
Generalization Bounds for Neural Networks via Approximate Description Length

Anonymous Author(s)

Affiliation

Address

email

Abstract

1 We investigate the sample complexity of networks with bounds on the magnitude
2 of its weights. In particular, we consider the class

$$\mathcal{N} = \{W_t \circ \rho \circ W_{t-1} \circ \rho \dots \circ \rho \circ W_1 : W_1, \dots, W_{t-1} \in M_{d \times d}, W_t \in M_{1,d}\}$$

3 where the spectral norm of each W_i is bounded by $O(1)$, the Frobenius norm
4 is bounded by R , and ρ is the sigmoid function $\frac{e^x}{1+e^x}$ or the smoothened ReLU
5 function $\ln(1+e^x)$. We show that for any depth t , if the inputs are in $[-1, 1]^d$,
6 the sample complexity of \mathcal{N} is $\tilde{O}\left(\frac{dR^2}{\epsilon^2}\right)$. This bound is optimal up to log-factors,

7 and substantially improves over the previous state of the art of $\tilde{O}\left(\frac{d^2 R^2}{\epsilon^2}\right)$, that was
8 established in a recent line of work [9, 4, 7, 5, 2, 8].

9 We furthermore show that this bound remains valid if instead of considering the
10 magnitude of the W_i 's, we consider the magnitude of $W_i - W_i^0$, where W_i^0 are
11 some reference matrices, with spectral norm of $O(1)$. By taking the W_i^0 to be the
12 matrices at the onset of the training process, we get sample complexity bounds that
13 are sub-linear in the number of parameters, in many *typical* regimes of parameters.

14 To establish our results we develop a new technique to analyze the sample complex-
15 ity of families \mathcal{H} of predictors. We start by defining a new notion of a randomized
16 approximate description of functions $f : \mathcal{X} \rightarrow \mathbb{R}^d$. We then show that if there is a
17 way to approximately describe functions in a class \mathcal{H} using d bits, then $\frac{d}{\epsilon^2}$ examples
18 suffices to guarantee uniform convergence. Namely, that the empirical loss of all
19 the functions in the class is ϵ -close to the true loss. Finally, we develop a set of
20 tools for calculating the approximate description length of classes of functions
21 that can be presented as a composition of linear function classes and non-linear
22 functions.

23 1 Introduction

24 We analyze the sample complexity of networks with bounds on the magnitude of their weights. Let
25 us consider a prototypical case, where the input space is $\mathcal{X} = [-1, 1]^d$, the output space is \mathbb{R} , the
26 number of layers is t , all hidden layers has d neurons, and the activation function is $\rho : \mathbb{R} \rightarrow \mathbb{R}$. The
27 class of functions computed by such an architecture is

$$\mathcal{N} = \{W_t \circ \rho \circ W_{t-1} \circ \rho \dots \circ \rho \circ W_1 : W_1, \dots, W_{t-1} \in M_{d \times d}, W_t \in M_{1,d}\}$$

28 As the class \mathcal{N} is defined by $(t-1)d^2 + d = O(d^2)$ parameters, classical results (e.g. [1]) tell
29 us that order of d^2 examples are sufficient and necessary in order to learn a function from \mathcal{N} (in a
30 standard worst case analysis). However, modern networks often succeed to learn with substantially

less examples. One way to provide alternative results, and a potential explanation to the phenomena, is to take into account the magnitude of the weights. This approach was a success story in the days of SVM [3] and Boosting [10], provided a nice explanation to generalization with sub-linear (in the number of parameters) number of examples, and was even the deriving force behind algorithmic progress. It seems just natural to adopt this approach in the context of modern networks. For instance, it is natural to consider the class

$$\mathcal{N}_R = \{W_t \circ \rho \circ W_{t-1} \circ \rho \dots \circ \rho \circ W_1 : \forall i, \|W_i\|_F \leq R, \|W_i\| \leq O(1)\}$$

where $\|W\| = \max_{\|x\|=1} \|Wx\|$ is the spectral norm and $\|W\|_F = \sqrt{\sum_{i,j=1}^d W_{ij}^2}$ is the Frobenius norm. This class has been analyzed in several recent works [9, 4, 7, 5, 2, 8]. Best known results show a sample complexity of $\tilde{O}\left(\frac{d^2 R^2}{\epsilon^2}\right)$ (for the sake of simplicity, in the introduction, we ignore the dependence on the depth in the big-O notation). In this paper we prove, for various activations, a stronger bound of $\tilde{O}\left(\frac{dR^2}{\epsilon^2}\right)$, which is optimal, up to log factors, for constant depth networks.

How good is this bound? Does it finally provide sub-linear bound in typical regimes of the parameters? To answer this question, we need to ask how large R is. While this question of course don't have a definite answer, empirical studies (e.g. [12]) show that it is usually the case that the norm (spectral, Frobenius, and others) of the weight matrices is at the same order of magnitude as the norm of the matrix in the onset of the training process. In most standard training methods, the initial matrices are random matrices with independent (or almost independent) entries, with mean zero and variance of order $\frac{1}{d}$. The Frobenius norm of such a matrix is of order \sqrt{d} . Hence, the magnitude of R is of order \sqrt{d} . Going back to our $\tilde{O}\left(\frac{dR^2}{\epsilon^2}\right)$ bound, we get a sample complexity of $\tilde{O}\left(\frac{d^2}{\epsilon^2}\right)$, which is unfortunately still linear in the number of parameters.

Since our bound is almost optimal, we can ask whether this is the end of the story? Should we abandon the aforementioned approach to network sample complexity? A more refined examination of the training process suggests another hope for this approach. Indeed, the training process doesn't start from the zero matrix, but rather form a random initialization matrix. Thus, it stands to reason that instead of considering the magnitude of the weight matrices W_i , we should consider the magnitude of $W_i - W_i^0$, where W_i^0 is the initial weight matrix. Indeed, empirical studies [6] show that the Frobenius norm of $W_i - W_i^0$ is often order of magnitude smaller than the Frobenius norm of W_i . Following this perspective, it is natural to consider the class

$$\mathcal{N}_R(W_1^0, \dots, W_t^0) = \{W_t \circ \rho \circ W_{t-1} \circ \rho \dots \circ \rho \circ W_1 : \|W_i - W_i^0\| \leq O(1), \|W_i - W_i^0\|_F \leq R\}$$

For some fixed matrices, W_1^0, \dots, W_t^0 of spectral norm¹ $O(1)$. It is natural to expect that considering balls around the initial W_i^0 's instead of zero, shouldn't change the sample complexity of the class at hand. In other words, we can expect that the sample complexity of $\mathcal{N}_R(W_1^0, \dots, W_t^0)$ should be approximately the sample complexity of \mathcal{N}_R . Namely, we expect a sample complexity of $\tilde{O}\left(\frac{dR^2}{\epsilon^2}\right)$. Such a bound would finally be sub-linear, as in practice, it is often the case that $R^2 \ll d$.

This approach was pioneered by [4] who considered the class

$$\mathcal{N}_R^{2,1}(W_1^0, \dots, W_t^0) = \{W_t \circ \rho \circ W_{t-1} \circ \rho \dots \circ \rho \circ W_1 : \|W_i - W_i^0\| \leq O(1), \|W_i - W_i^0\|_{2,1} \leq R\}$$

where $\|W\|_{2,1} = \sum_{i=1}^d \sqrt{\sum_{j=1}^d W_{ij}^2}$. For this class they proved a sample complexity bound of $\tilde{O}\left(\frac{dR^2}{\epsilon^2}\right)$. Since, $\|W\|_{2,1} \leq \sqrt{d}\|W\|_F$, this implies a sample complexity bound of $\tilde{O}\left(\frac{d^2 R^2}{\epsilon^2}\right)$ on $\mathcal{N}_R(W_1^0, \dots, W_t^0)$, which is still not sublinear². In this paper we finally prove a sub-linear sample complexity bound of $\tilde{O}\left(\frac{dR^2}{\epsilon^2}\right)$ on $\mathcal{N}_R(W_1^0, \dots, W_t^0)$.

To prove our results, we develop a new technique for bounding the sample complexity of function classes. Roughly speaking, we define a notion of approximate description of a function, and count

¹The bound of $O(1)$ on the spectral norm of the W_i^0 's and $W_i - W_i^0$ is again motivated by the practice of neural networks – the spectral norm of W_i^0 , with standard initializations, is $O(1)$, and empirical studies [6, 12] show that the spectral norm of $W_i - W_i^0$ is usually very small.

²We note that $\|W\|_{2,1} = \Theta(\sqrt{d})$ even if W is a random matrix with variance that is calibrated so that $\|W\|_F = \Theta(1)$ (namely, each entry has variance $\frac{1}{d^2}$).

how many bits are required in order to give an approximate description for the functions in the class under study. We then show that this number, called the *approximate description length (ADL)*, gives an upper bound on the sample complexity. The advantage of our method over existing techniques is that it behaves nicely with compositions. That is, once we know the approximate description length of a class \mathcal{H} of functions from \mathcal{X} to \mathbb{R}^d , we can also bound the ADL of $\rho \circ \mathcal{H}$, as well as $\mathcal{L} \circ \mathcal{H}$, where \mathcal{L} is a class of linear functions. This allows us to utilize the compositional structure of neural networks.

2 Preliminaries

Notation We denote by $\text{med}(x_1, \dots, x_k)$ the median of $x_1, \dots, x_k \in \mathbb{R}$. For vectors $\mathbf{x}^1, \dots, \mathbf{x}^k \in \mathbb{R}^d$ we denote $\text{med}(\mathbf{x}^1, \dots, \mathbf{x}^k) = (\text{med}(x_1^1, \dots, x_1^k), \dots, \text{med}(x_d^1, \dots, x_d^k))$. We use \log to denote \log_2 , and \ln to denote \log_e . An expression of the form $f(n) \lesssim g(n)$ means that there is a universal constant $c > 0$ for which $f(n) \leq cg(n)$. For a finite set A and $f : A \rightarrow \mathbb{R}$ we let $\mathbb{E}_{x \in A} f = \frac{1}{|A|} \sum_{a \in A} f(a)$. We denote $\mathbb{B}_M^d = \{\mathbf{x} \in \mathbb{R}^d : \|\mathbf{x}\| \leq M\}$ and $\mathbb{B}^d = \mathbb{B}_1^d$. Likewise, we denote $\mathbb{S}^{d-1} = \{\mathbf{x} \in \mathbb{R}^d : \|\mathbf{x}\| = 1\}$. We denote the Frobenius norm of a matrix W by $\|W\|_F^2 = \sum_{ij} W_{ij}^2$, while the spectral norm is denoted by $\|W\| = \max_{\|\mathbf{x}\|=1} \|W\mathbf{x}\|$. For a pair of vectors $\mathbf{x}, \mathbf{y} \in \mathbb{R}^d$ we denote by $\mathbf{xy} \in \mathbb{R}^d$ their point-wise product $\mathbf{xy} = (x_1y_1, \dots, x_dy_d)$.

