

---

# Deep ReLU Networks Preserve Expected Length

---

**Boris Hanin\***  
Princeton University  
bhanin@princeton.edu

**Ryan Jeong\***  
University of Pennsylvania  
rsjeong@sas.upenn.edu

**David Rolnick\***  
McGill University & Mila  
drolnick@cs.mcgill.ca

## Abstract

Assessing the complexity of functions computed by a neural network helps us understand how the network will learn and generalize. One natural measure of complexity is how the network distorts length – if the network takes a unit-length curve as input, what is the length of the resulting curve of outputs? It has been widely believed that this length grows exponentially in network depth. We prove that in fact this is not the case: the expected length distortion does not grow with depth, and indeed shrinks slightly, for ReLU networks with standard random initialization. We also generalize this result by proving upper bounds both for higher moments of the length distortion and for the distortion of higher-dimensional volumes. These theoretical results are corroborated by our experiments.

## 1 Introduction

The utility of deep neural networks ultimately arises from their ability to learn functions that are sufficiently complex to fit training data and yet simple enough to generalize well. Understanding the precise sense in which functions computed by a given network have low or high complexity is therefore important for studying when the network will perform well. Despite the fundamental importance of this question, our mathematical understanding of the functions expressed and learned by different neural network architectures remains limited.

A popular way to measure the complexity of a neural network function is to compute how it distorts lengths. This may be done by considering a set of inputs lying along a curve and measuring the length of the corresponding curve of outputs. It has been claimed in prior literature that in a ReLU network this *length distortion* grows exponentially with the network’s depth [22, 23], and this has been used as a justification of the power of deeper networks. We prove that, in fact, for networks with the typical initialization used in practice, expected length distortion does not grow at all with depth.

Our main contributions are:

1. We prove that for ReLU networks initialized with the usual  $2/\text{fan-in}$  weight variance, the expected length distortion does not grow with depth at initialization, actually decreasing slightly with depth (Thm. 3.1) and exhibiting an interesting width dependency.
2. We prove bounds on higher moments of the length distortion, giving upper bounds that hold with high probability (Thm. 4.1). We also obtain similar results for the distortion in the volume of higher-dimensional manifolds of inputs (Thm. 4.2).
3. We empirically verify that our theoretical results accurately predict observed behavior for networks at initialization, while previous bounds are loose and fail to capture subtle architecture dependencies.

It is worth explaining why our conclusions differ from those of [22, 23]. First, prior authors prove only *lower bounds* on the expected length distortion, while we use different methodology to calculate

---

\*Equal contribution

tight upper bounds, allowing us to say that the expected length distortion *does not grow* with depth. As we show in Thm. 5.1 and Fig. 1, the prior bounds are in fact quite loose.

Second, Thm. 3(a) in [23] has a critical typo<sup>2</sup>, which has unfortunately been perpetuated in Thm. 1 of [22]. Namely, the “2” was omitted in the following statement (paraphrased from the original):

$$\mathbb{E}[\text{length distortion}] = \Omega \left[ \left( \frac{\sigma_w \sqrt{\text{width}}}{2\sqrt{\text{width} + 1}} \right)^{\text{depth}} \right],$$

where  $\sigma_w^2/\text{width}$  is the variance of the weight distribution. As we prove in Thm. 3.1, leaving out the 2 makes the statement false, incorrectly suggesting that length distortion explodes with depth for the standard He initialization  $\sigma_w = \sqrt{2}$ .

Finally, prior authors drew the conclusion that length distortion grows exponentially with depth by considering the behavior of ReLU networks with unrealistically large weights. If one multiplies by  $C$  the weights and biases of a ReLU network, one multiplies the length distortion by  $C^{\text{depth}}$ , so it should come as no surprise that there exist settings of the weights for which the distortion grows exponentially with depth (or where it decays exponentially). The value of our results comes in analyzing the behavior specifically at He initialization ( $\sigma_w = \sqrt{2}$ ) [17]. This initialization is the one used in practice for ReLU networks, since this is the weight variance that must be used if the outputs [14] and gradients [11] are to remain well-controlled. In the present work, we show that this is also the correct initialization for the expected length distortion to remain well-behaved.

## 2 Related Work

A range of complexity measures for functions computed by deep neural networks have been considered in the literature, dividing the prior work into at least three categories. In the first, the emphasis is on *worst-case* (or *best-case*) scenarios – what is the maximal possible complexity of functions computed by a given network architecture. These works essentially study the expressive power of neural networks and often focus on showing that deep networks are able to express functions that cannot be expressed by shallower networks. For example, it has been shown that it is possible to set the weights of a deep ReLU network such that the number of linear regions computed by the network grows exponentially in the depth [7, 9, 20, 26, 27]. Other works consider the degree of polynomials approximable by networks of different depths [19, 24] and the topological invariants of networks [5].

While such work has sometimes been used to explain the utility of different neural network architectures (especially deeper ones), a second strand of prior work has shown that a significant gap can exist between the functions expressible by a given architecture and those which may be learned in practice. Such *average-case* analyses have indicated that some readily expressible functions are provably difficult to learn [25] or vanishingly unlikely for random networks [13, 15, 16], and that some functions learned more easily by deep architectures are nonetheless expressible by shallow ones [3]. (While here we consider neural nets with ReLU activation, it is worth noting that for arithmetic circuits, worst-case and average-case scenarios may be more similar – in both scenarios, a significant gap exists between the matricization rank of the functions computed by deep and shallow architectures [6].) As noted earlier, average-case analyses in [22, 23] provided lower bounds on expected length distortion, while [21] presented a similar analysis for the curvature of output trajectories.

Finally, a number of authors have sought complexity measures that either empirically or theoretically correlate with generalization [18]. Such measures have been based on classification margins [4], network compressibility [2], and PAC-Bayes considerations [8].

## 3 Expected Length Distortion

### 3.1 Motivation

While neural networks are typically overparameterized and trained with little or no explicit regularization, the functions they learn in practice are often able to generalize to unseen data. This phenomenon indicates an implicit regularization that causes these learned functions to be surprisingly simple.

<sup>2</sup>We have confirmed this in personal correspondence with the authors, and it is simple to verify – the typo arises in the last step of the proof given in the Supplementary Material of [23].

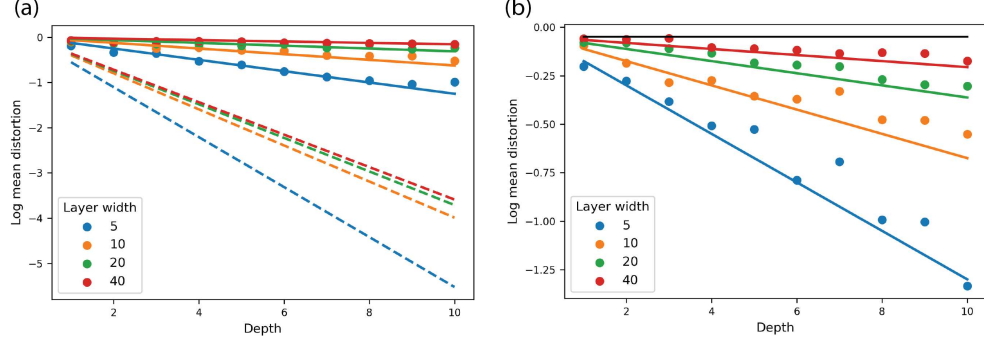


Figure 1: Mean length distortion as a function of depth, for randomly initialized ReLU networks of varying architecture. As described in Thm. 3.1, distortion not only fails to grow, but shrinks with depth, especially for networks of small width. (a) compares the predictions of our Thm. 5.1 (solid lines) to the lower bounds proven in prior work [23] (dashed lines) and the true empirical means (colored dots), which closely track our predictions. (b) zooms in on the upper part of (a), showing the empirical mean length distortion as a function of depth for different widths. The horizontal black line is  $y = \frac{\Gamma(3)}{\Gamma(2.5)\sqrt{2.5}}$ , the mean length distortion we predict when  $n_0 = n_L$  in the limit of infinite width for any fixed depth. In each network, width is constant in all hidden layers, while input and output dimension are both 5. The input curve is a fixed line segment of unit length. For each architecture, length distortion is calculated for 500 different initializations of the weights and biases of the network (with variance of the weights given by  $2/\text{fan-in}$ ). For further experimental details, see Appendix B.

Why this occurs is not well-understood theoretically, but there is a high level intuition. In a nutshell, randomly initialized neural networks will compute functions of low complexity. Moreover, as optimization by first order methods is a kind of greedy local search, network training is attracted to minima of the loss that are not too far from the initialization and hence will still be well-behaved.

While this intuition is compelling, a key challenge in making it rigorous is to devise appropriate notions of complexity that are small throughout training. The present work is intended to provide a piece of the puzzle, making precise the idea that, at the start of training, neural networks compute tame functions. Specifically, we demonstrate that neural networks at initialization have low distortion of length and volume, as defined in §3.2.

An important aspect of our analysis is that we study networks in typical, average-case scenarios. A randomly initialized neural network could, in principle, compute any function expressible by a network of that architecture, since the weights might with low probability take on any set of values. Some settings of weights will lead to functions of high complexity, but these settings may be unlikely to occur in practice, depending on the distribution over weights that is under consideration. As prior work has emphasized [22, 23], atypical distributions of weights can lead to exponentially high length distortion. We show that, in contrast, for deep ReLU networks with the standard initialization, the functions computed have low distortion in expectation and with high probability.

