\frac{d+1+\log(n)}{M^2 \max_j \sum_{i=1}^m x_{i,j}^2}}$, we can upper bound the above by

$$2M \sqrt{(d+1 + \log(n)) \max_j \sum_{i=1}^m x_{i,j}^2},$$

from which the result follows.

7.2 Proof of Thm. 3

The proof will build on the following few technical lemmas.

Lemma 4. For any matrix W , and any Schatten p -norm $\|\cdot\|_p$ such that $p < \infty$, there exists a rank-1 matrix \tilde{W} of the same size such that

$$\|\tilde{W}\| \leq \|W\|, \quad \|\tilde{W}\|_p \leq \|W\|_p, \quad \|W - \tilde{W}\| \leq (\|W\|_p^p - \|W\|^p)^{1/p}.$$

Proof. Let USV^\top denote the SVD decomposition of W , where $S = \text{diag}(s_1, s_2, \dots, s_r)$, and choose $\tilde{W} = \mathbf{u}_1 s_1 \mathbf{v}_1^\top$, where $\mathbf{u}_1, \mathbf{v}_1, s_1$ are top singular vectors and values of W . The first two inequalities in the lemma are easy to verify. As to the third inequality, using the unitaral invariance of the spectral norm, we have

$$\|W - \tilde{W}\| = \|USV^\top - \mathbf{u}_1 s_1 \mathbf{v}_1^\top\| = \|U \text{diag}(0, s_2, \dots, s_r) V^\top\| = s_2 \leq \left(\sum_{j=2}^r s_j^p \right)^{1/p} = \left(\sum_{j=1}^r s_j^p - s_1^p \right)^{1/p},$$

which equals $(\|W\|_p^p - \|W\|^p)^{1/p}$. \square

Lemma 5. Given network parameters $W_1^d = \{W_1, \dots, W_d\}$, let $\tilde{W}_1^d = \{W_1, \dots, W_{r-1}, \tilde{W}_r, W_{r+1}, \dots, W_d\}$ be the same parameters, where the parameter matrix W_r of the r -th layer (for some fixed $r \in \{1, \dots, d\}$) is changed to some other matrix \tilde{W}_r . Then

$$\sup_{\mathbf{x} \in \mathcal{X}} \|N_{W_1^d}(\mathbf{x}) - N_{\tilde{W}_1^d}(\mathbf{x})\| \leq B \left(\prod_{j=1}^d \|W_j\| \right) \frac{\|W_r - \tilde{W}_r\|}{\|W_r\|}.$$

Proof. By a simple calculation, we have that the Lipschitz constant of the function $N_{W_b^r}$ is at most $\prod_{j=b}^r \|W_j\|$.

Assume for now that $2 \leq r \leq d-1$. By definition, we have

$$N_{W_1^d}(\mathbf{x}) = N_{W_{r+1}^d}(\sigma_r(W_r \sigma_{r-1}(N_{W_1^{r-1}}(\mathbf{x}))))$$

and

$$N_{\tilde{W}_1^d}(\mathbf{x}) = N_{W_{r+1}^d}(\sigma_r(\tilde{W}_r \sigma_{r-1}(N_{W_1^{r-1}}(\mathbf{x})))) .$$

The Lipschitz constant of the function $N_{W_{r+1}^d}$ is at most $\prod_{j=r+1}^d \|W_j\|$, and the norm of $N_{W_1^r}(\mathbf{x})$ is at most

$\|\mathbf{x}\| \prod_{j=1}^r \|W_j\|$. Therefore, for any $\mathbf{x} \in \mathcal{X}$,

$$\begin{aligned}
\|N_{W_1^d}(\mathbf{x}) - N_{\tilde{W}_1^d}(\mathbf{x})\| &= \|N_{W_{r+1}^d}(\sigma_r(W_r \sigma_{r-1}(N_{W_1^{r-1}}(\mathbf{x})))) - N_{W_{r+1}^d}(\sigma_r(\tilde{W}_r \sigma_{r-1}(N_{W_1^{r-1}}(\mathbf{x}))))\| \\
&\leq \left(\prod_{j=r+1}^d \|W_j\| \right) \cdot \|\sigma_r(W_r \sigma_{r-1}(N_{W_1^{r-1}}(\mathbf{x}))) - \sigma_r(\tilde{W}_r \sigma_{r-1}(N_{W_1^{r-1}}(\mathbf{x})))\| \\
&\leq \left(\prod_{j=r+1}^d \|W_j\| \right) \cdot \|W_r \sigma_{r-1}(N_{W_1^{r-1}}(\mathbf{x})) - \tilde{W}_r \sigma_{r-1}(N_{W_1^{r-1}}(\mathbf{x}))\| \\
&\leq \left(\prod_{j=r+1}^d \|W_j\| \right) \cdot \|W_r - \tilde{W}_r\| \cdot \|\sigma_{r-1}(N_{W_1^{r-1}}(\mathbf{x}))\| \\
&\leq \left(\prod_{j=r+1}^d \|W_j\| \right) \cdot \|W_r - \tilde{W}_r\| \cdot \left(\prod_{j=1}^{r-1} \|W_j\| \right) \cdot \|\mathbf{x}\|,
\end{aligned}$$

from which the result follows after a simplification. The cases $r = 1$ and $r = d$ are handled in exactly the same manner. \square