Uniform Convergence and Covering Numbers Fix an instance space \mathcal{X} , a label space \mathcal{Y} and a loss $\ell : \mathbb{R}^d \times \mathcal{Y} \rightarrow [0, \infty)$. We say that ℓ is Lipschitz / Bounded / etc. if for any $y \in \mathcal{Y}$, $\ell(\cdot, y)$ is. Fix a class \mathcal{H} from \mathcal{X} to \mathbb{R}^d . For a distribution \mathcal{D} and a sample $S \in (\mathcal{X} \times \mathcal{Y})^m$ we define the *representativeness* of S as

$$\text{rep}_{\mathcal{D}}(S, \mathcal{H}) = \sup_{h \in \mathcal{H}} \ell_{\mathcal{D}}(h) - \ell_S(h) \text{ where } \ell_{\mathcal{D}}(h) = \mathbb{E}_{(x,y) \sim \mathcal{D}} \ell(h(x), y) \text{ and } \ell_S(h) = \frac{1}{m} \sum_{i=1}^m \ell(h(x_i), y_i)$$

We note that if $\text{rep}_{\mathcal{D}}(S, \mathcal{H}) \leq \epsilon$ then any algorithm that is guaranteed to return a function $\hat{h} \in \mathcal{H}$ will enjoy a generalization bound $\ell_{\mathcal{D}}(\hat{h}) \leq \ell_S(\hat{h}) + \epsilon$. In particular, the ERM algorithm will return a function whose loss is optimal, up to an additive factor of ϵ . We will focus on bounds on $\text{rep}_{\mathcal{D}}(S, \mathcal{H})$ when $S \sim \mathcal{D}^m$. To this end, we will rely on the connection between representativeness and the *covering numbers* of \mathcal{H} .

Definition 2.1. Fix a class \mathcal{H} of functions from \mathcal{X} to \mathbb{R}^d , an integer $m, \epsilon > 0$ and $1 \leq p \leq \infty$. We define $N_p(\mathcal{H}, m, \epsilon)$ as the minimal integer for which the following holds. For every $A \subset \mathcal{X}$ of size $\leq m$ there exists $\tilde{\mathcal{H}} \subset (\mathbb{R}^d)^{\mathcal{X}}$ such that $|\tilde{\mathcal{H}}| \leq N_p(\mathcal{H}, m, \epsilon)$ and for any $h \in \mathcal{H}$ there is $\tilde{h} \in \tilde{\mathcal{H}}$ with $\left(\mathbb{E}_{x \in A} \|h(x) - \tilde{h}(x)\|_{\infty}^p \right)^{\frac{1}{p}} \leq \epsilon$. For $p = 2$, we denote $N(\mathcal{H}, m, \epsilon) = N_2(\mathcal{H}, m, \epsilon)$.

We conclude with a lemma, which will be useful in this paper. The proof can be found in the supplementary material.

Lemma 2.2. Let $\ell : \mathbb{R}^d \times \mathcal{Y} \rightarrow \mathbb{R}$ be L -Lipschitz w.r.t. $\|\cdot\|_{\infty}$ and B -bounded. Assume that for any $0 < \epsilon \leq 1$, $\log(N(\mathcal{H}, m, \epsilon)) \leq \frac{n}{\epsilon^2}$. Then $\mathbb{E}_{S \sim \mathcal{D}^m} \text{rep}_{\mathcal{D}}(S, \mathcal{H}) \lesssim \frac{(L+B)\sqrt{n}}{\sqrt{m}} \log(m)$. Furthermore, with probability at least $1 - \delta$, $\text{rep}_{\mathcal{D}}(S, \mathcal{H}) \lesssim \frac{(L+B)\sqrt{n}}{\sqrt{m}} \log(m) + B\sqrt{\frac{2\ln(2/\delta)}{m}}$.

A Basic Inequality

Lemma 2.3. Let X_1, \dots, X_n be independent r.v. with that are σ -estimators to μ . Then $\Pr(|\text{med}(X_1, \dots, X_n) - \mu| > k\sigma) < \left(\frac{2}{k}\right)^n$.

3 Simplified Approximate Description Length

To give a soft introduction to our techniques, we first consider a simplified version of it. We next define the *approximate description length* of a class \mathcal{H} of functions from \mathcal{X} to \mathbb{R}^d , which quantifies the number of bits it takes to approximately describe a function from \mathcal{H} . We will use the following notion of approximation

113 **Definition 3.1.** A random vector $X \in \mathbb{R}^d$ is a σ -estimator to $\mathbf{x} \in \mathbb{R}^d$ if

$$\mathbb{E} X = \mathbf{x} \text{ and } \forall \mathbf{u} \in \mathbb{S}^{d-1}, \text{VAR}(\langle \mathbf{u}, X \rangle) = \mathbb{E} \langle \mathbf{u}, X - \mathbf{x} \rangle^2 \leq \sigma^2$$

114 A random function $\hat{f} : \mathcal{X} \rightarrow \mathbb{R}^d$ is a σ -estimator to $f : \mathcal{X} \rightarrow \mathbb{R}^d$ if for any $x \in \mathcal{X}$, $\hat{f}(x)$ is a
115 σ -estimator to $f(x)$.

116 A (σ, n) -compressor \mathcal{C} for a class \mathcal{H} takes as input a function $h \in \mathcal{H}$, and outputs a (random) function
117 $\mathcal{C}h$ such that (i) $\mathcal{C}h$ is a σ -estimator of h and (ii) it takes n bits to describe $\mathcal{C}h$. Formally,

118 **Definition 3.2.** A (σ, n) -compressor for \mathcal{H} is a pair $(\mathcal{C}, \Omega, \mu)$ where μ is a probability measure on Ω ,
119 and \mathcal{C} is a function $\mathcal{C} : \Omega \times \mathcal{H} \rightarrow (\mathbb{R}^d)^\mathcal{X}$ such that

- 120 1. For any $h \in \mathcal{H}$ and $x \in \mathcal{X}$, $(\mathcal{C}_\omega h)(x)$, $\omega \sim \mu$ is a σ -estimator of $h(x)$.
- 121 2. There are functions $E : \Omega \times \mathcal{H} \rightarrow \{\pm 1\}^n$ and $D : \{\pm 1\}^n \rightarrow (\mathbb{R}^d)^\mathcal{X}$ for which $\mathcal{C} = D \circ E$

122 **Definition 3.3.** We say that a class \mathcal{H} of functions from \mathcal{X} to \mathbb{R}^d has approximate description length
123 n if for any set there exists an $(1, n)$ -compressor for \mathcal{H}

124 It is not hard to see that if $(\mathcal{C}, \Omega, \mu)$ is a (σ, n) -compressor for \mathcal{H} , then

$$(\mathcal{C}_{\omega_1, \dots, \omega_k} h)(x) := \frac{\sum_{i=1}^k (\mathcal{C}_{\omega_i} h)(x)}{k}$$

125 is a $(\frac{\sigma}{\sqrt{k}}, kn)$ -compressor for \mathcal{H} . Hence, if the approximate description length of \mathcal{H} is n , then for
126 any $1 \geq \epsilon > 0$ there exists an $(\epsilon, n \lceil \epsilon^{-2} \rceil)$ -compressor for \mathcal{H} .

127 We next connect the approximate description length, to covering numbers and representativeness. We
128 separate it into two lemmas, one for $d = 1$ and one for general d , as for $d = 1$ we can prove a slightly
129 stronger bound.

130 **Lemma 3.4.** Fix a class \mathcal{H} of functions from \mathcal{X} to \mathbb{R} with approximate description length n . Then,
131 $\log(N(\mathcal{H}, m, \epsilon)) \leq n \lceil \epsilon^{-2} \rceil$. Hence, if $\ell : \mathbb{R}^d \times \mathcal{Y} \rightarrow \mathbb{R}$ is L -Lipschitz and B -bounded, then for any
132 distribution \mathcal{D} on $\mathcal{X} \times \mathcal{Y}$, $\mathbb{E}_{S \sim \mathcal{D}^m} \text{rep}_{\mathcal{D}}(S, \mathcal{H}) \lesssim \frac{(L+B)\sqrt{n}}{\sqrt{m}} \log(m)$. Furthermore, with probability
133 at least $1 - \delta$, $\text{rep}_{\mathcal{D}}(S, \mathcal{H}) \lesssim \frac{(L+B)\sqrt{n}}{\sqrt{m}} \log(m) + B \sqrt{\frac{2 \ln(2/\delta)}{m}}$

134 **Lemma 3.5.** Fix a class \mathcal{H} of functions from \mathcal{X} to \mathbb{R}^d with approximate description length n . Then,
 $\log(N_\infty(\mathcal{H}, m, \epsilon)) \leq \log(N(\mathcal{H}, m, \epsilon)) \leq n \lceil 16\epsilon^{-2} \rceil \lceil \log(dm) \rceil$

135 Hence, if $\ell : \mathbb{R}^d \times \mathcal{Y} \rightarrow \mathbb{R}$ is L -Lipschitz w.r.t. $\|\cdot\|_\infty$ and B -bounded, then for any distribution \mathcal{D}
136 on $\mathcal{X} \times \mathcal{Y}$, $\mathbb{E}_{S \sim \mathcal{D}^m} \text{rep}_{\mathcal{D}}(S, \mathcal{H}) \lesssim \frac{(L+B)\sqrt{n \log(dm)}}{\sqrt{m}} \log(m)$. Furthermore, with probability at least
137 $1 - \delta$, $\text{rep}_{\mathcal{D}}(S, \mathcal{H}) \lesssim \frac{(L+B)\sqrt{n \log(dm)}}{\sqrt{m}} \log(m) + B \sqrt{\frac{2 \ln(2/\delta)}{m}}$