### 3.2 Definitions

Let  $L \geq 1$  be a positive integer and fix positive integers  $n_0, \dots, n_L$ . We consider a fully connected feed-forward ReLU network  $\mathcal{N}$  with input dimension  $n_0$ , output dimension  $n_L$ , and hidden layer widths  $n_1, \dots, n_{L-1}$ . Suppose that the weights and biases of  $\mathcal{N}$  are *independent* and Gaussian with the weights  $W_{ij}^{(\ell)}$  between layers  $\ell - 1$  and  $\ell$  and biases  $b_j^{(\ell)}$  in layer  $\ell$  satisfying:

$$W_{ij}^{(\ell)} \sim G(0, 2/n_{\ell-1}), \quad b_j^{(\ell)} \sim G(0, C_b). \quad (3.1)$$

Here,  $C_b > 0$  is any fixed constant and  $G(\mu, \sigma^2)$  denotes a Gaussian with mean 0 and variance  $\sigma^2$ . For any  $M \subseteq \mathbb{R}^{n_0}$ , we denote by  $\mathcal{N}(M) \subseteq \mathbb{R}^{n_L}$  the (random) image of  $M$  under the map  $x \mapsto \mathcal{N}(x)$ . Our primary object of the study will be the size of the output  $\mathcal{N}(M)$  relative to that of

$M$ . Specifically, when  $M$  is a 1-dimensional curve, we define

$$\text{length distortion} = \frac{\text{len}(\mathcal{N}(M))}{\text{len}(M)}.$$

Note that while *a priori* this random variable depends on  $M$ , we will find that its statistics depend only on the architecture of  $\mathcal{N}$  (see Thms. 3.1 and 5.1).

### 3.3 Results

We prove that the expected length distortion does not grow with depth – in fact, it slightly decreases. Our result may be informally stated thus:

**Theorem 3.1** (Length distortion: Mean). *Consider a ReLU network of depth  $L$ , input dimension  $n_0$ , output dimension  $n_L$ , and hidden layer widths  $n_\ell$ , with weights given by standard He normal initialization [17]. The expected length distortion is upper bounded by  $\sqrt{n_L/n_0}$ . More precisely:*

$$\mathbb{E}[\text{length distortion}] \approx C \left( \frac{n_L}{n_0} \right)^{1/2} \exp \left[ -\frac{5}{8} \sum_{\ell=1}^{L-1} \frac{1}{n_\ell} \right],$$

where  $C$  is a constant close to 1 depending only on the output dimension  $n_L$ .

For a formal statement (including the exact values of constants) and proof, see Thm. 5.1 and App. C.




Figure 2: Mean length distortion as a function of the ratio of output to input dimension, for ReLU networks with various architectures (e.g. [10, 10, 10] denotes three hidden layers, each of width 10). All networks are randomly initialized as in Figure 1. We sample 100 pairs of input dimension  $n_0$  and output dimension  $n_L$ , each at most 50, such that the ratio of output to input dimension is distinct for each such pair. For each pair, 200 different network initializations are tested and the resulting length distortion is calculated; the log of the empirical mean is plotted. The dashed black line plots  $\log(y) = \frac{1}{2} \log(x)$ , the approximate prediction by Theorem 3.1.

Note that our theorem also applies for initializations other than He normal; the result in such cases is simply less interesting. Suppose we initialize the weights in each layer  $\ell$  as i.i.d. normal with variance  $2c^2/n_{\ell-1}$  instead of  $2/n_{\ell-1}$ . This is equivalent to taking a He normal initialization and multiplying the weights by  $c$ . Multiplying the weights and biases in a depth- $L$  ReLU network by  $c$  simply multiplies the output (and hence the length distortion) by  $c^L$ , so the expected distortion is  $c^L$  times the value given in Thm. 3.1. This emphasizes why an exponentially large distortion necessarily occurs if one sets the weights too large.

Our experiments align with the theorem. In Figure 1, we compute the empirical mean of the length distortion for randomly initialized deep ReLU networks of various architectures with fixed input and output dimension. The figure confirms three of our key predictions: (1) the expected length distortion does not grow with depth, and actually decreases slightly; (2) the decrease happens faster for narrower networks (since  $1/n_\ell$  is larger for smaller  $n_\ell$ ); and (3) for equal input and output dimension  $n_0 = n_L$ , there is an upper bound of 1 (in fact, the tighter upper bound of  $C$  is approximately 0.9515 for  $n_L = 5$ , as shown in the figure). The figure also shows the prior bounds in [23] to be quite loose. For further experimental details, see Appendix B.

In Fig. 2, we instead fix the set of hidden layer widths, but vary input and output dimensions. The results confirm that indeed the expected length distortion grows as the square root of the ratio of output to input dimension.

### 3.4 Intuitive explanation

The purpose of this section is to give an intuitive but technical explanation for why, in Theorem 3.1, we find that the distortion  $\text{len}(\mathcal{N}(M))/\text{len}(M)$  of the length of a 1D curve  $M$  under the neural network  $\mathcal{N}$  is typically not much larger than 1, even for deep networks.

Our starting point is that, near a given point  $x \in M$ , the length of the image under  $\mathcal{N}$  of a small portion of  $M$  with length  $dx$  is approximately given by  $\|J_x u\| dx$ , where  $J_x$  is the input-output Jacobian of  $\mathcal{N}$  at the input  $x$  and  $u$  is the unit tangent vector to  $M$  at  $x$ . In fact, in Lemma C.1 of the Supp. Material, we use a simple argument to give upper bounds on the moments of  $\text{len}(\mathcal{N}(M))/\text{len}(M)$  in terms of the moments of the norm  $\|J_x u\|$  of the Jacobian-vector product  $J_x u$ .




Figure 3: Length distortion as a function of  $\sum_{\ell=1}^L n_{\ell}^{-1}$ , showing both mean and standard deviation across initializations. We test several types of network architecture – with constant width 20, alternating between widths 30 and 10, and with the first (respectively, second) half of the layers of width 30 and the rest width 10. Each architecture type is tested for several depths. For each such network, we use  $n_0 = n_L = 5$  and compute length distortion for 100 initializations on a fixed line segment. As predicted in Thm. 4.1, the mean length distortion decreases with the sum of width reciprocals. Empirical standard deviation does not, in general, increase with greater depth, remaining modest throughout, as is consistent with the upper bound on variance in Thm. 4.1.

$$\mathbb{E}[\|J_{\ell,x} e_1\|] = 1 + O(n_{\ell}^{-1}).$$

This is a consequence of initializing weights to have variance  $2/\text{fan-in}$ . Put together, the two preceding displayed equations suggest writing

$$\|J_x u\| = \exp \left[ \sum_{\ell=1}^L \log(\|J_{\ell,x} e_1\|) \right]$$

The terms being summed in the exponent are independent and the argument of each logarithm scales like  $1 + O(n_{\ell}^{-1})$ . With probability  $1 - \delta$  we have for all  $\ell$  that  $\log(\|J_{\ell,x} e_1\|) \leq c n_{\ell}^{-1}$ , where  $c$  is a constant depending only on  $\delta$ . Thus, all together,

$$\|J_x u\| \leq \exp \left[ c \sum_{\ell=1}^L n_{\ell}^{-1} \right]$$

with high probability. In particular, in the simple case where  $n_{\ell}$  are proportional to a single large constant  $n$ , we find the typical size of  $\|J_x u\|$  is exponential in  $L/n$  rather than exponential in  $L$ , as in the worst case. If  $\mathcal{N}$  is wider than it is deep, so that  $L/n$  is bounded above, the typical size of  $\|J_x u\|$  and hence of  $\text{len}(\mathcal{N}(M))/\text{len}(M)$  remains bounded. Theorem 5.1 makes this argument precise. Moreover, let us note that articles like [7, 10, 21] show low distortion of angles and volumes in very wide networks. Our results, however, apply directly to more realistic network widths.

Thus, upper bounds on the length distortion reduce to studying the Jacobian  $J_x$ , which in a network with  $L$  hidden layers can be written as a product  $J_x = J_{L,x} \cdots J_{1,x}$  of the layer  $\ell - 1$  to layer  $\ell$  Jacobians  $J_{\ell,x}$ . In a worst-case analysis, the left singular vectors of  $J_{\ell,x}$  would align with the right singular vectors of  $J_{\ell+1,x}$ , causing the resulting product  $J_x$  to have a largest singular value that grows exponentially with  $L$ . However, at initialization, this is provably not the case with high probability. Indeed, the Jacobians  $J_{\ell,x}$  are independent (see Lemma C.3) and their singular vectors are therefore incoherent. We find in particular (Lemma C.2) the following equality in distribution:

$$\|J_x u\| \stackrel{d}{=} \prod_{\ell=1}^L \|J_{\ell,x} e_1\|,$$

where  $e_1$  is the first standard unit vector and the terms in the product are independent. On average each term in this product is close to 1:

## 4 Further Results

### 4.1 Higher moments

We have seen that the mean length distortion is upper-bounded by 1, and indeed by a function of the architecture that decreases with the depth of the network. However, a small mean is insufficient by itself to ensure that typical networks have low length distortion. For this, we now show that in fact all moments of the length distortion are well-controlled. Specifically, the variance is bounded above by the ratio of output to input dimension, and higher moments are upper-bounded by a function that grows very slowly with depth. Our results may be informally stated thus:

**Theorem 4.1** (Length distortion: Higher moments). *Consider, as before, a ReLU network of depth  $L$ , input dimension  $n_0$ , output dimension  $n_L$ , and hidden layer widths  $n_\ell$ , with weights given by He normal initialization. We have the following bounds on higher moments of the length distortion:*

$$\text{Var}[\text{length distortion}] \leq \frac{n_L}{n_0} \quad \text{and} \quad \mathbb{E}[(\text{length distortion})^m] \leq \left(\frac{n_L}{n_0}\right)^{\frac{m}{2}} \exp \left[ cm^2 \sum_{\ell=1}^L n_\ell^{-1} \right]$$

for some universal constant  $c > 0$ .

A formal statement and proof of this result are given in Theorem 5.1 and Appendix C.

We consider the mean and variance of length distortion for a wide range of network architectures in Figure 3. The figure confirms that variance is modest for all architectures and does not increase with depth, and that mean length distortion is consistently decreasing with the sum of layer width reciprocals, as predicted by Theorem 4.1.