Lemma 6. *Suppose that $N_{W_1^d}$ is such that $\prod_{j=1}^d \|W_j\| \geq \Gamma$ and $\prod_{j=1}^d \|W_j\|_p \leq M$. Then for any $r \in \{1, \dots, d\}$,*

$$\min_{j \in \{1, \dots, r\}} \frac{\|W_j\|_p}{\|W_j\|} \leq \left(\frac{M}{\Gamma} \right)^{1/r}.$$

Proof. Fixing some r , and using the stated assumptions as well as the fact that $\|W\|_p \geq \|W\|$ for any p , we have

$$\frac{M}{\Gamma} \geq \frac{\prod_{j=1}^d \|W_j\|_p}{\prod_{j=1}^d \|W_j\|} = \prod_{j=1}^d \frac{\|W_j\|_p}{\|W_j\|} \geq \prod_{j=1}^r \frac{\|W_j\|_p}{\|W_j\|} \geq \left(\min_{j \in \{1, \dots, r\}} \frac{\|W_j\|_p}{\|W_j\|} \right)^r.$$

Taking the r -th root from both sides, the result follows. \square

With these lemmas in hand, we can now turn to prove Thm. 3. Combining Lemma 4 and Lemma 5, we have that there indeed exists a network $N_{\tilde{W}_1^d}$ with a rank-1 matrix in layer r , such that

$$\begin{aligned}
\sup_{\mathbf{x} \in \mathcal{X}} \|N_{W_1^d}(\mathbf{x}) - N_{\tilde{W}_1^d}(\mathbf{x})\| &\leq B \left(\prod_{j=1}^d \|W_j\| \right) \cdot \frac{(\|W_r\|_p^p - \|W_r\|^p)^{1/p}}{\|W_r\|} \\
&= B \left(\prod_{j=1}^d \|W_j\| \right) \left(\frac{\|W_r\|_p^p}{\|W_r\|^p} - 1 \right)^{1/p}. \tag{14}
\end{aligned}$$

By the definition of $N_{W_1^d}$, we have that $\|N_{W_1^d}(\mathbf{x})\| \leq \|\mathbf{x}\| \prod_{j=1}^d \|W_j\|$. Since we assume there exists \mathbf{x} (of norm at most B) such that $\|N_{W_1^d}(\mathbf{x})\| \geq \Gamma$, it follows that $\prod_{j=1}^d \|W_j\| \geq \Gamma/B$. Using this and the assumption that $\prod_{j=1}^d \|W_j\|_p \leq M$, and plugging into Lemma 6, it follows that for any $r \in \{1, \dots, d\}$,

$$\min_{j \in \{1, \dots, r\}} \frac{\|W_j\|_p}{\|W_j\|} \leq \left(\frac{MB}{\Gamma} \right)^{1/r}.$$

Substituting into Eq. (14), we get that

$$\begin{aligned} \sup_{\mathbf{x} \in \mathcal{X}} \|N_{W_1^d}(\mathbf{x}) - N_{\tilde{W}_1^d}(\mathbf{x})\| &\leq B \left(\prod_{j=1}^d \|W_j\| \right) \left(\left(\frac{MB}{\Gamma} \right)^{p/r} - 1 \right)^{1/p} \\ &= B \left(\prod_{j=1}^d \|W_j\| \right) \left(\exp \left(\frac{p \log(MB/\Gamma)}{r} \right) - 1 \right)^{1/p}. \end{aligned}$$

Suppose for now that r is such that $r \geq p \log(MB/\Gamma)$. Using the fact that $\exp(z) \leq 1 + 2z$ for any $z \in [0, 1]$, it follows that the above is at most

$$B \left(\prod_{j=1}^d \|W_j\| \right) \left(\frac{2p \log(MB/\Gamma)}{r} \right)^{1/p}.$$

It remains to consider the case where $r < p \log(MB/\Gamma)$. However, in this regime the theorem trivially holds: Let \tilde{W}_r in the network $N_{\tilde{W}_1^d}$ be the all-zeros matrix (which is rank zero and ensures that $N_{\tilde{W}_1^d}(\mathbf{x}) = \mathbf{0}$ for all \mathbf{x}), and we have by definition that

$$\sup_{\mathbf{x} \in \mathcal{X}} \|N_{W_1^d}(\mathbf{x}) - N_{\tilde{W}_1^d}(\mathbf{x})\| = \sup_{\mathbf{x}: \|\mathbf{x}\| \leq B} \|N_{W_1^d}(\mathbf{x})\| \leq B \prod_{j=1}^d \|W_j\| < B \left(\prod_{j=1}^d \|W_j\| \right) \left(\frac{2p \log(MB/\Gamma)}{r} \right)^{1/p}.$$

7.3 Proof of Thm. 4

To prove the theorem, we use a straightforward covering number argument, beginning with a few definitions.