138 3.1 Linear Functions

139 We next bound the approximate description length of linear functions with bounded Frobenius norm.

140 **Theorem 3.6.** Let class $\mathcal{L}_{d_1, d_2, M} = \{\mathbf{x} \in \mathbb{B}^{d_1} \mapsto W\mathbf{x} : W \text{ is } d_2 \times d_1 \text{ matrix with } \|W\|_F \leq M\}$
141 has approximate description length

$$n \leq \left\lceil \frac{1}{4} + 2M^2 \right\rceil 2 \lceil \log(2d_1 d_2 (M + 1)) \rceil$$

142 Hence, if $\ell : \mathbb{R}^{d_2} \times \mathcal{Y} \rightarrow \mathbb{R}$ is L -Lipschitz w.r.t. $\|\cdot\|_\infty$ and B -bounded, then for any distribution \mathcal{D}
143 on $\mathcal{X} \times \mathcal{Y}$

$$\mathbb{E}_{S \sim \mathcal{D}^m} \text{rep}_{\mathcal{D}}(S, \mathcal{L}_{d_1, d_2, M}) \lesssim \frac{(L+B)\sqrt{M^2 \log(d_1 d_2 M) \log(d_2 m)}}{\sqrt{m}} \log(m)$$

144 Furthermore, with probability at least $1 - \delta$,

$$\text{rep}_{\mathcal{D}}(S, \mathcal{L}_{d_1, d_2, M}) \lesssim \frac{(L+B)\sqrt{M^2 \log(d_1 d_2 M) \log(d_2 m)}}{\sqrt{m}} \log(m) + B \sqrt{\frac{2 \ln(2/\delta)}{m}}$$

We remark that the above bounds on the representativeness coincides with standard bounds ([11] for instance), up to log factors. The advantage of these bound is that they remain valid for *any output dimension* d_2 .

In order to prove theorem 3.6 we will use a randomized sketch of a matrix.

Definition 3.7. Let $\mathbf{w} \in \mathbb{R}^d$ be a vector. A random sketch of \mathbf{w} is a random vector $\hat{\mathbf{w}}$ that is samples as follows. Choose i w.p. $p_i = \frac{w_i^2}{2\|\mathbf{w}\|^2} + \frac{1}{2d}$. Then, w.p. $\frac{w_i}{p_i} - \lfloor \frac{w_i}{p_i} \rfloor$ let $b = 1$ and otherwise $b = 0$. Finally, let $\hat{\mathbf{w}} = \left(\lfloor \frac{w_i}{p_i} \rfloor + b \right) \mathbf{e}_i$. A random k -sketch of \mathbf{w} is an average of k -independent random sketches of \mathbf{w} . A random sketch and a random k -sketch of a matrix is defined similarly, with the standard matrix basis instead of the standard vector basis.

The following useful lemma shows that an sketch \mathbf{w} is a $\sqrt{\frac{1}{4} + 2\|\mathbf{w}\|^2}$ -estimator of \mathbf{w} .

Lemma 3.8. Let $\hat{\mathbf{w}}$ be a random sketch of $\mathbf{w} \in \mathbb{R}^d$. Then, (1) $\mathbb{E} \hat{\mathbf{w}} = \mathbf{w}$ and (2) for any $\mathbf{u} \in \mathbb{S}^{d-1}$, $\mathbb{E} (\langle \mathbf{u}, \hat{\mathbf{w}} \rangle - \langle \mathbf{u}, \mathbf{w} \rangle)^2 \leq \mathbb{E} \langle \mathbf{u}, \hat{\mathbf{w}} \rangle^2 \leq \frac{1}{4} + 2\|\mathbf{w}\|^2$

Proof. (of theorem 3.6) We construct a compressor for $\mathcal{L}_{d_1, d_2, M}$ as follows. Given W , we will sample a k -sketch \hat{W} of W for $k = \lceil \frac{1}{4} + 2M^2 \rceil$, and will return the function $\mathbf{x} \mapsto \hat{W}\mathbf{x}$. We claim that that $W \mapsto \hat{W}$ is a $(1, 2k \lceil \log(2d_1 d_2 (M+1)) \rceil)$ -compressor for $\mathcal{L}_{d_1, d_2, M}$. Indeed, to specify a sketch of W we need $\lceil \log(d_1 d_2) \rceil$ bits to describe the chosen index, as well as $\log(2d_1 d_2 M + 2)$ bits to describe the value in that index. Hence, $2k \lceil \log(2d_1 d_2 (M+1)) \rceil$ bits suffices to specify a k -sketch. It remains to show that for $\mathbf{x} \in \mathbb{B}^{d_1}$, $\hat{W}\mathbf{x}$ is a 1-estimator of $W\mathbf{x}$. Indeed, by lemma 3.8, $\mathbb{E} \hat{W} = W$ and therefore $\mathbb{E} \hat{W}\mathbf{x} = W\mathbf{x}$. Likewise, for $\mathbf{u} \in \mathbb{S}^{d_2-1}$. We have

$$\mathbb{E} \left(\langle \mathbf{u}, \hat{W}\mathbf{x} \rangle - \langle \mathbf{u}, W\mathbf{x} \rangle \right)^2 = \mathbb{E} \left(\langle \hat{W}, \mathbf{x}\mathbf{u}^T \rangle - \langle W, \mathbf{x}\mathbf{u}^T \rangle \right)^2 \leq \frac{\frac{1}{4} + 2M^2}{k} \leq 1$$

□

3.2 Simplified Depth 2 Networks

To demonstrate our techniques, we consider the following class of functions. We let the domain \mathcal{X} to be \mathbb{B}^d . We fix an activation function $\rho : \mathbb{R} \rightarrow \mathbb{R}$ that is assumed to be a polynomial $\rho(x) = \sum_{i=0}^k a_i x^i$ with $\sum_{n=1}^n |a_n| = 1$. For any $W \in M_{d,d}$ we define $h_W(\mathbf{x}) = \frac{1}{\sqrt{d}} \sum_{i=1}^d \rho(\langle \mathbf{w}_i, \mathbf{x} \rangle)$ Finally, we let $\mathcal{H} = \{h_W : \forall i, \|\mathbf{w}_i\| \leq \frac{1}{2}\}$. In order to build compressors for classes of networks, we will utilize to compositional structure of the classes. Specifically, we have that $\mathcal{H} = \Lambda \circ \rho \circ \mathcal{F}$ where $\mathcal{F} = \{x \mapsto W\mathbf{x} : W \text{ is } d \times d \text{ matrix with } \|\mathbf{w}_i\| \leq 1 \text{ for all } i\}$ and $\Lambda(\mathbf{x}) = \frac{1}{\sqrt{d}} \sum_{i=1}^d x_i$.

As \mathcal{F} is a subset of $\mathcal{L}_{d,d,\sqrt{d}}$, we know that there exists a $(1, O(d \log(d)))$ -compressor for it. We will use this compressor to build a compressor to $\rho \circ \mathcal{F}$, and then to $\Lambda \circ \rho \circ \mathcal{F}$. We will start with the latter, linear case, which is simpler

Lemma 3.9. Let X be a σ -estimator to $\mathbf{x} \in \mathbb{R}^{d_1}$. Let $A \in M_{d_2, d_1}$ be a matrix of spectral norm $\leq r$. Then, AX is a $(r\sigma)$ -estimator to $A\mathbf{x}$. In particular, if \mathcal{C} is a $(1, n)$ -compressor to a class \mathcal{H} of functions from \mathcal{X} to \mathbb{R}^d . Then

$$\mathcal{C}'_\omega(\Lambda \circ h) = \Lambda \circ \mathcal{C}_\omega h$$

is a $(1, n)$ -compressor to $\Lambda \circ \mathcal{H}$

We next consider the composition of \mathcal{F} with the non-linear ρ . As opposed to composition with a linear function, we cannot just generate a compression version using \mathcal{F} 's compressor and then compose with ρ . Indeed, if X is a σ -estimator to \mathbf{x} , it is not true in general that $\rho(X)$ is an estimator of $\rho(\mathbf{x})$. For instance, consider the case that $\rho(x) = x^2$, and $X = (X_1, \dots, X_d)$ is a vector of independent standard Gaussians. X is a 1-estimator of $0 \in \mathbb{R}^d$. On the other hand, $\rho(X) = (X_1^2, \dots, X_d^2)$ is not an estimator of $0 = \rho(0)$. We will therefore take a different approach. Given $f \in \mathcal{F}$, we will sample k independent estimators $\{C_{\omega_i} f\}_{i=1}^k$ from \mathcal{F} 's compressor, and define the compressed version of $\sigma \circ h$ as $\mathcal{C}'_{\omega_1, \dots, \omega_k} f = \sum_{i=0}^d a_i \prod_{j=0}^i C_{\omega_j} f$. This construction is analyzed in the following lemma

187 **Lemma 3.10.** *If \mathcal{C} is a $(\frac{1}{2}, n)$ -compressor of a class \mathcal{H} of functions from \mathcal{X} to $[-\frac{1}{2}, \frac{1}{2}]^d$. Then \mathcal{C}' is*
 188 *a $(1, n)$ -compressor of $\rho \circ \mathcal{H}$*

189 Combining theorem 3.6 and lemmas 3.9, 3.10 we have:

190 **Theorem 3.11.** *\mathcal{H} has approximation length $\lesssim d \log(d)$. Hence, if $\ell : \mathbb{R} \times \mathcal{Y} \rightarrow \mathbb{R}$ is L -Lipschitz*
 191 *and B -bounded, then for any distribution \mathcal{D} on $\mathcal{X} \times \mathcal{Y}$*

$$\mathbb{E}_{S \sim \mathcal{D}^m} \text{rep}_{\mathcal{D}}(S, \mathcal{H}) \lesssim \frac{(L+B)\sqrt{d \log(d)}}{\sqrt{m}} \log(m)$$

192 Furthermore, with probability at least $1 - \delta$,

$$\text{rep}_{\mathcal{D}}(S, \mathcal{H}) \lesssim \frac{(L+B)\sqrt{d \log(d)}}{\sqrt{m}} \log(m) + B \sqrt{\frac{2 \ln(2/\delta)}{m}}$$

193 Lemma 3.10 is implied by the following useful lemma:

194 **Lemma 3.12.** 1. *If X is a σ -estimator of \mathbf{x} then aX is a $(|a|\sigma)$ -estimator of aX*