### 4.2 Higher-dimensional volumes

Another natural generalization to consider is how ReLU networks distort higher-dimensional volumes: Given a  $d$ -dimensional input  $M$ , the output  $\mathcal{N}(M)$  will in general also be  $d$ -dimensional, and we can therefore consider the *volume distortion*  $\text{vol}_d(\mathcal{N}(M))/\text{vol}_d(M)$ . The theorems presented in §3.3 when  $d = 1$  can be extended to all  $d$ , and the results may be informally stated thus:

**Theorem 4.2** (Volume distortion). *Consider a ReLU network of depth  $L$ , input dimension  $n_0$ , output dimension  $n_L$ , and hidden layer widths  $n_\ell$ , with weights given by He normal initialization. Both the squared mean and the variance of volume distortion are upper-bounded by:*

$$\left(\frac{n_L}{n_0}\right)^d \exp \left[ -\binom{d}{2} \sum_{\ell=1}^L n_\ell^{-1} \right].$$

A formal statement and proof of this result are given in Theorem 5.2 and Appendix D.

## 5 Formal Statements

In this section, we provide formal statements of our results. Full proofs are given in the Supp. Material.

### 5.1 One-dimensional manifolds

Fix a smooth one-dimensional manifold  $M \subseteq \mathbb{R}^{n_0}$  (i.e. a curve) with unit length. Each point on  $M$  represents a possible input to our ReLU network  $\mathcal{N}$ , which we recall has input dimension  $n_0$ , output dimension  $n_L$ , and hidden layer widths  $n_1, \dots, n_{L-1}$ . The function  $x \mapsto \mathcal{N}(x)$  computed by  $\mathcal{N}$  is continuous and piecewise linear. Thus, the image  $\mathcal{N}(M) \subseteq \mathbb{R}^{n_L}$  of  $M$  under this map is a piecewise smooth curve in  $\mathbb{R}^{n_L}$  with a finite number of pieces, and its length  $\text{len}(\mathcal{N}(M))$  is a random variable. Our first result states that, provided  $\mathcal{N}$  is wider than it is deep, this distortion is small with high probability at initialization.

**Theorem 5.1.** *Let  $\mathcal{N}$  be a fully connected ReLU network with input dimension  $n_0$ , output dimension  $n_L$ , hidden layer widths  $n_1, \dots, n_{L-1}$ , and independent centered Gaussian weights/biases with the*

weights having variance 2/fan-in (as in (3.1)). Let  $M$  be a 1-dimensional curve of unit length in  $\mathbb{R}^{n_0}$ . Then, the mean length  $\mathbb{E}[\text{len}(\mathcal{N}(M))]$  equals

$$\left(\frac{n_L}{n_0}\right)^{1/2} \frac{\Gamma\left(\frac{n_L+1}{2}\right)}{\Gamma\left(\frac{n_L}{2}\right) \left(\frac{n_L}{2}\right)^{1/2}} \times \exp\left[-\frac{5}{8} \sum_{\ell=1}^{L-1} n_\ell^{-1} + O\left(\sum_{\ell=1}^{L-1} n_\ell^{-2}\right)\right], \quad (5.1)$$

where implied constants in  $O(\cdot)$  are universal and  $\Gamma(\cdot)$  denotes the Gamma function. Moreover:

$$\text{Var}[\text{len}(\mathcal{N}(M))] \leq \frac{n_L}{n_0}. \quad (5.2)$$

Finally, there exist universal constants  $c_1, c_2 > 0$  such that if  $m < c_1 \min\{n_1, \dots, n_{L-1}\}$ , then

$$\mathbb{E}[(\text{len}(\mathcal{N}(M)))^m] \leq \left(\frac{n_L}{n_0}\right)^{m/2} \exp\left[c_2 m^2 \sum_{\ell=1}^L n_\ell^{-1}\right]. \quad (5.3)$$

Theorem 5.1 is proved in §C. Several comments are in order. First, we have assumed for simplicity that  $M$  has unit length. For curves of arbitrary length, both the equality (5.1) and the bounds (5.2) and (5.3) all hold and are derived in the same way provided  $\text{len}(\mathcal{N}(M))$  is replaced by the distortion  $\text{len}(\mathcal{N}(M))/\text{len}(M)$  per unit length.

Second, in the expression (5.1), note that the exponent tends to 0 as  $n_\ell \rightarrow \infty$  for any fixed  $L$ . More precisely, when  $n_\ell = n$  is large and constant, it scales as  $-5L/8n$ . This shows that, for fixed  $n_0, n_L, n$ , the mean  $\mathbb{E}[\text{len}(\mathcal{N}(M))]$  is actually decreasing with  $L$ . This somewhat surprising phenomenon is borne out in Fig. 1 and is a consequence of the fact that wide fully connected ReLU nets are strongly contracting (see e.g. Thms. 3 and 4 in [10]).

Third, the pre-factor  $(n_L/n_0)^{1/2}$  encodes the fact that, on average over initialization, for any vector of inputs  $x \in \mathbb{R}^{n_0}$ , we have (see e.g. Cor. 1 in [14])

$$\mathbb{E}[\|\mathcal{N}(x)\|^2] \approx \frac{n_L}{n_0} \|x\|^2, \quad (5.4)$$

where the approximate equality becomes exact if the bias variance  $C_b$  is set to 0 (see (3.1)). In other words, at the start of training, the network  $\mathcal{N}$  rescales inputs by an average factor  $(n_L/n_0)^{1/2}$ . This overall scaling factor also dilates  $\text{len}(\mathcal{N}(M))$  relative to  $\text{len}(M)$ .

Fourth, we stated Theorem 5.1 for Gaussian weights and biases (see (3.1)), but the result holds for any weight distribution that is symmetric around 0, has variance 2/fan-in, and has finite higher moments. The only difference is that the constant 5/8 may need to be slightly enlarged (e.g. around (C.6)).

Finally, all the estimates (5.1), (5.2), and (5.3) are consistent with the statement that, up to the scaling factor  $(n_L/n_0)^{1/2}$ , the random variable  $\text{len}(\mathcal{N}(M))$  is bounded above by a log-normal distribution:

$$\text{len}(\mathcal{N}(M)) \leq \left(\frac{n_L}{n_0}\right)^{1/2} \exp\left[\frac{1}{2} G\left(\frac{\beta}{2}, \beta\right)\right],$$

where  $\beta = c \sum_{\ell=1}^L n_\ell^{-1}$  and  $c$  is a fixed constant.

## 5.2 Higher-dimensional manifolds

The results in the previous section generalize to the case when  $M$  has higher dimension. To consider this case, fix a positive integer  $d \leq \min\{n_0, \dots, n_L\}$  and a smooth  $d$ -dimensional manifold  $M \subseteq \mathbb{R}^{n_0}$  of unit volume  $\text{vol}_d(M) = 1$ . Note that if  $d > \min\{n_0, \dots, n_L\}$ , then  $\mathcal{N}(M)$  is at most  $(d-1)$ -dimensional and its  $d$ -dimensional volume vanishes. Its image  $\mathcal{N}(M) \subseteq \mathbb{R}^{n_L}$  is a piecewise smooth manifold of dimension at most  $d$ . The following result, proved in §D, gives an upper bound on the typical value of the  $d$ -dimensional volume of  $\mathcal{N}(M)$ .

**Theorem 5.2.** *Let  $\mathcal{N}$  be a fully connected ReLU network with input dimension  $n_0$ , output dimension  $n_L$ , hidden layer widths  $n_1, \dots, n_{L-1}$  and independent centered Gaussian weights/biases with the variance of the weights given by 2/fan-in (see (3.1)). Let  $M$  be a  $d$ -dimensional smooth submanifold of  $\mathbb{R}^{n_0}$  with unit volume and  $d \leq \min\{n_0, \dots, n_L\}$ . Then, both the squared mean and the variance of the  $d$ -dimensional volume  $\text{vol}_d(\mathcal{N}(M))$  of  $\mathcal{N}(M)$  is bounded above by*

$$\left(\frac{n_L}{n_0}\right)^d \exp\left[-\binom{d}{2} \sum_{\ell=1}^L n_\ell^{-1}\right] \quad (5.5)$$

## 6 Proof Sketch

The purpose of this section is to explain the main steps to obtaining the mean and variance estimates (5.1) and (5.2) from Theorem 5.1. In several places, we will gloss over some mathematical subtleties that we deal with in the detailed proof of Theorem 5.1 given in Appendix C. We will also content ourselves here with proving a slightly weaker estimate than (5.1) and show instead simply that

$$\mathbb{E} [\text{len}(\mathcal{N}(M))] \leq \left( \frac{n_L}{n_0} \right)^{1/2} \quad \text{and} \quad \text{Var} [\text{len}(\mathcal{N}(M))] \leq \frac{n_L}{n_0}. \quad (6.1)$$

We refer the reader to Appendix C for a derivation of the more refined result stated in Theorem 5.1. Since we have  $\mathbb{E}[X]^2 \leq \mathbb{E}[X^2]$  and  $\text{Var}[X] \leq \mathbb{E}[X^2]$ , both estimates in (6.1) follow from

$$\mathbb{E} [\text{len}(\mathcal{N}(M))^2] \leq \frac{n_L}{n_0}. \quad (6.2)$$

To obtain this bound, we begin by relating moments of  $\text{len}(\mathcal{N}(M))$  to those of the input-output Jacobian of  $\mathcal{N}$  at a single point for which we need some notation. Namely, let us choose a unit speed parameterization of  $M$ ; that is, fix a smooth mapping

$$\gamma : \mathbb{R} \rightarrow \mathbb{R}^{n_0}, \quad \gamma(t) = (\gamma_1(t), \dots, \gamma_{n_0}(t))$$

for which  $M$  is the image under  $\gamma$  of the interval  $[0, 1]$ . That  $\gamma$  has unit speed means that for every  $t \in [0, 1]$  we have  $\|\gamma'(t)\| = 1$ , where  $\gamma'(t) := (\gamma'_1(t), \dots, \gamma'_{n_0}(t))$ . Then, the mapping  $\Gamma := \mathcal{N} \circ \gamma$  (for  $\Gamma : \mathbb{R} \rightarrow \mathbb{R}^{n_L}$ ) gives a parameterization of the curve  $\mathcal{N}(M)$ . Note that this parameterization is not unit speed. Rather the length of  $\mathcal{N}(M)$  is given by

$$\text{len}(\mathcal{N}(M)) = \int_0^1 \|\Gamma'(t)\| dt. \quad (6.3)$$

Intuitively, the integrand  $\|\Gamma'(t)\| dt$  computes, at a given  $t$ , the length of  $\Gamma([t, t + dt])$  as  $dt \rightarrow 0$ . The following Lemma uses (6.3) to bound the moments of  $\text{len}(\mathcal{N}(M))$  in terms of the moments of the norm of the input-output Jacobian  $J_x$ , defined for any  $x \in \mathbb{R}^{n_0}$  by

$$J_x := \left( \frac{\partial \mathcal{N}_i}{\partial x_j}(x) \right)_{\substack{1 \leq i \leq n_L \\ 1 \leq j \leq n_0}}, \quad (6.4)$$

where  $\mathcal{N}_i$  is the  $i$ -th component of the network output.