Given any function class \mathcal{F} , a metric d on the elements of \mathcal{F} , and $\epsilon > 0$, we let the covering number $\mathcal{N}(\mathcal{F}, d, \epsilon)$ denote the minimal number n of functions f_1, f_2, \dots, f_n in \mathcal{F} , such that for all $f \in \mathcal{F}$, $\min_{i=1, \dots, n} d(f_i, f) \leq \epsilon$. In particular, fix some set of data points $\mathbf{x}_1, \dots, \mathbf{x}_m$, and define the empirical L_2 distance

$$\hat{d}_m(f, f') = \sqrt{\frac{1}{m} \sum_{i=1}^m (f(\mathbf{x}_i) - f'(\mathbf{x}_i))^2}.$$

Also, given a function class \mathcal{F} and $\mathbf{x}_1, \dots, \mathbf{x}_m$, we let

$$\hat{\mathcal{G}}_m(\mathcal{F}) := \mathbb{E}_\eta \left[\sup_{f \in \mathcal{F}} \frac{1}{m} \sum_{i=1}^m \eta_i f(\mathbf{x}_i) \right]$$

denote the (empirical) Gaussian complexity of \mathcal{F} , where η_1, \dots, η_m are i.i.d. standard Gaussian random variables. It is well known that $\hat{\mathcal{R}}_m(\mathcal{H})$ and $\hat{\mathcal{G}}_m(\mathcal{H})$ are equivalent up to a $c\sqrt{\log(m)}$ factor [Ledoux and Talagrand, 1991, pg. 97]. By Sudakov's minoration theorem (see theorem 3.18 in Ledoux and Talagrand [1991]), we have that for all $\alpha > 0$

$$\log(\mathcal{N}(\mathcal{H}, \hat{d}_m, \alpha)) \leq c \left(\frac{\sqrt{m} \cdot \hat{\mathcal{G}}_m(\mathcal{H})}{\alpha} \right)^2 \quad (15)$$

for a universal constant $c > 0$.

With these definitions in hand, we turn to prove the theorem. We first note that $\mathcal{F}_L \circ \mathcal{H}$ is equivalent to the class $\{Lg(\cdot) + a : g \in \mathcal{F}_{1,0} \circ \mathcal{H}\}$, and therefore

$$\begin{aligned} \hat{\mathcal{R}}_m(\mathcal{F}_{L,a} \circ \mathcal{H}) &= \mathbb{E}_\epsilon \left[\sup_{g \in \mathcal{F}_{L,a} \circ \mathcal{H}} \frac{1}{m} \sum_{i=1}^m \epsilon_i g(\mathbf{x}_i) \right] = \mathbb{E}_\epsilon \left[\sup_{g \in \mathcal{F}_{1,0} \circ \mathcal{H}} \frac{1}{m} \sum_{i=1}^m \epsilon_i (Lg(\mathbf{x}_i) + a) \right] \\ &= \mathbb{E}_\epsilon \left[L \cdot \sup_{g \in \mathcal{F}_{1,0} \circ \mathcal{H}} \left(\frac{1}{m} \sum_{i=1}^m \epsilon_i g(\mathbf{x}_i) \right) + \frac{a}{m} \sum_{i=1}^m \epsilon_i \right] \\ &= L \cdot \mathbb{E}_\epsilon \left[\sup_{g \in \mathcal{F}_{1,0} \circ \mathcal{H}} \frac{1}{m} \sum_{i=1}^m \epsilon_i g(\mathbf{x}_i) \right] = L \cdot \hat{\mathcal{R}}_m(\mathcal{F}_{1,0} \circ \mathcal{H}). \end{aligned} \quad (16)$$

Therefore, it is enough to consider $\hat{\mathcal{R}}_m(\mathcal{F}_{1,0} \circ \mathcal{H})$. To simplify notation in what follows, we will drop the 1, 0 subscript from \mathcal{F} .

We first argue that

$$\log(\mathcal{N}(\mathcal{F}, \hat{d}_m, \epsilon)) \leq c' \left(1 + \frac{R}{\epsilon} \right), \quad (17)$$

again for some numerical constant $c' > 0$. To prove this, we note that for any functions f, f' , it holds that $\hat{d}_m(f, f') \geq \hat{d}_\infty(f, f') := \sup_{\mathbf{x}} |f(\mathbf{x}) - f'(\mathbf{x})|$, so it is enough to upper bound $\mathcal{N}(\mathcal{H}, \hat{d}_\infty, \epsilon)$. To do so, we first notice that the range of any $f \in \mathcal{F}$ is in $[-R, R]$. Discretize $[-R, R] \times [-R, R]$ into a two-dimensional grid $\mathcal{U}_x \times \mathcal{U}_y$, where

$$\mathcal{U}_x := \{-R, -R + \epsilon, -R + 2\epsilon, \dots, -R + \lfloor 2R/\epsilon \rfloor \epsilon, R\} \quad , \quad \mathcal{U}_y := \{0, \pm\epsilon, \pm 2\epsilon, \dots, \pm \lfloor R/\epsilon \rfloor \epsilon, \pm R\}.$$

Given any $f \in \mathcal{F}$, construct the piecewise-linear function f' as follows: For any input $x \in \mathcal{U}_x$, let $f'(x)$ be the point in \mathcal{U}_y nearest to $f(x)$ (breaking ties arbitrarily), and let the rest of f' be constructed as a linear interpolation of these points on \mathcal{U}_x . It is easily verified that $\sup_{x \in [-R, R]} |f(x) - f'(x)| \leq \epsilon$. Moreover, note that for two neighboring points x, x' in \mathcal{U}_x , the points $f'(x), f'(x')$ on \mathcal{U}_y must be neighboring or the same. Therefore, each such function f' can be parameterized by a vector of the form $\{-, 0, +\}^{|\mathcal{U}_x|-1}$, which specifies whether (starting from the origin) f' goes up, down, or remains the same on each of its linear segments. The number of such functions is at most $3^{|\mathcal{U}_x|-1} \leq 3^{2R/\epsilon+1}$, and therefore $\mathcal{N}(\mathcal{F}, \hat{d}_\infty, \epsilon) \leq 3^{2R/\epsilon+1}$. Recalling that $\hat{d}_m(f, f')$ majorizes $\hat{d}_\infty(f, f')$, we get Eq. (17).