195 2. *Suppose that for $n = 1, 2, 3, \dots$ X_n is a σ_n -estimator of $\mathbf{x}_n \in \mathbb{R}^d$. Assume furthermore*
 196 *that $\sum_{n=1}^{\infty} \mathbf{x}_n$ and $\sum_{n=1}^{\infty} \sigma_n$ converge to $\mathbf{x} \in \mathbb{R}^d$ and $\sigma \in [0, \infty)$. Then, $\sum_{n=1}^{\infty} X_n$ is a*
 197 *σ -estimator of \mathbf{x}*

198 3. *Suppose that $\{X_i\}_{i=1}^k$ are independent σ_i -estimators of $\mathbf{x}_i \in \mathbb{R}^d$. Then $\prod_{i=1}^k X_i$ is a*
 199 *σ' -estimator of $\prod_{i=1}^k \mathbf{x}_i$ for $\sigma'^2 = \prod_{i=1}^k (\sigma_i^2 + \|\mathbf{x}_i\|_{\infty}^2) - \prod_{i=1}^k \|\mathbf{x}_i\|_{\infty}^2$*

200 We note that the bounds in the above lemma are all tight.

201 4 Approximation Description Length

202 In this section we refine the definition of approximate description length that were given in section 3.
 203 We start with the encoding of the compressed version of the functions. Instead of standard strings,
 204 we will use what we call *bracketed string*. The reason for that often, in order to create a compressed
 205 version of a function, we concatenate compressed versions of other functions. This results with
 206 strings with a nested structure. For instance, consider the case that a function h is encoded by the
 207 concatenation of h_1 and h_2 . Furthermore, assume that h_1 is encoded by the string 01, while h_2 is
 208 encoded by the concatenation of h_3, h_4 and h_5 that are in turn encoded by the strings 101, 0101 and
 209 1110. The encoding of h will then be $[[01][[101][0101][1110]]]$. We note that in section 3 we could
 210 avoid this issue since the length of the strings and the recursive structure were fixed, and did not
 211 depend on the function we try to compress. Formally, we define

212 **Definition 4.1.** *A bracketed string is a rooted tree S , such that (i) the children of each edge are*
 213 *ordered, (ii) there are no nodes with a single child, and (iii) the leaves are labeled by $\{0, 1\}$. The*
 214 *length, $\text{len}(S)$ of S is the number of its leaves.*

215 Let S be a bracketed string. There is a linear order on its leaves that is defined as follows. Fix a pair
 216 of leaves, v_1 and v_2 , and let u be their LCA. Let u_1 (resp. u_2) be the child of u that lie on the path to
 217 v_1 (resp. v_2). We define $v_1 < v_2$ if $u_1 < u_2$ and $v_1 > v_2$ otherwise (note that necessarily $u_1 \neq u_2$).
 218 Let v_1, \dots, v_n be the leaves of T , ordered according to the above order, and let b_1, \dots, b_n be the
 219 corresponding bits. The string associated with T is $s = b_1 \dots b_n$. We denote by \mathcal{S}_n the collection of
 220 bracketed strings of length $\leq n$, and by $\mathcal{S} = \bigcup_{n=1}^{\infty} \mathcal{S}_n$ the collection of all bracketed strings.

221 The following lemma shows that in log-scale, the number of bracketed strings of length $\leq n$ differ
 222 from standard strings of length $\leq n$ by only a constant factor

223 **Lemma 4.2.** $|\mathcal{S}_n| \leq 32^n$

224 We next revisit the definition of a compressor for a class \mathcal{H} . The definition of compressor will now
 225 have a third parameter, n_s , in addition to σ and n . We will make three changes in the definition.
 226 The first, which is only for the sake of convenience, is that we will use bracketed strings rather than
 227 standard strings. The second change, is that the length of the encoding string will be bounded only

in expectation. The final change is that the compressor can now output a *seed*. That is, given a function $h \in \mathcal{H}$ that we want to compress, the compressor can generate both a non-random seed $E_s(h) \in \mathcal{S}_{n_s}$ and a random encoding $E(\omega, h) \in \mathcal{S}$ with $\mathbb{E}_{\omega \sim \mu} \text{len}(E(\omega, h)) \leq n$. Together, $E_s(h)$ and $E(\omega, h)$ encode a σ -estimator. Namely, there is a function $D : \mathcal{S}_{n_s} \times \mathcal{S} \rightarrow (\mathbb{R}^d)^{\mathcal{X}}$ such that $D(E_s(h), E(\omega, h))$, $\omega \sim \mu$ is a σ -estimator of h . The advantage of using seeds is that it will allow us to generate many independent estimators, at a lower cost. In the case that $n \ll n_s$, the cost of generating k independent estimators of $h \in \mathcal{H}$ is $n_s + kn$ bits (in expectation) instead of $k(n_s + n)$ bits. Indeed, we can encode k estimators by a single seed $E_s(h)$ and k independent “regular” encodings $E(\omega_k, h), \dots, E(\omega_k, h)$. The formal definition is given next.

Definition 4.3. A (σ, n_s, n) -compressor for \mathcal{H} is a 5-tuple $\mathcal{C} = (E_s, E, D, \Omega, \mu)$ where μ is a probability measure on Ω , and E_s, E, D are functions $E_s : \mathcal{H} \rightarrow \mathcal{T}^{n_s}$, $E : \Omega \times \mathcal{H} \rightarrow \mathcal{T}$, and $D : \mathcal{T}^{n_s} \times \mathcal{T} \rightarrow (\mathbb{R}^d)^{\mathcal{X}}$ such that for any $h \in \mathcal{H}$ and $x \in \mathcal{X}$ (1) $D(E_s(h), E(\omega, h))$, $\omega \sim \mu$ is a σ -estimator of h and (2) $\mathbb{E}_{\omega \sim \mu} \text{len}(E(\omega, h)) \leq n$

We finally revisit the definition of approximate description length. We will add an additional parameter, to accommodate the use of seeds. Likewise, the approximate description length will now be a function of m – we will say that \mathcal{H} has approximate description length $(n_s(m), n(m))$ if there is a $(1, n_s(m), n(m))$ -compressor for the restriction of \mathcal{H} to any set $A \subset \mathcal{X}$ of size at most m . Formally:

Definition 4.4. We say that a class \mathcal{H} of functions from \mathcal{X} to \mathbb{R}^d has approximate description length $(n_s(m), n(m))$ if for any set $A \subset \mathcal{X}$ of size $\leq m$ there exists a $(1, n_s(m), n(m))$ -compressor for $\mathcal{H}|_A$

It is not hard to see that if \mathcal{H} has approximate description length $(n_s(m), n(m))$, then for any $1 \geq \epsilon > 0$ and a set $A \subset \mathcal{X}$ of size $\leq m$, there exists an $(\epsilon, n_s(m), n(m) \lceil \epsilon^{-2} \rceil)$ -compressor for $\mathcal{H}|_A$. We next connect the approximate description length, to covering numbers and representativeness. The proofs are similar the the proofs of lemmas 3.4 and 3.5.

Lemma 4.5. Fix a class \mathcal{H} of functions from \mathcal{X} to \mathbb{R} with approximate description length $(n_s(m), n(m))$. Then, $\log(N(\mathcal{H}, m, \epsilon)) \lesssim n_s(m) + \frac{n(m)}{\epsilon^2}$. Hence, if $\ell : \mathbb{R}^d \times \mathcal{Y} \rightarrow \mathbb{R}$ is L -Lipschitz and B -bounded, then for any distribution \mathcal{D} on $\mathcal{X} \times \mathcal{Y}$

$$\mathbb{E}_{S \sim \mathcal{D}^m} \text{rep}_{\mathcal{D}}(S, \mathcal{H}) \lesssim \frac{(L + B) \sqrt{n_s(m) + n(m)}}{\sqrt{m}} \log(m)$$

Furthermore, with probability at least $1 - \delta$,

$$\text{rep}_{\mathcal{D}}(S, \mathcal{H}) \lesssim \frac{(L + B) \sqrt{n_s(m) + n(m)}}{\sqrt{m}} \log(m) + B \sqrt{\frac{2 \ln(2/\delta)}{m}}$$

Lemma 4.6. Fix a class \mathcal{H} of functions from \mathcal{X} to \mathbb{R}^d with approximate description length $(n_s(m), n(m))$. Then, $\log(N(\mathcal{H}, m, \epsilon)) \leq \log(N_{\infty}(\mathcal{H}, m, \epsilon)) \lesssim n_s(m) + \frac{n(m) \log(dm)}{\epsilon^2}$. Hence, if $\ell : \mathbb{R}^d \times \mathcal{Y} \rightarrow \mathbb{R}$ is L -Lipschitz w.r.t. $\|\cdot\|_{\infty}$ and B -bounded, then for any distribution \mathcal{D} on $\mathcal{X} \times \mathcal{Y}$

$$\mathbb{E}_{S \sim \mathcal{D}^m} \text{rep}_{\mathcal{D}}(S, \mathcal{H}) \lesssim \frac{(L + B) \sqrt{n_s(m) + n(m) \log(dm)}}{\sqrt{m}} \log(m)$$

Furthermore, with probability at least $1 - \delta$,

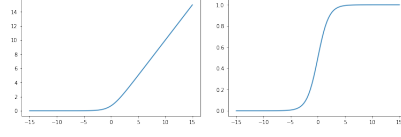
$$\text{rep}_{\mathcal{D}}(S, \mathcal{H}) \lesssim \frac{(L + B) \sqrt{n_s(m) + n(m) \log(dm)}}{\sqrt{m}} \log(m) + B \sqrt{\frac{2 \ln(2/\delta)}{m}}$$

We next analyze the behavior of the approximate description length under various operations

Lemma 4.7. Let $\mathcal{H}_1, \mathcal{H}_2$ be classes of functions from \mathcal{X} to \mathbb{R}^d with approximate description length of $(n_s^1(m), n^1(m))$ and $(n_s^2(m), n^2(m))$. Then $\mathcal{H}_1 + \mathcal{H}_2$ has approximate description length of $(n_s^1(m) + n_s^2(m), 2n^1(m) + 2n^2(m))$

Lemma 4.8. Let \mathcal{H} be a class of functions from \mathcal{X} to \mathbb{R}^d with approximate description length of $(n_s(m), n(m))$. Let A be $d_2 \times d_1$ matrix. Then $A \circ \mathcal{H}_1$ has approximate description length $(n_s(m), \lceil \|A\|^2 \rceil n(m))$

Figure 1: The functions $\ln(1 + e^x)$ and $\frac{e^x}{1+e^x}$



Definition 4.9. Denote by $\mathcal{L}_{d_1, d_2, r, R}$ the class of all $d_2 \times d_1$ matrices of spectral norm at most r and Frobenius norm at most R .