**Lemma 6.1.** *We have*

$$\mathbb{E} [\text{len}(\mathcal{N}(M))^2] \leq \int_0^1 \mathbb{E} [\|J_{\gamma(t)} \gamma'(t)\|^2] dt. \quad (6.5)$$

*Sketch of Proof.* Taking powers in (6.3) and interchanging the expectation and integrals, we obtain

$$\mathbb{E} [\text{len}(\mathcal{N}(M))^2] = \int_0^1 \int_0^1 \mathbb{E} [\|\Gamma'(t_1)\| \|\Gamma'(t_2)\|] dt_1 dt_2. \quad (6.6)$$

Applying the inequality  $ab \leq \frac{1}{2}(a^2 + b^2)$ , for  $a, b \in \mathbb{R}$ , to the integrand in (6.6), we conclude

$$\mathbb{E} [\text{len}(\mathcal{N}(M))^2] \leq \int_0^1 \mathbb{E} [\|\Gamma'(t)\|^2] dt. \quad (6.7)$$

The chain rule yields  $\Gamma'(t) = J_{\gamma(t)} \gamma'(t)$ . Substituting this into (6.7) completes the proof.  $\square$

Lemma 6.1, while elementary, appears to be new, allowing us to bound global quantities such as length in terms of local quantities such as the moments of  $\|J_x u\|$ . In particular, having obtained the expression (6.5), we have reduced bounding the second moment of  $\text{len}(\mathcal{N}(M))$  to bounding the second moment of the random variable  $\|J_x u\|$  for a fixed  $x \in \mathbb{R}^{n_0}$  and unit vector  $u \in \mathbb{R}^{n_0}$ . This is a simple matter since in our setting the distribution of weights and biases is Gaussian and hence, in distribution,  $\|J_x u\|$  is equal to a product of independent random variables:



**Proposition 6.2.** *For any  $x \in \mathbb{R}^{n_0}$  and any unit vector  $u \in \mathbb{R}^{n_0}$ , the random variable  $\|J_x u\|^2$  is equal in distribution to a product of independent scaled chi-squared random variables*

$$\|J_x u\|^2 \stackrel{d}{=} \frac{n_L}{n_0} \left( \prod_{\ell=1}^{L-1} \frac{2}{n_\ell} \chi_{\mathbf{n}_\ell}^2 \right) \cdot \frac{1}{n_L} \chi_{n_L}^2,$$

where the number of degrees of freedom in the  $\ell$ th term of the product (for  $\ell = 1, \dots, L-1$ ) is given by an independent binomial  $\mathbf{n}_\ell \stackrel{d}{=} \text{Bin}(n_\ell, 1/2)$  with  $n_\ell$  trials and success probability  $1/2$ . The number of degrees of freedom in the final term is deterministic and given by  $n_L$ .

This Proposition has been derived a number of times in the literature (see Theorem 3 in [11] and Fact 7.2 in [1]). We provide an alternative proof based on Proposition 2 in [12] in the Supplementary Material. Note that the distribution of  $\|J_x u\|$  is the same at every  $x$  and  $u$ . Thus, fixing  $x \in \mathbb{R}^{n_0}$  and a unit vector  $u \in \mathbb{R}^{n_0}$ , we find

$$\mathbb{E} [\text{len}(\mathcal{N}(M))^2] \leq \mathbb{E} [\|J_x u\|^2].$$

To prove (6.2), we note that  $\mathbb{E} [\chi_{\mathbf{n}}^2] = n/2$  and apply Lemma 6.2 to find, as desired,

$$\mathbb{E} [\|J_x u\|^2] = \frac{n_L}{n_0} \left( \prod_{\ell=1}^{L-1} \mathbb{E} \left[ \frac{2}{n_\ell} \chi_{\mathbf{n}_\ell}^2 \right] \right) \mathbb{E} [n_L^{-1} \chi_{n_L}^2] = \frac{n_L}{n_0} \prod_{\ell=1}^{L-1} \left\{ \frac{2}{n_\ell} \mathbb{E} [\mathbf{n}_\ell] \right\} = \frac{n_L}{n_0}.$$

□

## 7 Conclusion

We show that deep ReLU networks with appropriate initialization do not appreciably distort lengths and volumes, contrary to prior assumptions of exponential growth. Specifically, we provide an exact expression for the mean length distortion, which is bounded above by 1 and decreases slightly with increasing depth. We also prove that higher moments of the length distortion admit well-controlled upper bounds, and generalize this to distortion of higher dimensional volumes. We show empirically that our theoretical results closely match observations, unlike previous loose lower bounds.

There remain several promising directions for future work. First, we have proved statements for networks at initialization, and our results do not necessarily hold after training; analyzing such behavior formally would require consideration of the loss function, optimizer, and data distribution. Second, our results, as stated, do not apply to convolutional, residual, or recurrent networks. We envision generalizations for networks with skip connections being straightforward, while formulations for convolutional and recurrent networks will likely be more complicated. In Appendix A, we provide preliminary results suggesting that expected length distortion decreases modestly with depth in both convolutional networks and those with skip connections. Third, we believe that various other measures of complexity for neural networks, such as curvature, likely demonstrate similar behavior to length and volume distortion in average-case scenarios with appropriate initialization. While somewhat parallel results for counting linear regions were shown in [15, 16], there remains much more to be understood for other notions of complexity. Finally, we look forward to work that explicitly ties properties such as low length distortion to improved generalization, as well as new learning algorithms that leverage this line of theory in order to better control inductive biases to fit real-world tasks.

## References

- [1] Zeyuan Allen-Zhu, Yuanzhi Li, and Zhao Song. A convergence theory for deep learning via over-parameterization. In *International Conference on Machine Learning (ICML)*, 2019.
- [2] Sanjeev Arora, Rong Ge, Behnam Neyshabur, and Yi Zhang. Stronger generalization bounds for deep nets via a compression approach. In *International Conference on Machine Learning (ICML)*, 2018.
- [3] Lei Jimmy Ba and Rich Caruana. Do deep nets really need to be deep? In *Conference on Neural Information Processing Systems (NeurIPS)*, 2014.

- [4] Peter Bartlett, Dylan J Foster, and Matus Telgarsky. Spectrally-normalized margin bounds for neural networks. *Preprint arXiv:1706.08498*, 2017.
- [5] Monica Bianchini and Franco Scarselli. On the complexity of neural network classifiers: A comparison between shallow and deep architectures. *IEEE Transactions on Neural Networks and Learning Systems*, 25(8):1553–1565, 2014.
- [6] Nadav Cohen, Or Sharir, and Amnon Shashua. On the expressive power of deep learning: A tensor analysis. In *Conference on Learning Theory (COLT)*, 2016.
- [7] Amit Daniely. Depth separation for neural networks. In *Conference on Learning Theory (COLT)*, 2017.
- [8] Gintare Karolina Dziugaite and Daniel M Roy. Computing nonvacuous generalization bounds for deep (stochastic) neural networks with many more parameters than training data. *Preprint arXiv:1703.11008*, 2017.
- [9] Ronen Eldan and Ohad Shamir. The power of depth for feedforward neural networks. In *Conference on Learning Theory (COLT)*, 2016.
- [10] Raja Giryes, Guillermo Sapiro, and Alex M Bronstein. Deep neural networks with random Gaussian weights: A universal classification strategy? *IEEE Transactions on Signal Processing*, 64(13):3444–3457, 2016.
- [11] Boris Hanin. Which neural net architectures give rise to exploding and vanishing gradients? In *Conference on Neural Information Processing Systems (NeurIPS)*, 2018.
- [12] Boris Hanin and Mihai Nica. Products of many large random matrices and gradients in deep neural networks. *Communications in Mathematical Physics*, pages 1–36, 2019.
- [13] Boris Hanin and Mihai Nica. Finite depth and width corrections to the neural tangent kernel. In *International Conference on Learning Representations (ICLR)*, 2020.
- [14] Boris Hanin and David Rolnick. How to start training: The effect of initialization and architecture. In *Conference on Neural Information Processing Systems (NeurIPS)*, 2018.
- [15] Boris Hanin and David Rolnick. Complexity of linear regions in deep networks. In *International Conference on Machine Learning (ICML)*, 2019.
- [16] Boris Hanin and David Rolnick. Deep ReLU networks have surprisingly few activation patterns. In *Conference on Neural Information Processing Systems (NeurIPS)*, 2020.
- [17] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Delving deep into rectifiers: Surpassing human-level performance on ImageNet classification. In *IEEE International Conference on Computer Vision (ICCV)*, 2015.
- [18] Yiding Jiang, Behnam Neyshabur, Hossein Mobahi, Dilip Krishnan, and Samy Bengio. Fantastic generalization measures and where to find them. *Preprint arXiv:1912.02178*, 2019.
- [19] Henry W Lin, Max Tegmark, and David Rolnick. Why does deep and cheap learning work so well? *Journal of Statistical Physics*, 168(6):1223–1247, 2017.
- [20] Guido Montúfar, Razvan Pascanu, Kyunghyun Cho, and Yoshua Bengio. On the number of linear regions of deep neural networks. In *Conference on Neural Information Processing Systems (NeurIPS)*, 2014.
- [21] Ben Poole, Subhaneil Lahiri, Maithra Raghu, Jascha Sohl-Dickstein, and Surya Ganguli. Exponential expressivity in deep neural networks through transient chaos. In *Conference on Neural Information Processing Systems (NeurIPS)*, 2016.
- [22] Ilan Price and Jared Tanner. Trajectory growth lower bounds for random sparse deep ReLU networks. *Preprint arXiv:1911.10651*, 2019.
- [23] Maithra Raghu, Ben Poole, Jon Kleinberg, Surya Ganguli, and Jascha Sohl-Dickstein. On the expressive power of deep neural networks. In *International Conference on Machine Learning (ICML)*, 2017.
- [24] David Rolnick and Max Tegmark. The power of deeper networks for expressing natural functions. In *International Conference on Learning Representations*, 2018.
- [25] Shai Shalev-Shwartz, Ohad Shamir, and Shaked Shammah. Failures of gradient-based deep learning. In *International Conference on Machine Learning (ICML)*, 2017.