Next, we argue that

$$\mathcal{N}(\mathcal{F} \circ \mathcal{H}, \hat{d}_m, \epsilon) \leq \mathcal{N}\left(\mathcal{F}, \hat{d}_m, \frac{\epsilon}{2}\right) \cdot \mathcal{N}\left(\mathcal{H}, \hat{d}_m, \frac{\epsilon}{2}\right). \quad (18)$$

To see this, pick any $f \in \mathcal{F}$ and $h \in \mathcal{H}$, and let f', h' be the respective closest functions in the cover of \mathcal{F} and \mathcal{H} (at scale $\epsilon/2$). By the triangle inequality and the easily verified fact that f' is 1-Lipschitz, we have

$$\begin{aligned} \hat{d}_m(fh, f'h') &\leq \hat{d}_m(fh, f'h) + \hat{d}_m(f'h, f'h') \leq \hat{d}_\infty(fh, f'h) + \hat{d}_m(f'h, f'h') \\ &\leq \sup_{x \in [-R, R]} |f(x) - f'(x)| + \sqrt{\frac{1}{m} \sum_{i=1}^m (f'h(\mathbf{x}_i) - f'h'(\mathbf{x}_i))^2} \\ &\leq \frac{\epsilon}{2} + \sqrt{\frac{1}{m} \sum_{i=1}^m (h(\mathbf{x}_i) - h'(\mathbf{x}_i))^2} \\ &= \frac{\epsilon}{2} + \hat{d}_m(h, h') \leq \frac{\epsilon}{2} + \frac{\epsilon}{2} = \epsilon. \end{aligned}$$

Therefore, we can cover $\mathcal{F} \circ \mathcal{H}$ (at scale ϵ) by taking $f'(h'(\cdot))$ for all possible choices of f', h' from the covers of \mathcal{F}, \mathcal{H} (at scale $\epsilon/2$), leading to Eq. (18).

Combining Eq. (15), Eq. (17) and Eq. (18), we get that

$$\log(\mathcal{N}(\mathcal{F} \circ \mathcal{H}, \hat{d}_m, \epsilon)) \leq c'' \left(1 + \frac{R}{\epsilon} + \left(\frac{\sqrt{m} \cdot \hat{\mathcal{G}}_m(\mathcal{H})}{\epsilon} \right)^2 \right),$$

for some numerical constant $c'' > 0$. Finally, we use Dudley's entropy integral, which together with the equation above, implies the following for some numerical constant $c > 0$ (possibly changing from row to row):

$$\begin{aligned} \hat{\mathcal{R}}_m(\mathcal{F} \circ \mathcal{H}) &\leq c \inf_{\alpha \geq 0} \left\{ \alpha + \frac{1}{\sqrt{m}} \int_{\alpha}^{\sup_{g \in \mathcal{F} \circ \mathcal{H}} \hat{d}_m(g, 0)} \sqrt{\log(\mathcal{N}(\mathcal{F} \circ \mathcal{H}, \hat{d}_m, \epsilon))} d\epsilon \right\} \\ &\leq c \inf_{\alpha \geq 0} \left\{ \alpha + \frac{1}{\sqrt{m}} \int_{\alpha}^R \sqrt{\log(\mathcal{N}(\mathcal{F} \circ \mathcal{H}, \hat{d}_m, \epsilon))} d\epsilon \right\} \\ &\leq c \inf_{\alpha \geq 0} \left\{ \alpha + \frac{1}{\sqrt{m}} \int_{\alpha}^R \sqrt{1 + \frac{R}{\epsilon} + \left(\frac{\sqrt{m} \cdot \hat{\mathcal{G}}_m(\mathcal{H})}{\epsilon} \right)^2} d\epsilon \right\} \\ &\leq c \inf_{\alpha \geq 0} \left\{ \alpha + \frac{1}{\sqrt{m}} \int_{\alpha}^R \left(1 + \sqrt{\frac{R}{\epsilon}} + \left| \frac{\sqrt{m} \cdot \hat{\mathcal{G}}_m(\mathcal{H})}{\epsilon} \right| \right) d\epsilon \right\} \\ &\leq c \inf_{\alpha \geq 0} \left\{ \alpha + \frac{1}{\sqrt{m}} \int_0^R \left(1 + \sqrt{\frac{R}{\epsilon}} \right) d\epsilon + \hat{\mathcal{G}}_m(\mathcal{H}) \int_{\alpha}^R \frac{1}{\epsilon} d\epsilon \right\} \\ &\leq c \inf_{\alpha \geq 0} \left\{ \alpha + \frac{R}{\sqrt{m}} + \hat{\mathcal{G}}_m(\mathcal{H}) \log \left(\frac{R}{\alpha} \right) \right\}. \end{aligned}$$