Lemma 4.10. Let \mathcal{H} be a class of functions from \mathcal{X} to \mathbb{R}^{d_1} with approximate description length $(n_s(m), n(m))$. Assume furthermore that for any $x \in \mathcal{X}$ and $h \in \mathcal{H}$ we have that $\|h(x)\| \leq B$. Then, $\mathcal{L}_{d_1, d_2, r, R} \circ \mathcal{H}$ has approximate description length

$$(n_s(m), n(m)O(r^2 + 1) + O((d_1 + B^2)(R^2 + 1)\log(Rd_1d_2 + 1)))$$

Definition 4.11. A function $f : \mathbb{R} \rightarrow \mathbb{R}$ is B -strongly-bounded if for all $n \geq 1$, $\|f^{(n)}\|_\infty \leq n!B^n$. Likewise, f is strongly-bounded if it is B -strongly-bounded for some B

We note that

Lemma 4.12. If f is B -strongly-bounded then f is analytic and its Taylor coefficients around any point are bounded by B^n

The following lemma gives an example to a strongly bounded sigmoid function, as well as a strongly bounded smoothened version of the ReLU (see figure 1).

Lemma 4.13. The functions $\ln(1 + e^x)$ and $\frac{e^x}{1+e^x}$ are strongly-bounded

Lemma 4.14. Let \mathcal{H} be a class of functions from \mathcal{X} to \mathbb{R}^d with approximate description length of $(n_s(m), n(m))$. Let $\rho : \mathbb{R} \rightarrow \mathbb{R}$ be B -strongly-bounded. Then, $\rho \circ \mathcal{H}$ has approximate description length of

$$(n_s(m) + O(n(m)B^2 \log(md)), O(n(m)B^2 \log(d)))$$

5 Sample Complexity of Neural Networks

Fix the instance space \mathcal{X} to be the ball of radius \sqrt{d} in \mathbb{R}^d (in particular $[-1, 1]^d \subset \mathcal{X}$) and a B -strongly-bounded activation ρ . Fix matrices $W_i^0 \in M_{d_i, d_{i-1}}$, $i = 1, \dots, t$. Consider the following class of depth- t networks

$$\mathcal{N}_{r, R}(W_1^0, \dots, W_t^0) = \{W_t \circ \rho \circ W_{t-1} \circ \rho \dots \circ \rho \circ W_1 : \|W_i - W_i^0\| \leq r, \|W_i - W_i^0\|_F \leq R\}$$

We note that

$$\mathcal{N}_{r, R}(W_1^0, \dots, W_t^0) = \mathcal{N}_{r, R}(W_t^0) \circ \dots \circ \mathcal{N}_{r, R}(W_1^0)$$

The following lemma analyzes the cost, in terms of approximate description length, when moving from a class \mathcal{H} to $\mathcal{N}_{r, R}(W^0) \circ \mathcal{H}$.

Lemma 5.1. Let \mathcal{H} be a class of functions from \mathcal{X} to \mathbb{R}^{d_1} with approximate description length $(n_s(m), n(m))$ and $\|h(x)\| \leq M$ for any $x \in \mathcal{X}$ and $h \in \mathcal{H}$. Fix $W^0 \in M_{d_2, d_1}$. Then, $\mathcal{N}_{r, R}(W_t^0) \circ \mathcal{H}$ has approximate description length of

$$(n_s(m) + n'(m)B^2 \log(md_2), n'(m)B^2 \log(d_2))$$

for

$$n'(m) = n(m)O(r^2 + \|W^0\|^2 + 1) + O((d_1 + M^2)(R^2 + 1)\log(Rd_1d_2 + 1))$$

The lemma follows by combining lemmas 4.7, 4.8, 4.10 and 4.14. We note that in the case that $d_1, d_2 \leq d$, $M = O(\sqrt{d_1})$, $B, r, \|W^0\| = O(1)$ (and hence $R = O(\sqrt{d})$) and $R \geq 1$ we get that $\mathcal{N}_{r, R}(W^0) \circ \mathcal{H}$ has approximate description length of

$$(n_s(m) + O(n(m)\log(md)), O(n(m)\log(d)) + O(d_1R^2\log^2(d)))$$

By induction, the approximate description length of $\mathcal{N}_{r, R}(W_1^0, \dots, W_t^0)$ is

$$(dR^2O(\log(d))^t \log(md), dR^2O(\log(d))^{t+1})$$

References

- [1] Martin Anthony and Peter Bartlett. *Neural Network Learning: Theoretical Foundations*. Cambridge University Press, 1999.
- [2] Sanjeev Arora, Rong Ge, Behnam Neyshabur, and Yi Zhang. Stronger generalization bounds for deep nets via a compression approach. In *ICML*, 2018.
- [3] P. L. Bartlett and S. Mendelson. Rademacher and Gaussian complexities: Risk bounds and structural results. *Journal of Machine Learning Research*, 3:463–482, 2002.
- [4] Peter L Bartlett, Dylan J Foster, and Matus J Telgarsky. Spectrally-normalized margin bounds for neural networks. In *Advances in Neural Information Processing Systems*, pages 6240–6249, 2017.
- [5] Noah Golowich, Alexander Rakhlin, and Ohad Shamir. Size-independent sample complexity of neural networks. In *COLT*, 2018.
- [6] Vaishnavh Nagarajan and J Zico Kolter. Generalization in deep networks: The role of distance from initialization. *arXiv preprint arXiv:1901.01672*, 2019.
- [7] Behnam Neyshabur, Srinadh Bhojanapalli, and Nathan Srebro. A pac-bayesian approach to spectrally-normalized margin bounds for neural networks. In *ICLR*, 2018.
- [8] Behnam Neyshabur, Zhiyuan Li, Srinadh Bhojanapalli, Yann LeCun, and Nathan Srebro. The role of over-parametrization in generalization of neural networks. In *ICLR*, 2019.
- [9] Behnam Neyshabur, Ryota Tomioka, and Nathan Srebro. Norm-based capacity control in neural networks. In *Conference on Learning Theory*, pages 1376–1401, 2015.
- [10] R.E. Schapire, Y. Freund, P. Bartlett, and W.S. Lee. Boosting the margin: A new explanation for the effectiveness of voting methods. In *Machine Learning: Proceedings of the Fourteenth International Conference*, pages 322–330, 1997. To appear, *The Annals of Statistics*.
- [11] Shai Shalev-Shwartz and Shai Ben-David. *Understanding machine learning: From theory to algorithms*. Cambridge university press, 2014.
- [12] Ilya Sutskever, James Martens, George Dahl, and Geoffrey Hinton. On the importance of initialization and momentum in deep learning. In *International conference on machine learning*, pages 1139–1147, 2013.

6 Omitted proofs

Lemma 6.1. [11] Let $\ell : \mathbb{R}^d \times \mathcal{Y} \rightarrow \mathbb{R}$ be B -bounded. Then

$$\mathbb{E}_{S \sim \mathcal{D}^m} \text{rep}_{\mathcal{D}}(S, \mathcal{H}) \leq B2^{-M+1} + \frac{12B}{\sqrt{m}} \sum_{k=1}^M 2^{-k} \sqrt{\ln(N(\ell \circ \mathcal{H}, m, B2^{-k}))}$$

Furthermore, with probability at least $1 - \delta$,

$$\text{rep}_{\mathcal{D}}(S, \mathcal{H}) \leq B2^{-M+1} + \frac{12B}{\sqrt{m}} \sum_{k=1}^M 2^{-k} \sqrt{\ln(N(\ell \circ \mathcal{H}, m, B2^{-k}))} + B \sqrt{\frac{2 \ln(2/\delta)}{m}}$$

Proof. (of lemma 2.2) Denote

$$A = B2^{-M+1} + \frac{12B}{\sqrt{m}} \sum_{k=1}^M 2^{-k} \sqrt{\ln(N(\ell \circ \mathcal{H}, m, B2^{-k}))}$$

We will show that $A \lesssim \frac{(L+B)\sqrt{n}}{\sqrt{m}} \log(m)$. We have, if $\frac{B2^{-k}}{L} \leq 1$,

$$\ln(N(\ell \circ \mathcal{H}, m, B2^{-k})) \leq \ln\left(N\left(\mathcal{H}, m, \frac{B}{L}2^{-k}\right)\right) \leq \frac{nL^2 2^{2k}}{B^2} + n$$

Hence,

$$A \leq B2^{-M+1} + \frac{12B}{\sqrt{m}} \sum_{k=1}^M \frac{\sqrt{n}L}{B} + \frac{12B}{\sqrt{m}} \sum_{k=1}^M 2^{-k} \sqrt{n} \leq B2^{-M+1} + \frac{12(LM + B)\sqrt{n}}{\sqrt{m}}$$

Choosing $M = \log(\sqrt{\frac{m}{n}})$ we get,

$$A \leq \frac{12(L \log(\sqrt{\frac{m}{n}}) + B)\sqrt{n} + B\sqrt{n}}{\sqrt{m}}$$

□

Proof. (of lemma 2.3) We have that $\Pr(|X_i - \mu| > k\sigma) \leq \frac{1}{k^2}$. It follows that the probability that $\geq \frac{n}{2}$ of X_1, \dots, X_n fall outside of the segment $(\mu - k\sigma, \mu + k\sigma)$ is bounded by

$$\binom{n}{\lceil n/2 \rceil} \left(\frac{1}{k^2}\right)^{\lceil n/2 \rceil} < 2^n \left(\frac{1}{k^2}\right)^{\lceil n/2 \rceil} \leq \left(\frac{2}{k}\right)^n$$