- [26] Matus Telgarsky. Representation benefits of deep feedforward networks. *Preprint arXiv:1509.08101*, 2015.
- [27] Matus Telgarsky. Benefits of depth in neural networks. In *Conference on Learning Theory (COLT)*, 2016.

## A Experiments for Convolutional and Residual Networks

We here provide preliminary experimental results for the mean length distortion in networks with convolutional layers and skip connections.

In Figure 4(a), we show results for networks with sequential convolutional layers (without pooling layers), initialized as before with weights i.i.d. normal with variance  $2/\text{fan-in}$ , and a final fully connected layer. We use input dimension  $n_0 = 16 \times 16 \times 3 = 768$  and output dimension 5. Our results indicate that the mean length distortion is, as expected, approximately equal to  $\sqrt{n_L/n_0} \approx 0.08$ , and that it decays slightly with depth.

In Figure 4(b), we consider residual networks. Here, the overall network  $\mathcal{N}$  is defined in terms of residual modules  $\mathcal{N}_1, \mathcal{N}_2, \dots, \mathcal{N}_L$  and scales  $\eta_1, \eta_2, \dots, \eta_L$  according to:

$$\begin{aligned} \mathcal{N}(x) = & x + \eta_1 \mathcal{N}_1(x) + \eta_2 \mathcal{N}_2(x + \eta_1 \mathcal{N}_1(x)) + \dots \\ & + \eta_L \mathcal{N}_L(x + \eta_1 \mathcal{N}_1(x) + \eta_2 \mathcal{N}_2(x + \eta_1 \mathcal{N}_1(x)) + \dots). \end{aligned}$$

We set all residual modules  $\mathcal{N}_\ell$  to be two-layer, fully connected ReLU networks. In keeping with [14], we initialize all weights i.i.d. normal with variance  $2/\text{fan-in}$ , while setting  $\eta_1 = \eta_2 = \dots = \eta_L = 1/L$ . Our results suggest that mean length distortion again decays modestly overall, with a somewhat sharper decrease for small depths.

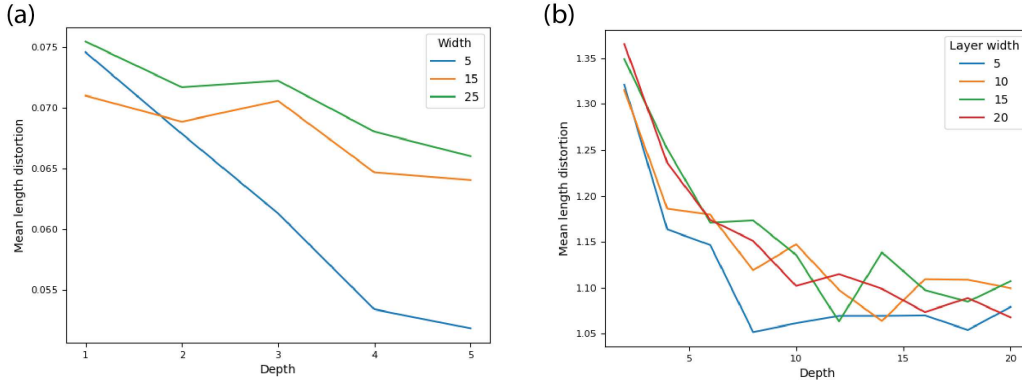


Figure 4: Mean length distortion as a function of depth, for networks with convolutional layers and skip connections, initialized using He normal initialization [17]. (a) shows results for networks having convolutional layers, where the input dimension  $n_0$  equals  $16 \times 16 \times 3 = 768$  and output dimension  $n_L$  equals 5. Here width corresponds to the number of  $3 \times 3$  kernels in each layer. As expected, the mean length distortion is  $\sqrt{n_L/n_0} \approx 0.08$  and decays modestly with depth. (b) shows results for networks with skip connections occurring between even layers (i.e. each residual block is a fully-connected ReLU network of depth 2), again taking a line segment of unit length as the input curve. Here, depth corresponds to the total number of layers in the network, so that a depth-20 network includes 10 residual blocks. Both plots (a) and (b) show the empirical mean over 100 initializations of the weights and biases.

## B Experimental Details

For all experiments, weights were initialized from i.i.d. normal distributions with variance  $2/\text{fan-in}$  and bias variance 0.1. We run several experiments that involve computing the length distortion of a given line segment in  $\mathbb{R}^{n_0}$ . We remark on how this was done, for which we rely on the notion of linear

regions and bent hyperplanes associated with ReLU networks, explored in [16, 26]. Specifically, ReLU networks partition input space into a collection of convex polytopes, which we call linear regions, as (generically) distinct linear functions are defined on each region corresponding to a different subset of the hidden neurons being activated (where a neuron is activated if its pre-activation is nonnegative). The boundaries of these linear regions are given by sets of points for which a particular neuron has preactivation equal to 0.

Say we are given a line segment  $\ell$  with endpoints  $\mathbf{p}_1, \mathbf{p}_2 \in \mathbb{R}^{n_0}$  parametrized by unit speed on  $[0, 1]$ , given by the set  $\{\mathbf{p}_1 + t(\mathbf{p}_2 - \mathbf{p}_1) : t \in [0, 1]\}$ , for which we aim to compute the length distortion. Let  $\mathbf{w}^\ell = \mathbf{p}_2 - \mathbf{p}_1$ .

We first approximate the intersections of the line segment with linear region boundaries using a binary search subroutine. Specifically, initialize the set  $\mathcal{S} = [0, 0.5, 1]$ , which will contain the parameter values for these approximations. For each iteration of the subroutine, we run the following procedure: for any three consecutive points  $t_1, t_2, t_3 \in \mathcal{S}$ , let  $\mathbf{x}_1 = \mathbf{p}_1 + t_1(\mathbf{p}_2 - \mathbf{p}_1)$ ,  $\mathbf{x}_2 = \mathbf{p}_1 + t_2(\mathbf{p}_2 - \mathbf{p}_1)$ ,  $\mathbf{x}_3 = \mathbf{p}_1 + t_3(\mathbf{p}_2 - \mathbf{p}_1)$ , and consider the values  $\frac{||\mathcal{N}(\mathbf{x}_2) - \mathcal{N}(\mathbf{x}_1)||}{||\mathbf{x}_2 - \mathbf{x}_1||}$  and  $\frac{||\mathcal{N}(\mathbf{x}_3) - \mathcal{N}(\mathbf{x}_2)||}{||\mathbf{x}_3 - \mathbf{x}_2||}$ . These values are equal if the three points are in the same linear region, in which case we eliminate  $t_1$  and  $t_3$  from  $\mathcal{S}$ . If not equal, then there exists a linear region boundary in the segment from  $\mathbf{x}_1$  to  $\mathbf{x}_3$ ; we add the points  $\frac{t_1+t_2}{2}$  and  $\frac{t_2+t_3}{2}$  to  $\mathcal{S}$ . We iterate this procedure (15 times in our implementation); for the final iteration, we ensure that only the last point associated with a given linear region is in  $\mathcal{S}$ .

Now take the set  $\mathcal{S} = [t_1, t_2, \dots, t_n]$  returned by the binary search procedure; we proceed to compute the parameter values denoting the exact intersections of the line segment with region boundaries, which we shall store in  $\mathcal{S}^*$ . For consecutive points  $t_i, t_{i+1} \in \mathcal{S}$ ,  $i = 1, \dots, n-1$ , determine the set of activated neurons for both points at  $\mathbf{x}_i = \mathbf{p}_1 + t_i(\mathbf{p}_2 - \mathbf{p}_1)$ ,  $\mathbf{x}_{i+1} = \mathbf{p}_1 + t_{i+1}(\mathbf{p}_2 - \mathbf{p}_1)$ , and find the neuron at which these sets differ; we solve a linear equation to determine the value  $t^* \in [0, 1]$  for which this neuron is 0, and replace  $t_i$  with the exact value  $t_i^*$  in  $\mathcal{S}^*$ . Finally, we append 0 and 1 to the respective ends of  $\mathcal{S}^* = [t_0^* = 0, t_1^*, \dots, t_n^* = 1]$ .