Choosing in particular $\alpha = R/\sqrt{m}$, we get the upper bound

$$\hat{\mathcal{R}}_m(\mathcal{F} \circ \mathcal{H}) \leq c \left(\frac{R}{\sqrt{m}} + \log(m) \cdot \hat{\mathcal{G}}_m(\mathcal{H}) \right)$$

for some $c > 0$. Plugging this into Eq. (16), and upper bounding $\hat{\mathcal{G}}_m(\mathcal{H})$ by $c\sqrt{\log(m)}\tilde{\mathcal{R}}_m(\mathcal{H})$ (see [Ledoux and Talagrand, 1991]), the result follows.

7.4 Proof of Thm. 5

It is enough to prove the bound for any fixed $r \in \{1, \dots, d\}$, and then take the infimum over any such r .

Given \mathcal{H} and r , construct a new hypothesis class $\tilde{\mathcal{H}}$ by replacing each network $h \in \mathcal{H}$ by the network \tilde{h} as defined in Thm. 3 (namely, where the parameter matrix in the r -th layer is replaced by a rank-1 matrix). We will use the notation \tilde{h}_h to clarify the dependence of \tilde{h} on h . According to that theorem, as well as the

definition of Rademacher complexity, we have

$$\begin{aligned}
\hat{\mathcal{R}}_m(\ell \circ \mathcal{H}) &= \mathbb{E}_\epsilon \left[\sup_{h \in \mathcal{H}} \frac{1}{m} \sum_{i=1}^m \epsilon_i \ell(h(\mathbf{x}_i)) \right] \\
&= \mathbb{E}_\epsilon \left[\sup_{h \in \mathcal{H}} \left(\frac{1}{m} \sum_{i=1}^m \epsilon_i \ell(\tilde{h}_h(\mathbf{x}_i)) + \frac{1}{m} \sum_{i=1}^m \epsilon_i (\ell(h(\mathbf{x}_i)) - \ell(\tilde{h}_h(\mathbf{x}_i))) \right) \right] \\
&\leq \mathbb{E}_\epsilon \left[\sup_{h \in \mathcal{H}} \left(\frac{1}{m} \sum_{i=1}^m \epsilon_i \ell(\tilde{h}_h(\mathbf{x}_i)) + \sup_{\mathbf{x} \in \mathcal{X}} |\ell(h(\mathbf{x})) - \ell(\tilde{h}_h(\mathbf{x}))| \right) \right] \\
&\leq \mathbb{E}_\epsilon \left[\sup_{\tilde{h} \in \tilde{\mathcal{H}}} \left(\frac{1}{m} \sum_{i=1}^m \epsilon_i \ell(\tilde{h}(\mathbf{x}_i)) + \frac{1}{\gamma} \cdot \sup_{\mathbf{x} \in \mathcal{X}} |h(\mathbf{x}) - \tilde{h}_h(\mathbf{x})| \right) \right] \\
&\leq \hat{\mathcal{R}}_m(\ell \circ \tilde{H}) + \frac{B \left(\prod_{j=1}^d M(j) \right)}{\gamma} \left(\frac{2p \log \left(\frac{B}{\Gamma} \prod_{j=1}^d M_p(j) \right)}{r} \right)^{1/p}. \tag{19}
\end{aligned}$$

We now reach the crucial observation which lies at the heart of the proof. Consider any network $N_{\tilde{W}_1^d}$ in $\tilde{\mathcal{H}}$, and let $s\mathbf{u}\mathbf{v}^\top$ be the SVD decomposition of its (rank-1) parameter matrix \tilde{W}_r in layer r (where $s, \mathbf{u}, \mathbf{v}$ are also leading singular value and vectors of W_r , by construction). By definition, we have that the composition of $N_{\tilde{W}_1^d}$ with any loss ℓ_j equals

$$\mathbf{x} \mapsto \ell_j(W_d \sigma_{d-1}(W_{d-1} \sigma_{d-2}(\dots \sigma_r(s\mathbf{u}\mathbf{v}^\top \sigma_{r-1}(\dots \sigma_1(W_1 \mathbf{x}))))).$$

This function is equivalent to the composition of the function

$$\mathbf{x} \mapsto s\mathbf{v}^\top \sigma_{r-1}(\dots \sigma_1(W_1 \mathbf{x}))$$

with the *univariate* function

$$x \mapsto \ell_j(W_d \sigma_{d-1}(W_{d-1} \sigma_{d-2}(\dots \sigma_r(\mathbf{u}x)))) .$$