□

Proof. (of lemma 3.4) Fix a set $A \subset \mathcal{X}$. Let $(\mathcal{C}, \Omega, \mu)$ be a $(n \lceil \epsilon^{-2} \rceil, \epsilon)$ -compressor for \mathcal{H} . Let $\tilde{\mathcal{H}}$ be the range of \mathcal{C} . Note that $|\tilde{\mathcal{H}}| \leq 2^{n \lceil \epsilon^{-2} \rceil}$. Fix $h \in \mathcal{H}$. It is enough to show that there is $\tilde{h} \in \tilde{\mathcal{H}}$ with $\mathbb{E}_{x \in A} (h(x) - \tilde{h}(x))^2 \leq \epsilon^2$. Indeed,

$$\mathbb{E}_{\omega \sim \mu} \mathbb{E}_{x \in A} (h(x) - (\mathcal{C}_{\omega} h)(x))^2 = \mathbb{E}_{x \in A} \mathbb{E}_{\omega \sim \mu} (h(x) - (\mathcal{C}_{\omega} h)(x))^2 \leq \epsilon^2.$$

Hence, there exists $\tilde{h} \in \tilde{\mathcal{H}}$ for which $\mathbb{E}_{x \in A} (h(x) - \tilde{h}(x))^2 \leq \epsilon^2$

□

Proof. (of lemma 3.5) Denote $k = \lceil \log(dm) \rceil$. Fix a set $A \subset \mathcal{X}$. Let \mathcal{C} be a $(n \lceil 16\epsilon^{-2} \rceil, \frac{\epsilon}{4})$ -compressor for \mathcal{H} . Define

$$(\mathcal{C}'_{\omega_1, \dots, \omega_k} h)(x) = \text{med}((\mathcal{C}_{\omega_1} h)(x), \dots, (\mathcal{C}_{\omega_k} h)(x))$$

345 Let $\tilde{\mathcal{H}}$ be the range of \mathcal{C}' . Note that $|\tilde{\mathcal{H}}| \leq 2^{kn \lceil 16\epsilon^{-2} \rceil}$. Fix $h \in \mathcal{H}$. It is enough to show that there is
 346 $\tilde{h} \in \tilde{\mathcal{H}}$ with $\max_{x \in A} \|h(x) - \tilde{h}(x)\|_\infty \leq \epsilon$. By lemma 2.3 we have that

$$\Pr_{\omega_1, \dots, \omega_k \sim \mu} (\exists x \in A, |(\mathcal{C}'_{\omega_1, \dots, \omega_k} h)(x) - h(x)| > \epsilon) < dm2^{-k} \leq 1$$

347 In particular, there exists $\tilde{h} \in \tilde{\mathcal{H}}$ for which $\max_{x \in A} \|h(x) - \tilde{h}(x)\|_\infty \leq \epsilon$ □

348 *Proof.* (of lemma 3.8) Items 1. is straight forward. To see item 2. note that

$$\begin{aligned} \mathbb{E} (\langle \mathbf{u}, \hat{\mathbf{w}} \rangle - \langle \mathbf{u}, \mathbf{w} \rangle)^2 &\leq \mathbb{E} \langle \mathbf{u}, \mathbf{w} \rangle^2 \\ &= \sum_i p_i \left(\left(\frac{w_i}{p_i} \right) \left(\left\lfloor \frac{w_i}{p_i} \right\rfloor + 1 \right)^2 + \left(1 - \left\lfloor \frac{w_i}{p_i} \right\rfloor \right) \left(\left\lfloor \frac{w_i}{p_i} \right\rfloor \right)^2 \right) u_i^2 \\ &= \sum_i p_i \left(\left(\left\lfloor \frac{w_i}{p_i} \right\rfloor \right)^2 + 2 \left\lfloor \frac{w_i}{p_i} \right\rfloor \left\lceil \frac{w_i}{p_i} \right\rceil + \left\lceil \frac{w_i}{p_i} \right\rceil^2 \right) u_i^2 \\ &= \sum_i p_i \left(\left(\left\lfloor \frac{w_i}{p_i} \right\rfloor + \left\lceil \frac{w_i}{p_i} \right\rceil \right)^2 + \left\lfloor \frac{w_i}{p_i} \right\rfloor - \left\lceil \frac{w_i}{p_i} \right\rceil^2 \right) u_i^2 \\ &= \sum_i p_i \left(\left(\frac{w_i}{p_i} \right)^2 + \left\lfloor \frac{w_i}{p_i} \right\rfloor \left(1 - \left\lfloor \frac{w_i}{p_i} \right\rfloor \right) \right) u_i^2 \\ &\leq \sum_i p_i \left(\left(\frac{w_i}{p_i} \right)^2 + \frac{1}{4} \right) u_i^2 \\ &\leq \frac{1}{4} \|\mathbf{u}\|_\infty^2 + \sum_i \frac{w_i^2 u_i^2}{p_i} \\ &\leq \frac{1}{4} + \sum_i \frac{w_i^2 u_i^2}{p_i} \end{aligned}$$

349 Now, since $p_i = \frac{w_i^2}{2\|\mathbf{w}\|^2} + \frac{1}{2d}$ we have

$$\sum_i \frac{w_i^2 u_i^2}{p_i} \leq \sum_i \frac{w_i^2 u_i^2}{\frac{w_i^2}{2\|\mathbf{w}\|^2}} = 2\|\mathbf{w}\|^2 \sum_i u_i^2 = 2\|\mathbf{w}\|^2$$

350 □

351 *Proof.* (of lemma 3.9) We have $\mathbb{E} AX = A \mathbb{E} X = A\mathbf{x}$. Furthermore, for any $\mathbf{u} \in \mathbb{S}^{d_2-1}$,

$$\mathbb{E} \langle \mathbf{u}, AX - A\mathbf{x} \rangle^2 = \mathbb{E} \langle A^T \mathbf{u}, X - \mathbf{x} \rangle^2 \leq \|A^T \mathbf{u}\|^2 \sigma^2 \leq r^2 \sigma^2$$

352 □

353 *Proof.* (of lemma 3.12) 1. and 2. are straight forward. We next prove 3. By replacing each X_i with
 354 $\frac{X_i}{\sigma_i}$ we can assume w.l.o.g. that $\sigma_1 = \dots = \sigma_k = 1$. We have

$$\begin{aligned}
 \mathbb{E}_{X_1, \dots, X_k} \left\langle \mathbf{u}, \prod_{i=1}^k X_i - \prod_{i=1}^k \mathbf{x}_i \right\rangle^2 &= \mathbb{E}_{X_1, \dots, X_k} \left\langle \mathbf{u}, \prod_{i=1}^k ((X_i - \mathbf{x}_i) + \mathbf{x}_i) - \prod_{i=1}^k \mathbf{x}_i \right\rangle^2 \\
 &= \mathbb{E}_{X_1, \dots, X_k} \left\langle \mathbf{u}, \sum_{A \subset [k]} \prod_{i \in A} (X_i - \mathbf{x}_i) \prod_{i \in A^c} \mathbf{x}_i - \prod_{i=1}^k \mathbf{x}_i \right\rangle^2 \\
 &= \mathbb{E}_{X_1, \dots, X_k} \left(\left\langle \mathbf{u}, \sum_{A \subset [k]} \prod_{i \in A} (X_i - \mathbf{x}_i) \prod_{i \in A^c} \mathbf{x}_i \right\rangle - \left\langle \mathbf{u}, \prod_{i=1}^k \mathbf{x}_i \right\rangle \right)^2 \\
 &= \mathbb{E}_{X_1, \dots, X_k} \sum_{A \subset [k]} \sum_{B \subset [k]} \left\langle \mathbf{u}, \prod_{i \in A} (X_i - \mathbf{x}_i) \prod_{i \in A^c} \mathbf{x}_i \right\rangle \left\langle \mathbf{u}, \prod_{i \in B} (X_i - \mathbf{x}_i) \prod_{i \in B^c} \mathbf{x}_i \right\rangle \\
 &\quad - 2 \mathbb{E}_{X_1, \dots, X_k} \sum_{A \subset [k]} \left\langle \mathbf{u}, \prod_{i=1}^k \mathbf{x}_i \right\rangle \left\langle \mathbf{u}, \prod_{i \in A} (X_i - \mathbf{x}_i) \prod_{i \in A^c} \mathbf{x}_i \right\rangle \\
 &\quad + \left\langle \mathbf{u}, \prod_{i=1}^k \mathbf{x}_i \right\rangle^2 \\
 &\stackrel{(1)}{=} \mathbb{E}_{X_1, \dots, X_k} \sum_{A \subset [k]} \left\langle \mathbf{u}, \prod_{i \in A} (X_i - \mathbf{x}_i) \prod_{i \in A^c} \mathbf{x}_i \right\rangle^2 - \left\langle \mathbf{u}, \prod_{i=1}^k \mathbf{x}_i \right\rangle^2 \\
 &\stackrel{(2)}{\leq} \sum_{A \subset [k], A \neq [k]} \left\| \mathbf{u} \prod_{i \in A^c} \mathbf{x}_i \right\|^2 \\
 &= \sum_{A \subset [k], A \neq \emptyset} \left\| \mathbf{u} \prod_{i \in A} \mathbf{x}_i \right\|^2 \\
 &\stackrel{(3)}{\leq} \sum_{A \subset [k], A \neq \emptyset} \prod_{i \in A} \|\mathbf{x}_i\|_\infty^2 \\
 &= \prod_{i=1}^k (1 + \|\mathbf{x}_i\|_\infty^2) - \prod_{i=1}^k \|\mathbf{x}_i\|_\infty^2
 \end{aligned}$$

355 (1) If $A \neq B$, then w.l.o.g. $k \in A \setminus B$. In this case we have

$$\begin{aligned}
 &\mathbb{E}_{X_1, \dots, X_k} \left\langle \mathbf{u}, \prod_{i \in A} (X_i - \mathbf{x}_i) \prod_{i \in A^c} \mathbf{x}_i \right\rangle \left\langle \mathbf{u}, \prod_{i \in B} (X_i - \mathbf{x}_i) \prod_{i \in A^c} \mathbf{x}_i \right\rangle \\
 &= \mathbb{E}_{X_1, \dots, X_{k-1}} \mathbb{E}_{X_k} \left\langle \mathbf{u}, \prod_{i \in A} (X_i - \mathbf{x}_i) \prod_{i \in A^c} \mathbf{x}_i \right\rangle \left\langle \mathbf{u}, \prod_{i \in B} (X_i - \mathbf{x}_i) \prod_{i \in B^c} \mathbf{x}_i \right\rangle \\
 &= \mathbb{E}_{X_1, \dots, X_{k-1}} \left\langle \mathbf{u}, \prod_{i \in B} (X_i - \mathbf{x}_i) \prod_{i \in B^c} \mathbf{x}_i \right\rangle \mathbb{E}_{X_k} \left\langle \mathbf{u}, \prod_{i \in A} (X_i - \mathbf{x}_i) \prod_{i \in A^c} \mathbf{x}_i \right\rangle \\
 &= \mathbb{E}_{X_1, \dots, X_{k-1}} \left\langle \mathbf{u}, \prod_{i \in B} (X_i - \mathbf{x}_i) \prod_{i \in B^c} \mathbf{x}_i \right\rangle \left\langle \mathbf{u}, \prod_{i \in A \setminus [k]} (X_i - \mathbf{x}_i) \overbrace{\mathbb{E}_{X_k} (X_k - \mathbf{x}_k)}^{=0} \prod_{i \in A^c} \mathbf{x}_i \right\rangle \\
 &= 0
 \end{aligned}$$