Computing the length distortion is reduced to summing the lengths of the output segments corresponding to consecutive pairs  $t_i, t_{i+1} \in \mathcal{S}^*$ . Namely, for each such pair, the network is given by a single linear function on the segment of inputs between  $\mathbf{x}_i = \mathbf{p}_1 + t_i(\mathbf{p}_2 - \mathbf{p}_1)$  and  $\mathbf{x}_{i+1} = \mathbf{p}_1 + t_{i+1}(\mathbf{p}_2 - \mathbf{p}_1)$ . We calculate the weight matrix  $\mathbf{W}_i$  of this linear function – the length of the corresponding output segment is the product of  $\mathbf{W}_i$  with the vector  $\mathbf{w}^\ell$ . The total length distortion is then given by the sum

$$\sum_{i=0}^{n-1} ||\mathbf{W}_i \mathbf{w}^\ell|| (t_{i+1}^* - t_i^*)$$

## C Proof of Theorem 5.1

Our first step in proving Theorem 5.1 is to obtain bounds on the moments of  $\text{len}(\mathcal{N}(M))$  in terms of the input-output Jacobian of  $\mathcal{N}$  at a single point, which we recall was defined in (6.4). To accomplish them, we recall the notation from §6. Namely, fix a smooth unit speed parameterization of  $M = \gamma([0, 1])$  with

$$\gamma : \mathbb{R} \rightarrow \mathbb{R}^{n_0}, \quad \gamma(t) = (\gamma_1(t), \dots, \gamma_{n_0}(t)).$$

The mapping

$$\Gamma := \mathcal{N} \circ \gamma, \quad \Gamma : \mathbb{R} \rightarrow \mathbb{R}^{n_L}$$

gives a parameterization of the curve  $\mathcal{N}(M)$ , and we have

$$\text{len}(\mathcal{N}(M)) = \int_0^1 ||\Gamma'(t)|| dt. \quad (\text{C.1})$$

Let us note an important but ultimately benign technicality. The Jacobian  $J_x$  of the map  $x \mapsto \mathcal{N}(x)$  is not defined at every  $x$  (namely those  $x$  where some neuron turns from on to off). Thus, *a priori*,  $\Gamma'(t)$  exists only at those  $t$  for which  $J_{\gamma(t)}$  exists. However, the formula (C.1) is still valid. Indeed, for any setting of weights and biases of  $\mathcal{N}$  the map  $\Gamma$  is Lipschitz. Thus, by Rademacher's theorem,  $\Gamma'(t)$  exists for almost every  $t \in [0, 1]$  and the length of the curve is given by integrating the norm of this almost-everywhere defined derivative. The following simple Lemma is a generalization of Lemma 6.1 and allows us to bound all moments of  $\text{len}(\mathcal{N}(M))$  in terms of the moments of the norm of the Jacobian vector product  $||J_x u||$ .

**Lemma C.1.** *For any integer  $m \geq 1$ , we have*

$$\mathbb{E} [\text{len}(\mathcal{N}(M))^m] \leq \int_0^1 \mathbb{E} [||J_{\gamma(t)} \gamma'(t)||^m] dt. \quad (\text{C.2})$$

*Proof.* Taking powers in (C.1) and using Tonelli's theorem to interchange the expectation and integrals, we obtain

$$\mathbb{E} [\text{len}(\mathcal{N}(M))^m] = \int_0^1 \cdots \int_0^1 \mathbb{E} \left[ \prod_{j=1}^m ||\Gamma'(t_j)|| \right] dt_1 \cdots dt_m. \quad (\text{C.3})$$

To proceed, let us recall a special case of the power mean inequality which says that for any  $a_i \geq 0$  we have

$$\prod_{i=1}^m a_i \leq \frac{1}{m} \sum_{i=1}^m a_i^m$$

Applying this to the integrand in (C.3), we conclude

$$\mathbb{E} [\text{len}(\mathcal{N}(M))^m] \leq \int_0^1 \mathbb{E} [||\Gamma'(t)||^m] dt. \quad (\text{C.4})$$

Next, fix  $t \in [0, 1]$ . For any neuron  $z$  in  $\mathcal{N}$ , denote by  $x \mapsto z(x)$  its pre-activation. Note that  $\mathbb{P}(z(\gamma(t)) = 0) = 0$  since our bias variance  $C_b$  is set to some fixed positive constant and hence the bias of each neuron has a continuous density. Therefore, with probability 1 over the randomness in the weights and biases of  $\mathcal{N}$ , there exists a neighborhood  $\mathcal{U}$  of  $\gamma(t)$  on which  $z(x)$  has constant sign for all  $x \in \mathcal{U}$ . The Jacobian  $J_{\gamma(t)}$  is therefore well-defined and the chain rule yields

$$\Gamma'(t) = J_{\gamma(t)} \gamma'(t). \quad (\text{C.5})$$

Substituting this into (C.4) completes the proof.  $\square$

Having obtained the expression (C.2), we have reduced bounding the moments of  $\text{len}(\mathcal{N}(M))$  to bounding the moments of the random variable  $||J_x u||$  for a fixed  $x \in \mathbb{R}^{n_0}$  and unit vector  $u \in \mathbb{R}^{n_0}$ . Prior work (e.g. Thm. 1 in [12]) shows that these moments satisfy

$$\mathbb{E} [||J_x u||^{2m}] = \left( \frac{n_L}{n_0} \right)^m \exp \left( 5 \binom{m}{2} \sum_{\ell=1}^{L-1} n_\ell^{-1} + O \left( \sum_{\ell=1}^{L-1} n_\ell^{-2} \right) \right),$$

provided

$$\binom{m}{2} < \min_{\ell=1, \dots, L-1} n_\ell. \quad (\text{C.6})$$

Substituting these moment estimates in (C.2) completes the derivation of (5.3). However, the results in [12] are subtle because they apply to any distribution of weights and biases. They also give the slightly sub-optimal restriction (C.6) that  $m^2$  must be smaller than a constant times the minimum of the  $n_\ell$ 's. In the special case where the distribution of weights and biases is Gaussian, we can do slightly better by computing the moments of  $||J_x u||$  more directly (note that in the statement of Theorem 5.1, we required only that  $m$  is smaller than a constant times the minimum of the  $n_\ell$ 's). We will need this anyway to derive the slightly more refined estimates in (5.1) and (5.2). Let us therefore check that, in distribution,  $||J_x u||$  is equal to a product of independent random variables:

**Proposition C.2.** *For any  $x \in \mathbb{R}^{n_0}$  and any unit vector  $u \in \mathbb{R}^{n_0}$ , the random variable  $||J_x u||^2$  is equal in distribution to a product of independent scaled chi-squared random variables*

$$||J_x u||^2 \stackrel{d}{=} \frac{n_L}{n_0} \left( \prod_{\ell=1}^{L-1} \frac{2}{n_\ell} \chi_{\mathbf{n}_\ell}^2 \right) \cdot \frac{1}{n_L} \chi_{n_L}^2,$$

where the number of degrees of freedom in the  $\ell$ th term of the product (for  $\ell = 1, \dots, L-1$ ) is given by an independent binomial

$$\mathbf{n}_\ell \stackrel{d}{=} \text{Bin}(n_\ell, 1/2)$$

with  $n_\ell$  trials and success probability  $1/2$ . The number of degrees of freedom in the final term is deterministic and given by  $n_L$ .

*Proof.* Consider a ReLU network  $\mathcal{N}$  with input dimension 1, output dimension  $n_L$ , and hidden layer widths  $n_1, \dots, n_{L-1}$ . Suppose the weight matrices  $W^{(\ell)}$  and bias vectors  $b^{(\ell)}$  are independent with i.i.d. Gaussian components:

$$W_{ij}^{(\ell)} \sim G(0, 2/n_{\ell-1}), \quad b_j^{(\ell)} \sim G(0, C_b),$$

where  $C_b > 0$  is any fixed constant. For a fixed network input  $x \in \mathbb{R}^{n_0}$ , with probability 1, the input-output Jacobian  $J_x$  is well-defined. Moreover, it can be written as

$$J_x = W^{(L)} D^{(L-1)} W^{(L-1)} \dots D^{(1)} W^{(1)},$$

where  $W^{(\ell)}$  is the matrix of weights from layer  $\ell - 1$  to layer  $\ell$  and  $D^{(\ell)}$  is an  $n_\ell \times n_\ell$  diagonal matrix:

$$D^{(\ell)} = \text{Diag} \left( \mathbf{1}_{\{z_i^{(\ell)} \geq 0\}}, i = 1, \dots, n_\ell \right)$$

whose diagonal entries are 0 or 1 depending on whether the pre-activation  $z_i^{(\ell)}$  of neuron  $i$  in layer  $\ell$  is positive at our fixed input  $x$ . Next, according to Proposition 2 in [12], the marginal distribution of each  $D^{(\ell)}$  is that its diagonal entries are independent Bernoulli(1/2) random variables. Moreover, we have the following equality in distribution:

$$J_x \stackrel{d}{=} \eta W^{(L)} \mathcal{D}^{(L-1)} W^{(L-1)} \dots \mathcal{D}^{(1)} W^{(1)}$$

where  $\mathcal{D}^{(\ell)}$  are independent of each other (and of  $W^{(\ell)}$ ) resampled from the marginal distribution of  $D^{(\ell)}$  and  $\eta$  is an independent diagonal matrix with independent diagonal entries that are  $\pm 1$  with equal probability. In particular, for a fixed unit vector  $u \in \mathbb{R}^{n_0}$ , we have

$$\|J_x u\| \stackrel{d}{=} \|W^{(L)} \mathcal{D}^{(L-1)} W^{(L-1)} \dots \mathcal{D}^{(1)} W^{(1)} u\|.$$

We may rewrite this as

$$\|W^{(L)} \mathcal{D}^{(L-1)} W^{(L-1)} \dots \mathcal{D}^{(2)} W^{(2)} u^{(1)}\| \|\mathcal{D}^{(1)} W^{(1)} u\|, \quad (\text{C.7})$$

for  $u^{(1)} := \frac{\mathcal{D}^{(1)} W^{(1)} u}{\|\mathcal{D}^{(1)} W^{(1)} u\|}$ , where we interpret the expression (C.7) as equal to 0 if  $\mathcal{D}^{(1)}$  is the zero matrix. To complete the proof, we need the following standard observation.