Note that since $\|\mathbf{sv}^\top\| = s = \|W_r\|$ and $\|\mathbf{sv}^\top\|_p = s \leq \|W_r\|_p$, the former function is contained in \mathcal{H}_r as defined in the theorem; whereas the latter function has Lipschitz constant at most $\frac{1}{\gamma} \prod_{j=r+1}^d \|W_j\| \leq \frac{1}{\gamma} \prod_{j=r+1}^d M(j)$, and maps the input 0 to the same fixed output (which we'll denote as a) for all j . Therefore, we obtain that $\ell \circ \tilde{\mathcal{H}}$ is contained in the composition of \mathcal{H}_r (with functions whose output is bounded in $\pm B \prod_{j=1}^r M(j)$), and the class $\mathcal{F}_{\frac{1}{\gamma} \prod_{j=r+1}^d M(j), a}$ consisting of $\prod_{j=r+1}^d M(j)$ -Lipschitz functions. As a result, we can apply Thm. 4 and obtain that

$$\begin{aligned}
\hat{\mathcal{R}}_m(\ell \circ \hat{H}) &\leq \hat{\mathcal{R}}_m \left(\mathcal{H}_r \circ \mathcal{F}_{\frac{1}{\gamma} \prod_{j=r+1}^d M(j), a} \right) \\
&\leq \frac{c}{\gamma} \left(\prod_{j=r+1}^d M(j) \right) \left(\frac{B}{\sqrt{m}} \prod_{j=1}^r M(j) + \log^{3/2}(m) \cdot \hat{\mathcal{R}}_m(\mathcal{H}_r) \right) \\
&= \frac{c}{\gamma} \left(\prod_{j=1}^d M(j) \right) \left(\frac{B}{\sqrt{m}} + \frac{\log^{3/2}(m) \cdot \hat{\mathcal{R}}_m(\mathcal{H}_r)}{\prod_{j=1}^r M(j)} \right).
\end{aligned}$$

Plugging this back into Eq. (19) and simplifying a bit (also noting that $p^{1/p}$ can be upper bounded by a universal constant), we get that $\hat{\mathcal{R}}_m(\ell \circ \mathcal{H})$ is upper bounded by

$$\frac{c' B \prod_{j=1}^d M(j)}{\gamma} \left(\frac{\log^{3/2}(m) \cdot \hat{\mathcal{R}}_m(\mathcal{H}_r)}{B \prod_{j=1}^d M(j)} + \left(\frac{\log \left(\frac{B}{\Gamma} \prod_{j=1}^d M_p(j) \right)}{r} \right)^{1/p} + \frac{1}{\sqrt{m}} \right)$$

for an appropriate constant c' . As mentioned at the beginning of the proof, this upper bound holds for any fixed $r \in \{1, \dots, d\}$, from which the result follows.

7.5 Proof of Lemma 3

We will show that for any α, β, b, c, n as stated in the lemma, there always exists a choice of $r \in \{1, \dots, d\}$ such that

$$\min \left\{ \frac{cr^\alpha}{n} + \frac{b}{r^\beta}, \frac{d^\alpha}{n} \right\} \leq 3 \cdot \frac{b^{\frac{\alpha}{\alpha+\beta}}}{(n/c)^{\frac{\beta}{\alpha+\beta}}}.$$

Since the left hand side is also trivially at most $\frac{d^\alpha}{n}$, the result follows. We prove this inequality by a case analysis:

- If $(bn/c)^{\frac{1}{\alpha+\beta}} \in [1, d]$, pick $r = \lfloor (bn/c)^{\frac{1}{\alpha+\beta}} \rfloor \in \{1, 2, \dots, d\}$, in which case

$$\frac{cr^\alpha}{n} + \frac{b}{r^\beta} \leq \frac{c \left((bn/c)^{\frac{1}{\alpha+\beta}} \right)^\alpha}{n} + \frac{b}{\left(\frac{1}{2} (bn/c)^{\frac{1}{\alpha+\beta}} \right)^\beta} = \frac{(b)^{\frac{\alpha}{\alpha+\beta}}}{(n/c)^{\frac{\beta}{\alpha+\beta}}} + 2^\beta \frac{b^{\frac{\alpha}{\alpha+\beta}}}{(n/c)^{\frac{\beta}{\alpha+\beta}}} \leq 3 \frac{b^{\frac{\alpha}{\alpha+\beta}}}{(n/c)^{\frac{\beta}{\alpha+\beta}}}.$$

- If $(bn/c)^{\frac{1}{\alpha+\beta}} > d$, it follows that

$$\min \left\{ \frac{cr^\alpha}{n} + \frac{b}{r^\beta}, \frac{d^\alpha}{n} \right\} \leq \frac{d^\alpha}{n} < c \frac{(bn/c)^{\frac{\alpha}{\alpha+\beta}}}{n} = \frac{b^{\frac{\alpha}{\alpha+\beta}}}{(n/c)^{\frac{\beta}{\alpha+\beta}}}.$$

7.6 Proof of Corollary 1

A direct application of Thm. 1, as well as the fact that the loss ℓ is $1/\gamma$ Lipschitz, implies that

$$\hat{\mathcal{R}}_m(\ell \circ \mathcal{H}) \leq \mathcal{O} \left(\frac{B \prod_{j=1}^d M_F(j)}{\gamma} \sqrt{\frac{d}{m}} \right). \quad (20)$$