356 Similarly, if $A \neq \emptyset$, then w.l.o.g. $k \in A$. In this case we have

$$\begin{aligned}
 \mathbb{E}_{X_1, \dots, X_k} \left\langle \mathbf{u}, \prod_{i=1}^k \mathbf{x}_i \right\rangle \left\langle \mathbf{u}, \prod_{i \in A} (X_i - \mathbf{x}_i) \prod_{i \in A^c} \mathbf{x}_i \right\rangle &= \mathbb{E}_{X_1, \dots, X_{k-1}} \mathbb{E}_{X_k} \left\langle \mathbf{u}, \prod_{i=1}^k \mathbf{x}_i \right\rangle \left\langle \mathbf{u}, \prod_{i \in A} (X_i - \mathbf{x}_i) \prod_{i \in A^c} \mathbf{x}_i \right\rangle \\
 &= \mathbb{E}_{X_1, \dots, X_{k-1}} \left\langle \mathbf{u}, \prod_{i=1}^k \mathbf{x}_i \right\rangle \left\langle \mathbf{u}, \prod_{i \in A \setminus [k]} (X_i - \mathbf{x}_i) \overbrace{\mathbb{E}_{X_k} (X_k - \mathbf{x}_k)}^{=0} \prod_{i \in A^c} \mathbf{x}_i \right\rangle
 \end{aligned}$$

357 (2) Fix a set A that is w.l.o.g. $A = \{1, \dots, k'\}$. We note that if $X \in \mathbb{R}^d$ is a 1-estimator to 0,
 358 then for any vector $\mathbf{z} \in \mathbb{R}^d$

$$\mathbb{E}_X \|\mathbf{z}X\|^2 = \sum_{i=1}^d z_i^2 \mathbb{E}_X X_i^2 = \sum_{i=1}^d z_i^2 \mathbb{E}_X \langle \mathbf{e}_i, X \rangle^2 \leq \sum_{i=1}^d z_i^2 = \|\mathbf{z}\|^2$$

It follows that

$$\begin{aligned}
\mathbb{E}_{X_1, \dots, X_{k'-1}} \left\| \mathbf{u} \prod_{i=k'+1}^k \mathbf{x}_i \prod_{i=1}^{k'-1} (X_i - \mathbf{x}_i) \right\|^2 &= \mathbb{E}_{X_1, \dots, X_{k'-2}} \mathbb{E}_{X_{k'-1}} \left\| \mathbf{u} \prod_{i=k'+1}^k \mathbf{x}_i \prod_{i=1}^{k'-1} (X_i - \mathbf{x}_i) \right\|^2 \\
&\leq \mathbb{E}_{X_1, \dots, X_{k'-2}} \left\| \mathbf{u} \prod_{i=k'+1}^k \mathbf{x}_i \prod_{i=1}^{k'-2} (X_i - \mathbf{x}_i) \right\|^2 \\
&\vdots \\
&\leq \left\| \mathbf{u} \prod_{i=k'+1}^k \mathbf{x}_i \right\|^2 \\
&= \left\| \mathbf{u} \prod_{i \in A^c} \mathbf{x}_i \right\|^2
\end{aligned}$$

Hence,

$$\begin{aligned}
\mathbb{E}_{X_1, \dots, X_k} \left\langle \mathbf{u}, \prod_{i \in A} (X_i - \mathbf{x}_i) \prod_{i \in A^c} \mathbf{x}_i \right\rangle^2 &= \mathbb{E}_{X_1, \dots, X_{k'-1}} \mathbb{E}_{X_{k'}} \left\langle \mathbf{u} \prod_{i=k'+1}^k \mathbf{x}_i \prod_{i=1}^{k'-1} (X_i - \mathbf{x}_i), (X_{k'} - \mathbf{x}_{k'}) \right\rangle^2 \\
&\stackrel{X_{k'} \text{ is 1-estimator of } \mathbf{x}_{k'}}{\leq} \mathbb{E}_{X_1, \dots, X_{k'-1}} \left\| \mathbf{u} \prod_{i=k'+1}^k \mathbf{x}_i \prod_{i=1}^{k'-1} (X_i - \mathbf{x}_i) \right\|^2 \\
&\leq \left\| \mathbf{u} \prod_{i \in A^c} \mathbf{x}_i \right\|^2
\end{aligned}$$

(3) If $\mathbf{z} = \mathbf{u} \prod_{i \in A} \mathbf{x}_i$ then for any $j \in [d]$, $|z_j| \leq |u_j| \prod_{i \in A} \|\mathbf{x}_i\|_\infty$. Hence,

$$\|\mathbf{z}\|^2 \leq \prod_{i \in A} \|\mathbf{x}_i\|_\infty \sum_{j=1}^d u_j^2 = \prod_{i \in A} \|\mathbf{x}_i\|_\infty$$

□

Proof. (of lemma 4.2) By adding a pair of brackets around each bit, each bracketed string can be described by $2n - 1$ correctly matched pairs of brackets, and a string of length $\leq n$. As the number of ways to correctly match k pairs of brackets is the Catalan number $C_k = \frac{1}{k+1} \binom{2k}{k} \leq 2^{2k}$, we have,
 $|\mathcal{S}_n| \leq 2^{4n-2} 2^{n+1}$ □

Proof. (of lemma 4.10) Fix as set $A \subset \mathcal{X}$ of size m . We will construct a compressor to $\mathcal{L}_{d_1, d_2, r, R} \circ \mathcal{H}$ as follows. Given $h \in \mathcal{H}$ and $W \in \mathcal{L}_{d_1, d_2, r, R}$ we first pay a seed cost $n_s(m)$ to use \mathcal{H} 's compressor. Then, we use \mathcal{H} 's compressor to generate a $\sqrt{\frac{1}{k_1}}$ -estimator \hat{h} of h , at the cost of $k_1 n(m)$ bits. Then, we take \hat{W} to be a k_2 -sketch of W , at the costs of $k_2 O(\log(d_1 d_2 R + 1))$ bits. Finally, we output the estimator $\hat{h} \circ \hat{W}$. Fix $a \in A$. We must show that $\hat{W}X := \hat{W}\hat{h}(a)$ is a 1-estimator of $\mathbf{x} = h(a)$.

372 Indeed, for $\mathbf{u} \in \mathbb{S}^{d_2-1}$ we have,

$$\begin{aligned}
\mathbb{E}_X \mathbb{E}_W \langle \mathbf{u}, \hat{W}X - WX \rangle^2 &= \mathbb{E}_X \mathbb{E}_W \langle \mathbf{u}, \hat{W}X - WX \rangle^2 + 2 \langle \mathbf{u}, \hat{W}X - WX \rangle \langle \mathbf{u}, WX - W\mathbf{x} \rangle + \langle \mathbf{u}, WX - W\mathbf{x} \rangle^2 \\
&= \mathbb{E}_X \mathbb{E}_W \langle \mathbf{u}, \hat{W}X - WX \rangle^2 + 2 \mathbb{E}_X \left\langle \mathbf{u}, \overbrace{\mathbb{E}_W [\hat{W} - W]}^{=0} X \right\rangle \langle \mathbf{u}, WX - W\mathbf{x} \rangle + \mathbb{E}_X \mathbb{E}_W \langle \mathbf{u}, WX - W\mathbf{x} \rangle^2 \\
&= \mathbb{E}_X \mathbb{E}_W \langle \mathbf{u}, \hat{W}X - WX \rangle^2 + \langle \mathbf{u}, WX - W\mathbf{x} \rangle^2 \\
&= \mathbb{E}_X \mathbb{E}_W \langle \hat{W} - W, X\mathbf{u}^T \rangle^2 + \langle W^T \mathbf{u}, X - \mathbf{x} \rangle^2 \\
&\stackrel{\text{Lemma 3.8}}{\leq} \frac{2\|W\|_F^2 + 1}{k_2} \mathbb{E}_X \|X\|^2 + \frac{1}{k_1} \|W\mathbf{u}\|^2 \\
&\stackrel{(1)}{\leq} \frac{2\|W\|_F^2 + 1}{k_2} \left[\mathbb{E}_X \|X - \mathbf{x}\|^2 + \|\mathbf{x}\|^2 \right] + \frac{1}{k_1} \|W\|^2 \\
&\stackrel{(2)}{\leq} \frac{2\|W\|_F^2 + 1}{k_2} \left[\frac{1}{k_1} d_1 + \|\mathbf{x}\|^2 \right] + \frac{1}{k_1} \|W\|^2 \\
&\leq \frac{2R^2 + 1}{k_2} \left[\frac{1}{k_1} d_1 + B^2 \right] + \frac{1}{k_1} r^2
\end{aligned}$$

373 (1) We have

$$\begin{aligned}
\mathbb{E}_X \|X - \mathbf{x}\|^2 &= \mathbb{E}_X \|X\|^2 - 2 \langle X, \mathbf{x} \rangle + \|\mathbf{x}\|^2 \\
&= \mathbb{E}_X \|X\|^2 - 2 \langle \mathbb{E}_X X, \mathbf{x} \rangle + \|\mathbf{x}\|^2 \\
&= \mathbb{E}_X \|X\|^2 - \|\mathbf{x}\|^2
\end{aligned}$$

374 (2) We have

$$\begin{aligned}
\mathbb{E}_X \|X - \mathbf{x}\|^2 &= \sum_{i=1}^{d_1} \mathbb{E}(X_i - x_i)^2 \\
&= \sum_{i=1}^{d_1} \mathbb{E} \langle X - \mathbf{x}, \mathbf{e}_i \rangle^2 \\
&\leq \sum_{i=1}^{d_1} \frac{1}{k_1} = \frac{d_1}{k_1}
\end{aligned}$$

375 Finally, by choosing $k_1 = \lceil 2r^2 \rceil + 1$ and $k_2 = 2(d_1 + B^2)(2R^2 + 1)$ we get the result. \square

376 **Lemma 6.2.** Suppose that $\{X_n\}_{n=1}^\infty$ are independent σ -estimators to $\mathbf{x} \in \mathbb{R}^d$. Let $\rho(\mathbf{t}) =$
377 $\sum_{n=0}^\infty \mathbf{a}_n \mathbf{t}^n$. Let $U = \mathbf{a}_0 + \sum_{n=1}^\infty \hat{\mathbf{a}}_n Y_n$ where $Y_n = \prod_{i=1}^n X_i$ and $\hat{\mathbf{a}}_n = \frac{\mathbf{a}_n}{p_1}$ w.p. p_i and 0 otherwise.