**Lemma C.3.** *Suppose  $W$  is an  $n \times n'$  matrix with i.i.d. Gaussian entries and  $u$  is a random unit vector in  $\mathbb{R}^{n'}$  that is independent of  $W$  but otherwise has any distribution. Then  $Wu$  is independent of  $u$  and is equal in distribution to  $Wv$  where  $v$  is any fixed unit vector in  $\mathbb{R}^{n'}$ .*

*Proof.* For any fixed orthogonal matrix  $\mathcal{O}$ , it is a standard fact that  $W\mathcal{O}$  is equal in distribution to  $W$ . Thus, for any measurable sets  $A \subseteq \mathbb{R}^n$  and  $B \subseteq \mathbb{R}^{n'}$ , since  $u, W$  are independent we have

$$\mathbb{P}(Wu \in A, u \in B) = \int_{S^{n'-1}} \mathbb{P}(Wu_0 \in A) d\mathbb{P}_u(u_0),$$

where  $\mathbb{P}_u(u_0)$  is the distribution of  $u$ . Fix any  $u_0 \in S^{n'-1}$  and let  $\mathcal{O} = \mathcal{O}(u_0)$  be any orthogonal matrix so that  $u_0 = \mathcal{O}e_1$  with  $e_1 = (1, 0, \dots, 0)$  is the first standard unit vector. Then, since  $W\mathcal{O}$  is equal in distribution to  $W$ , we have

$$\mathbb{P}(Wu_0 \in A) = \mathbb{P}(We_1 \in A),$$

which is independent of  $u_0$ . We therefore find

$$\mathbb{P}(Wu \in A, u \in B) = \mathbb{P}(We_1 \in A) \mathbb{P}(u \in B),$$

as desired.  $\square$

We are now in a position to complete the proof of Proposition C.2. Combining Lemma C.3 with (C.7), we find that, in distribution  $\|J_x u\|$  equals

$$\|W^{(L)} \mathcal{D}^{(L-1)} W^{(L-1)} \dots \mathcal{D}^{(2)} W^{(2)} u\| \|\mathcal{D}^{(1)} W^{(1)} u\|. \quad (\text{C.8})$$

Note that these two terms are independent. Repeating this argument, we obtain that  $\|J_x u\|$  is equal in distribution to the product

$$\|W^{(L)} u\| \|\mathcal{D}^{(L-1)} W^{(L-1)} u\| \dots \|\mathcal{D}^{(1)} W^{(1)} u\|. \quad (\text{C.9})$$

The terms in this product are independent. To complete the proof note that for  $\ell = 1, \dots, L-1$  the number of non-zero entries in  $\mathcal{D}^{(\ell)}$  is a binomial random variable  $\mathbf{n}_\ell$  with  $n_\ell$  trials, each with probability of success  $1/2$  and that

$$\|\mathcal{D}^{(\ell)} W^{(\ell)} u\|^2$$

is precisely  $2/n_\ell$  times a sum of  $\mathbf{n}_\ell$  squares of independent standard Gaussians. Thus, for  $\ell = 1, \dots, L-1$ ,

$$\|\mathcal{D}^{(\ell)} W^{(\ell)} u\|^2 \stackrel{d}{=} \frac{2}{n_{\ell-1}} \chi_{\mathbf{n}_\ell}^2. \quad (\text{C.10})$$

Similarly

$$\|W^{(L)} u\|^2 \stackrel{d}{=} \frac{2}{n_{L-1}} \chi_{n_L}^2. \quad (\text{C.11})$$

Substituting (C.10) and (C.11) into (C.9) completes the proof.  $\square$

Evaluating the moments of  $\text{len}(\mathcal{N}(M))$  is now a matter of computing the moments of some scaled chi-squared random variables with a random number of degrees of freedom. For instance, recalling that

$$\mathbb{E}[\chi_k^2] = k$$

and applying Lemma C.2 as well as the tower property of the expectation we find

$$\begin{aligned} \mathbb{E}[\|J_x u\|^2] &= \frac{n_L}{n_0} \left( \prod_{\ell=1}^{L-1} \mathbb{E} \left[ \frac{2}{n_\ell} \chi_{\mathbf{n}_\ell}^2 \right] \right) \mathbb{E}[n_L^{-1} \chi_{n_L}^2] \\ &= \frac{n_L}{n_0} \prod_{\ell=1}^{L-1} \left\{ \frac{2}{n_\ell} \mathbb{E}[\mathbb{E}[\chi_{\mathbf{n}_\ell}^2 \mid \mathbf{n}_\ell]] \right\} \\ &= \frac{n_L}{n_0} \prod_{\ell=1}^{L-1} \left\{ \frac{2}{n_\ell} \mathbb{E}[\mathbf{n}_\ell] \right\} \\ &= \frac{n_L}{n_0}. \end{aligned}$$

Substituting this into (C.2) with  $m = 2$  yields

$$\text{Var}[\text{len}(\mathcal{N}(M))] \leq \frac{n_L}{n_0},$$

yielding (5.2). To prove (5.1), we need to estimate  $\mathbb{E}[\text{len} \mathcal{N}(M)]$ . By taking expectations in (C.1) and using (C.5), we find

$$\mathbb{E}[\text{len} \mathcal{N}(M)] = \int_0^1 \mathbb{E}[\|J_{\gamma(t)} \gamma'(t)\|] dt = \mathbb{E}[\|J_x u\|],$$

where we've used that by Lemma C.2, the distribution of  $\|J_x u\|$  is the same for every  $x \in \mathbb{R}^{n_0}$  and every unit vector  $u$ . Moreover, again using Lemma C.2, we see that  $\mathbb{E}[\|J_x u\|]$  equals

$$\left( \frac{n_0}{n_L} \right)^{1/2} \prod_{\ell=1}^{L-1} \mathbb{E} \left[ \left( \frac{2}{n_\ell} \chi_{\mathbf{n}_\ell}^2 \right)^{1/2} \right] \mathbb{E} \left[ \left( \frac{1}{n_L} \chi_{n_L}^2 \right)^{1/2} \right].$$

To simplify this, a direct computation shows that

$$\mathbb{E} \left[ \left( k^{-1} \chi_k^2 \right)^{1/2} \right] = \frac{\Gamma \left( \frac{k+1}{2} \right)}{\Gamma \left( \frac{k}{2} \right) \left( \frac{k}{2} \right)^{1/2}},$$

where  $\Gamma(\cdot)$  is the Gamma function. Next, for any positive random variable  $X$  with  $\mathbb{E}[X] = 1$  we have

$$\begin{aligned}\mathbb{E}[X^{1/2}] &= \mathbb{E}[(1 + (X - 1))^{1/2}] \\ &= 1 - \frac{1}{8}\text{Var}[X] + O(|X - 1|^3).\end{aligned}$$

Recall that

$$\text{Var}[\chi_k^2] = 2k.$$

Using this and the law of total variance yields

$$\begin{aligned}\mathbb{E}\left[\left(\frac{2}{n_\ell}\chi_{\mathbf{n}_\ell}^2\right)^{1/2}\right] &= 1 - \frac{1}{2n_\ell^2}\text{Var}[\chi_{\mathbf{n}_\ell}^2] + O(n_\ell^{-2}) \\ &= 1 - \frac{1}{2n_\ell^2}\left[\mathbb{E}[\text{Var}[\chi_{\mathbf{n}_\ell}^2 \mid \mathbf{n}_\ell]]\right. \\ &\quad \left.+ \text{Var}[\mathbb{E}[\chi_{\mathbf{n}_\ell}^2 \mid \mathbf{n}_\ell]]\right] + O(n_\ell^{-2}) \\ &= 1 - \frac{1}{2n_\ell^2}[\mathbb{E}[2\mathbf{n}_\ell] + \text{Var}[\mathbf{n}_\ell]] + O(n_\ell^{-2}) \\ &= 1 - \frac{1}{2n_\ell^2}\left[n_\ell + \frac{n_\ell}{4}\right] + O(n_\ell^{-2}) \\ &= 1 - \frac{5}{8n_\ell} + O(n_\ell^{-2}).\end{aligned}$$

Thus, we find that  $\mathbb{E}[\text{len}(\mathcal{N}(M))]$  equals

$$\left(\frac{n_L}{n_0}\right)^{1/2} \frac{\Gamma\left(\frac{n_L+1}{2}\right)}{\Gamma\left(\frac{n_L}{2}\right)\left(\frac{n_L}{2}\right)^{1/2}} \times \exp\left[-\frac{5}{8}\sum_{\ell=1}^{L-1}n_\ell^{-1} + O\left(\sum_{\ell=1}^{L-1}n_\ell^{-2}\right)\right],$$

as claimed. Finally, let us check the higher moment estimates (5.3). By Lemma C.2 we have the following estimate

$$\begin{aligned}\|J_x u\|^2 &= \frac{n_L}{n_0} \left(\prod_{\ell=1}^{L-1} \frac{2}{n_\ell}\chi_{\mathbf{n}_\ell}^2\right) n_L^{-1}\chi_{n_L}^2 \\ &\leq \frac{n_L}{n_0} \exp\left[\sum_{\ell=1}^{L-1}\left\{\frac{2}{n_\ell}\chi_{\mathbf{n}_\ell}^2 - 1\right\}\right] n_L^{-1}\chi_{n_L}^2,\end{aligned}$$

where we used that  $x = x - 1 + 1 \leq e^{x-1}$ . For any  $m$ , we have

$$\begin{aligned}\mathbb{E}\left[\left(\frac{1}{n_L}\chi_{n_L}^2\right)^m\right] &= \left(1 + \frac{2}{n_L}\right) \cdots \left(1 + \frac{2m-2}{n_L}\right) \\ &\leq \exp\left[\sum_{j=0}^{m-1} \frac{2j}{n_L}\right] \leq \exp\left[\frac{m^2}{n_L}\right].\end{aligned}$$