On the other hand, plugging Eq. (12) into Thm. 1, (using $p = 2$, letting each \mathcal{W}_j be the space of all matrices, and choosing $M(j) = M_F(j)$ for all j , noting that $\frac{\|W_j\|}{M(j)} \leq \frac{\|W_j\|_F}{M_F(j)} \leq 1$), we get that

$$\hat{\mathcal{R}}_m(\ell \circ \mathcal{H}) \leq \mathcal{O} \left(\frac{B \prod_{j=1}^d M_F(j)}{\gamma} \left(\min_{r \in \{1, \dots, d\}} \left\{ \frac{\log^{3/2}(m) \sqrt{r}}{\sqrt{m}} + \sqrt{\frac{\log \left(\frac{B}{\Gamma} \prod_{j=1}^d M_F(j) \right)}{r}} \right\} \right) \right). \quad (21)$$

Upper bounding $\hat{\mathcal{R}}_m(\ell \circ \mathcal{H})$ by the minimum of Eq. (20) and Eq. (21), and using Lemma 3 with $\alpha = \beta = \frac{1}{2}$, $b = \sqrt{\log \left(\frac{B}{\Gamma} \prod_{j=1}^d M_F(j) \right)}$, $c = \log^{3/2}(m)$, and $n = \sqrt{m}$, the result follows.

7.7 Proof of Corollary 2

Consider the class

$$\mathcal{H}_r = \left\{ N_{W_1^T} : \begin{array}{l} N_{W_1^T} \text{ maps to } \mathbb{R} \\ \forall j \in \{1 \dots r-1\}, \frac{\|W_j^T\|_{2,1}}{\|W_j\|} \leq L \\ \forall j \in \{1 \dots r\}, \max \left\{ \frac{\|W_j\|}{M(j)}, \frac{\|W_j\|_p}{M_p(j)} \right\} \leq 1 \end{array} \right\}.$$

Since $N_{W_1^T}$ in the definition of \mathcal{H}_r maps to \mathbb{R} , W_r is a vector, meaning that $\|W_r^T\|_{2,1} = \|W_r\|_2 = \|W_r\|$. Therefore, we can use Thm. 6 to bound $\hat{\mathcal{R}}_m(\mathcal{H}_r)$; in particular, for $r \geq 1$,

$$\hat{\mathcal{R}}_m(\mathcal{H}_r) \leq \mathcal{O} \left(\frac{B \log(h) \log(m) L \prod_{j=1}^r M(j) r^{3/2}}{\sqrt{m}} \right).$$

It therefore follows from Thm. 5 that $\hat{\mathcal{R}}_m(\ell \circ \mathcal{H})$ is at most

$$\mathcal{O} \left(\frac{BL \log(h) \log(m) \prod_{j=1}^d M(j)}{\gamma} \cdot \min_{r \in \{1, \dots, d\}} \left\{ \frac{cr^{3/2}}{n} + \frac{b}{r^{1/p}} \right\} \right),$$

where $c = \log^{3/2}(m)$, $n = \sqrt{m}$, and $b = \log \left(\frac{B}{\Gamma} \prod_{j=1}^d M_p(j) \right)^{1/p}$.

On the other hand, a direct application of Thm. 6 also implies that $\hat{\mathcal{R}}_m(\ell \circ \mathcal{H})$ is at most

$$\mathcal{O} \left(\frac{B \prod_{j=1}^d M(j)}{\gamma} \cdot \frac{\log(h) \log(m) L d^{3/2}}{\sqrt{m}} \right).$$

Combining the above two bounds, and applying Lemma 3, we get that $\hat{\mathcal{R}}_m(\ell \circ \mathcal{H})$ is at most

$$\mathcal{O} \left(\frac{BL \log(h) \log(m) \prod_{j=1}^d M(j)}{\gamma} \cdot \min \left\{ \frac{\log \left(\frac{B}{\Gamma} \prod_{j=1}^d M_p(j) \right)^{\frac{1}{\frac{2}{3}+p}} \left(\log^{3/2}(m) \right)^{1+\frac{1}{\frac{2}{3}+p}}}{m^{\frac{1}{2+3p}}}, \frac{d^{3/2}}{\sqrt{m}} \right\} \right).$$

7.8 Proof of Thm. 7

By definition of the Rademacher complexity, it is enough to lower bound the complexity of $\ell \circ \mathcal{H}'$ where \mathcal{H}' is some subset of \mathcal{H} . In particular, consider the class \mathcal{H}' of neural networks over \mathbb{R}^h of the form

$$\mathbf{x} \mapsto M_p(d) \cdot M_p(d-1) \cdots M_p(2) \cdot W \mathbf{x},$$

where $W = \text{diag}(\|\mathbf{w}\|)$ is a diagonal matrix satisfying $\|\mathbf{w}\|_p \leq M_p(1)$ (here, $\|\cdot\|_p$ refers to the vector p -norm). Furthermore, suppose that $\ell(\mathbf{z}) = \frac{1}{\gamma} \max\{z_1, \dots, z_h\}$. Finally, we will choose $\mathbf{x}_1, \dots, \mathbf{x}_m$ in \mathbb{R}^h as $\mathbf{x}_i = B \mathbf{e}_{(i \bmod h)}$ for all i , where \mathbf{e}_t is the t -th standard basis vector. Letting $A_k = \{i \in \{1, \dots, m\} :$