378 Then U is σ' -estimator of $\rho(\mathbf{x})$ with $\sigma' = \sum_{n=1}^\infty \sqrt{\frac{\|\mathbf{a}_n\|_\infty^2}{p_n} ((\sigma^2 + \|\mathbf{x}\|_\infty^2)^n + (1 - p_n)d\|\mathbf{x}\|_\infty^{2n})}$.

379 **Remark 6.3.** In particular, if $\|\mathbf{a}_n\|_\infty \leq B^n$, $\sqrt{\sigma^2 + \|\mathbf{x}\|_\infty^2} \leq \frac{1}{6B}$ and $p_n = \begin{cases} 1 & n \leq \lceil \frac{\log_3(d)}{2} \rceil \\ 4^{-n} & \text{otherwise} \end{cases}$,

380 We have $\sigma' \leq 1$ and $\mathbb{E} \max\{n : \hat{\mathbf{a}}_n \neq 0\} \leq \frac{\log_3(d)+4}{2}$. Indeed,

$$\begin{aligned}
\sum_{n=1}^\infty \sqrt{\frac{\|\mathbf{a}_n\|_\infty^2}{p_n} ((\sigma^2 + \|\mathbf{x}\|_\infty^2)^n + (1 - p_n)d\|\mathbf{x}\|_\infty^{2n})} &\leq \sum_{n=1}^\infty \sqrt{\frac{\|\mathbf{a}_n\|_\infty^2}{p_n} (\sigma^2 + \|\mathbf{x}\|_\infty^2)^n} + \sum_{n=1}^\infty \sqrt{\frac{\|\mathbf{a}_n\|_\infty^2}{p_n} (1 - p_n)d\|\mathbf{x}\|_\infty^{2n}} \\
&\leq \sum_{n=1}^\infty (2B)^n \sqrt{(\sigma^2 + \|\mathbf{x}\|_\infty^2)^n} + \sqrt{d} \sum_{n=\lceil \frac{\log_3(d)}{2} \rceil + 1}^\infty (2B)^n \|\mathbf{x}\|_\infty^n \\
&\leq \sum_{n=1}^\infty \left(\frac{1}{3}\right)^n + \sqrt{d} \sum_{n=\lceil \frac{\log_3(d)}{2} \rceil + 1}^\infty \left(\frac{1}{3}\right)^n \\
&\leq \sum_{n=1}^\infty \left(\frac{1}{3}\right)^n + \sum_{n=1}^\infty \left(\frac{1}{3}\right)^n = 1
\end{aligned}$$

381 and

$$\mathbb{E} \max\{n : \hat{\mathbf{a}}_n \neq 0\} \leq \left\lceil \frac{\log_3(d)}{2} \right\rceil + \sum_{n=\left\lceil \frac{\log_3(d)}{2} \right\rceil+1}^{\infty} 4^{-n} n \leq \left\lceil \frac{\log_3(d)}{2} \right\rceil + 1$$

382 *Proof.* By lemma 3.12 it is enough to show that for all n , $\hat{\mathbf{a}}_n Y_n$ is a
383 $\sqrt{\frac{\|\mathbf{a}_n\|_{\infty}^2}{p_n}} ((\sigma^2 + \|\mathbf{x}\|_{\infty}^2)^n + (1 - p_n)d\|\mathbf{x}\|_{\infty}^{2n})$ -estimator of $\mathbf{a}_n \mathbf{x}^n$. Indeed,

$$\begin{aligned} \text{VAR}(\langle \mathbf{u}, \hat{\mathbf{a}}_n Y_n \rangle) &= \mathbb{E}(\langle \mathbf{u}, \hat{\mathbf{a}}_n Y_n \rangle - \langle \mathbf{u}, \mathbf{a}_n \mathbf{x}^n \rangle)^2 \\ &= p_n \mathbb{E} \left(\left\langle \mathbf{u}, \frac{\mathbf{a}_n}{p_n} Y_n \right\rangle - \langle \mathbf{u}, \mathbf{a}_n \mathbf{x}^n \rangle \right)^2 + (1 - p_n) \langle \mathbf{u}, \mathbf{a}_n \mathbf{x}^n \rangle^2 \\ &= \frac{1}{p_n} \mathbb{E} \langle \mathbf{u}, \mathbf{a}_n Y_n \rangle^2 - 2 \mathbb{E} \langle \mathbf{u}, \mathbf{a}_n Y_n \rangle \langle \mathbf{u}, \mathbf{a}_n \mathbf{x}^n \rangle + p_n \langle \mathbf{u}, \mathbf{a}_n \mathbf{x}^n \rangle^2 + (1 - p_n) \langle \mathbf{u}, \mathbf{a}_n \mathbf{x}^n \rangle^2 \\ &= \frac{1}{p_n} \mathbb{E} \langle \mathbf{u}, \mathbf{a}_n Y_n \rangle^2 - \langle \mathbf{u}, \mathbf{a}_n \mathbf{x}^n \rangle^2 \\ &= \frac{1}{p_n} \mathbb{E} (\langle \mathbf{a}_n \mathbf{u}, Y_n \rangle^2 - \langle \mathbf{a}_n \mathbf{u}, \mathbf{x}^n \rangle^2) + \frac{1 - p_n}{p_n} \langle \mathbf{u}, \mathbf{a}_n \mathbf{x}^n \rangle^2 \\ &\stackrel{\text{lemma 3.12}}{\leq} \frac{\|\mathbf{a}_n \mathbf{u}\|_2^2}{p_n} ((\sigma^2 + \|\mathbf{x}\|_{\infty}^2)^n + (1 - p_n)\|\mathbf{x}^n\|_2^2) \\ &\leq \frac{\|\mathbf{a}_n\|_{\infty}^2}{p_n} ((\sigma^2 + \|\mathbf{x}\|_{\infty}^2)^n + (1 - p_n)d\|\mathbf{x}\|_{\infty}^{2n}) \end{aligned}$$

384

□

385 *Proof.* (of lemma 4.13) Consider the complex function $f(z) = \frac{e^z}{1+e^z}$. It is defined in the strip
386 $\{z = x + iy : |y| < \pi\}$. By Cauchy integral formula, for any $r < \pi$, $a \in \mathbb{R}$ and $n \geq 0$,

$$f^{(n)}(a) = \frac{n!}{2\pi i} \int_{|z-a|=r} \frac{f(z)}{(z-a)^{n+1}}$$

387 It follows that

$$\left| f^{(n)}(a) \right| \leq \frac{n!}{r^n} \max_{|z-a|=r} |f(z)| \leq \frac{n!}{r^n} \max_{x+iy: |y| < r} |f(x+iy)|$$

388 Now, if $|y| < r < \frac{\pi}{2}$, we have

$$|f(x+iy)| = \frac{e^x}{|1 + e^{iy}e^x|} \leq \frac{e^x}{|1 + \cos(y)e^x|} \leq \frac{e^x}{|1 + \cos(r)e^x|} \leq \frac{1}{\cos(r)}$$

389 This implies that $\frac{e^x}{1+e^x}$ is strongly bounded. Likewise, since $\frac{e^x}{1+e^x}$ is the derivative of $\ln(1+e^x)$, the
390 function $\ln(1+e^x)$ is strongly bounded as well. □

391 *Proof.* (of lemma 4.14) Fix a set $A \subset \mathcal{X}$ of size $\leq m$. Let $\epsilon^2 = \sigma^2 = \frac{1}{72B^2}$ and note that
392 $\sqrt{\sigma^2 + \epsilon^2} \leq \frac{1}{6B}$. To generate a 1-estimator to $\rho \circ h \in \rho \circ \mathcal{H}$ on A we first describe \tilde{h} , which forms
393 the seed, such that $\forall i \in [m]$, $\|\tilde{h}(x_i) - h(x_i)\|_{\infty} \leq \epsilon$. Then, we generate σ -estimators $\hat{h}_1, \hat{h}_2, \dots$,
394 to $h|_A$. Finally, we sample Bernoulli random variables Z_1, Z_2, \dots where the parameter of Z_n is

395 $p_n = \begin{cases} 1 & n \leq \left\lceil \frac{\log_3(d)}{2} \right\rceil \\ 4^{-n} & \text{otherwise} \end{cases}$. The final estimator is

$$\hat{g}(x) = \rho(\tilde{h}(x)) + \sum_{n=1}^{\infty} \frac{\rho^{(n)}(\tilde{h}(x))}{n!} \frac{Z_n}{p_n} Y_n \text{ where } Y_n = \prod_{i=1}^n (\hat{h}_i(x) - \tilde{h}(x))$$

396 By lemma 6.2 and the following remark, \hat{g} is 1-estimator of $\rho \circ h|_A$.

397 How many bits do we need in order to specify \hat{g} ? By lemma 4.6 the restriction of $\mathcal{H}|_A$ has an
398 ϵ -cover, w.r.t. the ∞ -norm, of log-size $\lesssim n_s(m) + \frac{n(m)\log(md)}{\epsilon^2}$. So the generation of the seed \tilde{h}
399 costs $n_s(m) + \frac{n(m)\log(md)}{\epsilon^2}$ bits. We also need to specify $N := \max\{n : Z_n \neq 0\}$, Z_1, \dots, Z_N
400 and $\hat{h}_1, \dots, \hat{h}_N$. This can be done by concatenating the descriptions of the pairs (Z_n, \hat{h}_n) for
401 $n = 1, \dots, N$. The bit cost of this is bounded (in expectation) by $\frac{\log_3(d)+4}{2} (\lceil 72B^2 \rceil n(m) + 1)$ □