Therefore,  $\|J_x u\|^{2m}$  is bounded above by

$$\left(\frac{n_L}{n_0}\right)^m \exp\left[m\sum_{\ell=1}^{L-1}\left\{\frac{2}{n_\ell}\chi_{\mathbf{n}_\ell}^2 - 1\right\} + \frac{m^2}{n_L}\right]. \quad (\text{C.12})$$

Finally, for any fixed positive integer  $n$  we may write

$$\frac{2}{n}\chi_{\mathbf{n}}^2 - 1 = \frac{2}{n}\sum_{k=1}^n\left\{\xi_k Z_k^2 - \frac{1}{2}\right\},$$



where  $\xi_k$  are independent Bernoulli(1/2) random variables and  $Z_k$  are independent standard Gaussians. A direct computation shows that the centered random variables  $\xi_k Z_k^2 - \frac{1}{2}$  are sub-exponential  $\text{SE}(\nu^2, \alpha)$  with some universal parameters  $\nu^2, \alpha$ . Thus, by the stability of sub-exponential random variables under summation we have

$$\frac{2}{n_\ell} \chi_{\mathbf{n}_\ell}^2 - 1 \in \text{SE} \left( \frac{4\nu^2}{n_\ell}, \frac{2\alpha}{n_\ell} \right).$$

Again using this property we conclude

$$\sum_{\ell=1}^{L-1} \frac{2}{n_\ell} \chi_{\mathbf{n}_\ell}^2 - 1 \in \text{SE} \left( 4\nu^2 \sum_{\ell=1}^{L-1} \frac{1}{n_\ell}, \frac{2\alpha}{n_*} \right),$$

where

$$n_* = \min \{n_1, \dots, n_{L-1}\}.$$

Therefore, by (C.12), we find that

$$\mathbb{E} \left[ \|J_x u\|^{2m} \right] \leq \left( \frac{n_L}{n_0} \right)^m \mathbb{E} [e^{mY}] \exp \left[ \frac{m^2}{n_L} \right],$$

where  $Y \in \text{SE} \left( 4\nu^2 \sum_{\ell=1}^{L-1} \frac{1}{n_\ell}, \frac{2\alpha}{n_*} \right)$ . By definition of sub-exponential random variables, we have

$$\mathbb{E} [e^{mY}] \leq \exp \left[ 4m^2 \nu^2 \sum_{\ell=1}^{L-1} \frac{1}{n_\ell} \right],$$

provided  $m < n_*/2\alpha$  for some universal constant  $c > 0$ . This completes the derivation of (5.3) and completes the proof of Theorem 5.1.  $\square$

## D Proof of Theorem 5.2

The proof of Theorem 5.2 follows closely the argument for Theorem 5.1. For a given input  $x \in \mathbb{R}^{n_0}$ , we will continue to write

$$J_{\mathcal{N}}(x) := (\partial_{x_j} \mathcal{N}_i(x))_{\substack{1 \leq j \leq n_0 \\ 1 \leq i \leq n_L}}$$

for the input-output Jacobian of the network function, which exists for Lebesgue almost every  $x \in \mathbb{R}^{n_0}$ . We will write  $\Pi_{T_x M} : \mathbb{R}^{n_0} \rightarrow T_x M$  for the orthogonal projection onto this tangent space. We have

$$\text{vol}_d(\mathcal{N}(M)) = \int_M (\det (\Pi_{T_x M} J_{\mathcal{N}}(x)^T J_{\mathcal{N}}(x) \Pi_{T_x M}))^{1/2} \text{vol}_d(dx),$$

where  $\text{vol}_d$  is the  $d$ -dimensional Hausdorff measure. Indeed, the integrand measures, at each  $x \in M$ , the volume of the image of an infinitesimal cube on  $M$  of volume  $\text{vol}_d(dx)$  centered at  $x$  under the map  $x \mapsto \mathcal{N}(x)$ . Arguing precisely as in the proof of Lemma C.1, we find that for any integer  $m \geq 1$

$$\mathbb{E} [\text{vol}_d(\mathcal{N}(M))^m] \leq \int_M \mathbb{E} [\det (\Pi_{T_x M} J_{\mathcal{N}}(x)^T J_{\mathcal{N}}(x) \Pi_{T_x M})]^{m/2} \text{vol}_d(dx) \quad (\text{D.1})$$

Fix  $x \in M$  and denote by

$$e_1(x), \dots, e_d(x)$$

an orthonormal basis of the tangent space of  $M$ . Then, by the Gram identity, we may write

$$\det (\Pi_{T_x M} J_{\mathcal{N}}(x)^T J_{\mathcal{N}}(x) \Pi_{T_x M}) = \|J_{\mathcal{N}}(x) e_1(x) \wedge \dots \wedge J_{\mathcal{N}}(x) e_d(x)\|^2, \quad (\text{D.2})$$

where we recall that the wedge product is the anti-symmetrization of the tensor product. Just as in the proof of Theorem 5.1, the key observation is that for Gaussian weights we have that  $\|J_{\mathcal{N}}(x) e_1(x) \wedge \dots \wedge J_{\mathcal{N}}(x) e_d(x)\|$  is a product of i.i.d. random variables. The formal statement is the following:

**Proposition D.1.** For any  $x \in \mathbb{R}^{n_0}$  and collection of orthonormal unit vectors  $u_1, \dots, u_d \in \mathbb{R}^{n_0}$ , the random variable

$$\|J_x u_1 \wedge \dots \wedge J_x u_d\|^2$$

is equal in distribution to a product of independent scaled chi-squared random variables:

$$\left(\frac{n_L}{n_0}\right)^d \left( \prod_{\ell=1}^{L-1} \prod_{j=1}^d \frac{2}{n_\ell} \chi_{\mathbf{n}_\ell - j + 1}^2 \right) \cdot \prod_{j=1}^d \frac{1}{n_L} \chi_{\mathbf{n}_L - j + 1}^2,$$

where all the terms in the product are independent and for  $\ell = 1, \dots, L-1$  we've written  $\mathbf{n}_\ell$  for independent Binomial random variables:

$$\mathbf{n}_\ell \stackrel{d}{=} \text{Bin}(n_\ell, 1/2)$$

with  $n_\ell$  trials and success probability  $1/2$ .

*Proof.* The proof is identical to that of Lemma C.2. The only difference is that we must invoke the following amplification of Lemma C.3:

**Lemma D.2.** Let  $u_1, \dots, u_k \in \mathbb{R}^{n'}$  be a collection of orthonormal vectors (i.e. a collection) with any distribution. Let  $W$  be an independent  $n \times n'$  matrix with i.i.d. Gaussian entries. Then

$$W(u_1 \wedge \dots \wedge u_k) = W u_1 \wedge \dots \wedge W u_k$$

is independent of  $u_1 \wedge \dots \wedge u_k$  and is equal in distribution to  $W v_1 \wedge \dots \wedge W v_k$  where  $v_1, \dots, v_k$  is any fixed collection of orthonormal vectors.

*Proof.* The proof is identical to that of Lemma C.3, except that we note that for that, given  $u_1, \dots, u_k$ , there exists an orthogonal matrix  $\mathcal{O} = \mathcal{O}(u_1, \dots, u_k)$  so that

$$u_j = \mathcal{O} e_j, \quad j = 1, \dots, k$$

where  $e_j$  are the standard basis vectors. □

□

With Lemma D.1 in hand, we complete the proof of Theorem 5.2 as follows. First, note that for any random variable  $X$  we have

$$\mathbb{E}[X] \leq \mathbb{E}[X^2]^{1/2}, \quad \text{Var}[X] \leq \mathbb{E}[X^2].$$

Hence, Theorem 5.2 will follow once we show that

$$\mathbb{E}[\text{vol}_d(\mathcal{N}(M))^2] \leq \left(\frac{n_L}{n_0}\right)^d \exp \left[ -\binom{d}{2} \sum_{\ell=1}^L n_\ell^{-1} \right].$$

Combining Proposition D.1 with (D.1) and (D.2), this estimate follows by showing that

$$\mathbb{E}[\|J_x u_1 \wedge \dots \wedge J_x u_d\|^2] \leq \left(\frac{n_L}{n_0}\right)^d \exp \left[ -\binom{d}{2} \sum_{\ell=1}^L n_\ell^{-1} \right]. \quad (\text{D.3})$$

To check this, recall that

$$\mathbb{E}[\chi_k^2] = k.$$

Hence,

$$\begin{aligned} \mathbb{E}[\|J_{\mathcal{N}}(x) e_1(x) \wedge \dots \wedge J_{\mathcal{N}}(x) e_d(x)\|^2] &= \left(\frac{n_L}{n_0}\right)^d \left( \prod_{\ell=1}^{L-1} \prod_{j=1}^d \frac{2}{n_\ell} \mathbb{E}[\chi_{\mathbf{n}_\ell - j + 1}^2] \right) \cdot \prod_{j=1}^d \frac{1}{n_L} \mathbb{E}[\chi_{\mathbf{n}_L - j + 1}^2] \\ &= \left(\frac{n_L}{n_0}\right)^d \prod_{\ell=1}^L \prod_{j=1}^d \left(1 - \frac{j-1}{n_\ell}\right) \\ &\leq \left(\frac{n_L}{n_0}\right)^d \exp \left[ -\binom{d}{2} \sum_{\ell=1}^L n_\ell^{-1} \right], \end{aligned}$$

where in the last line we used that  $1 + x \leq e^x$ . This completes the proof. □