$i \bmod h = k\}$, it holds that

$$\begin{aligned}\hat{\mathcal{R}}_m(\ell \circ \mathcal{H}') &= \mathbb{E}_\epsilon \sup_{W: \|W\|_p \leq M_p(1)} \frac{1}{m} \sum_{i=1}^m \epsilon_i \ell \left(\prod_{j=2}^d M_p(j) W \mathbf{x}_i \right) \\ &= \frac{B \prod_{j=2}^d M_p(j)}{\gamma m} \cdot \mathbb{E}_\epsilon \sup_{\mathbf{w}: \|\mathbf{w}\|_p \leq M_p(1)} \sum_{k=1}^h \max\{0, w_k\} \cdot \sum_{i \in A_k} \epsilon_i \\ &= \frac{B \prod_{j=1}^d M_p(j)}{\gamma m} \cdot \mathbb{E}_\epsilon \sup_{\mathbf{w}: \|\mathbf{w}\|_p \leq 1} \sum_{k=1}^h \max\{0, w_k\} \cdot \sum_{i \in A_k} \epsilon_i.\end{aligned}$$

In particular, by choosing $w_k = h^{-1/p} \cdot \text{sign} \left(\sum_{i \in A_k} \epsilon_i \right)$ for all k , we can lower bound the above by

$$\frac{B \prod_{j=1}^d M_p(j)}{\gamma m \cdot h^{1/p}} \cdot \mathbb{E}_\epsilon \left[\sum_{k=1}^h \max \left\{ 0, \sum_{i \in A_k} \epsilon_i \right\} \right] = \Omega \left(\frac{B \prod_{j=1}^d M_p(j)}{\gamma m \cdot h^{1/p}} \cdot \sum_{k=1}^h \sqrt{|A_k|} \right),$$

and since $|A_k| \geq \lfloor m/h \rfloor$ by its definition, we get a lower bound of

$$\Omega \left(\frac{B \prod_{j=1}^d M_p(j)}{\gamma m \cdot h^{1/p}} \cdot h \sqrt{\frac{m}{h}} \right) = \Omega \left(\frac{B \prod_{j=1}^d M_p(j) \cdot h^{\frac{1}{2} - \frac{1}{p}}}{\gamma \sqrt{m}} \right). \quad (22)$$

An alternative bound (which is better when $p < 2$) can be obtained by considering the class \mathcal{H}' of real-valued neural networks over \mathbb{R} , of the form

$$x \mapsto M_p(d) \cdot M_p(d-1) \cdots M_p(2) \cdot wx,$$

where $|w| \leq M_p(1)$. Furthermore, supposing that ℓ is the identity function times $\frac{1}{\gamma}$, and $\mathbf{x}_i = B$ for all i , it holds that

$$\begin{aligned}\hat{\mathcal{R}}_m(\ell \circ \mathcal{H}') &= \mathbb{E}_\epsilon \sup_{w: |w| \leq M_p(1)} \frac{1}{\gamma m} \sum_{i=1}^m \epsilon_i \prod_{j=2}^d M_p(j) B w \\ &= \frac{B \prod_{j=1}^d M_p(j)}{\gamma m} \cdot \mathbb{E}_\epsilon \left| \sum_{i=1}^m \epsilon_i \right| = \Omega \left(\frac{B \prod_{j=1}^d M_p(j)}{\sqrt{m}} \right).\end{aligned}$$

Taking the best of this lower bound and the lower bound in Eq. (22), the result follows.

References

Martin Anthony and Peter L Bartlett. *Neural network learning: Theoretical foundations*. cambridge university press, 2009.

Peter Bartlett, Dylan J Foster, and Matus Telgarsky. Spectrally-normalized margin bounds for neural networks. *arXiv preprint arXiv:1706.08498*, 2017.

- Peter L Bartlett and Shahar Mendelson. Rademacher and gaussian complexities: Risk bounds and structural results. *Journal of Machine Learning Research*, 3(Nov):463–482, 2002.
- Stéphane Boucheron, Gábor Lugosi, and Pascal Massart. *Concentration inequalities: A nonasymptotic theory of independence*. Oxford university press, 2013.
- Vladimir Koltchinskii and Dmitry Panchenko. Empirical margin distributions and bounding the generalization error of combined classifiers. *Annals of Statistics*, pages 1–50, 2002.
- Michel Ledoux and Michel Talagrand. *Probability in Banach Spaces*. Springer, 1991.
- Behnam Neyshabur, Ryota Tomioka, and Nathan Srebro. In search of the real inductive bias: On the role of implicit regularization in deep learning. *arXiv preprint arXiv:1412.6614*, 2014.
- Behnam Neyshabur, Ryota Tomioka, and Nathan Srebro. Norm-based capacity control in neural networks. In *Conference on Learning Theory*, pages 1376–1401, 2015.
- Behnam Neyshabur, Srinadh Bhojanapalli, David McAllester, and Nathan Srebro. A pac-bayesian approach to spectrally-normalized margin bounds for neural networks. *arXiv preprint arXiv:1707.09564*, 2017.
- Shai Shalev-Shwartz and Shai Ben-David. *Understanding machine learning: From theory to algorithms*. Cambridge university press, 2014.
- Chiyuan Zhang, Samy Bengio, Moritz Hardt, Benjamin Recht, and Oriol Vinyals. Understanding deep learning requires rethinking generalization. *arXiv preprint arXiv:1611.03530*, 2016.