Boris Hanin^{*1} David Rolnick^{*2}

Abstract

It is well-known that the expressivity of a neural network depends on its architecture, with deeper networks expressing more complex functions. In the case of networks that compute piecewise linear functions, such as those with ReLU activation, the number of distinct linear regions is a natural measure of expressivity. It is possible to construct networks with merely a single region, or for which the number of linear regions grows exponentially with depth; it is not clear where within this range most networks fall in practice, either before or after training. In this paper, we provide a mathematical framework to count the number of linear regions of a piecewise linear network and measure the volume of the boundaries between these regions. In particular, we prove that for networks at initialization, the average number of regions along any one-dimensional subspace grows linearly in the total number of neurons, far below the exponential upper bound. We also find that the average distance to the nearest region boundary at initialization scales like the inverse of the number of neurons. Our theory suggests that, even after training, the number of linear regions is far below exponential, an intuition that matches our empirical observations. We conclude that the practical expressivity of neural networks is likely far below that of the theoretical maximum, and that this gap can be quantified.

1. Introduction

A growing field of theory has sought to explain the broad success of deep neural networks via a mathematical characterization of the ability of these networks to approximate dif-



Figure 1. How many linear regions? This figure shows a twodimensional slice through the 784-dimensional input space of vectorized MNIST, as represented by a fully-connected ReLU network with three hidden layers of width 64 each. Colors denote different linear regions of the piecewise linear network.

ferent functions of input data. Many such works consider the expressivity of neural networks, showing that certain functions are more efficiently expressible by deep architectures than by shallow ones (e.g. Bianchini & Scarselli (2014); Montufar et al. (2014); Telgarsky (2015); Lin et al. (2017); Rolnick & Tegmark (2018)). It has, however, also been noted that many expressible functions are not efficiently learnable, at least by gradient descent (Shalev-Shwartz et al., 2018). More generally, the typical behavior of a network used in practice, the practical expressivity, may be very different from what is theoretically attainable. To adequately explain the power of deep learning, it is necessary to consider networks with parameters as they will naturally occur before, during, and after training.

Networks with a piecewise linear activation (e.g. ReLU, hard tanh) compute piecewise linear functions for which input space is divided into pieces, with the network computing a single linear function on each piece (see Figures 1-4). Fig-ure 2 shows how the complexity of these pieces, which we refer to as linear regions, changes in a deep ReLU net with two-dimensional inputs. Each neuron in the first layer splits the input space into two pieces along a hyperplane, fitting a different linear function to each of the pieces. Subsequent layers split the regions of the preceding layers. The local density of linear regions serves as a convenient proxy for the local complexity or smoothness of the network, with the ability to interpolate a complex data distribution seeming to require fitting many relatively small regions. The topic of

^{*}Equal contribution ¹Department of Mathematics, Texas A&M University and Facebook AI Research, New York ²University of Pennsylvania. Correspondence to: Boris Hanin <bhanin@tamu.edu>, David Rolnick <drolnick@seas.upenn.edu>.

Proceedings of the 36th International Conference on Machine Learning, Long Beach, California, PMLR 97, 2019. Copyright 2019 by the author(s).

counting linear regions is taken up by a number of authors (Telgarsky, 2015; Montufar et al., 2014; Serra et al., 2018; Raghu et al., 2017).

A worst case estimate is that every neuron in each new layer splits each of the regions present at the previous layer, giving a number of regions exponential in the depth. Indeed this is possible, as examined extensively e.g. in Montufar et al. (2014). An example of Telgarsky (2015) shows that a sawtooth function with 2^n teeth can be expressed exactly using only 3n + 4 neurons, as shown in Figure 3. However, even slightly perturbing this network (by adding noise to the weights and biases) ruins this beautiful structure and severely reduces the number of linear pieces, raising the question of whether typical neural networks actually achieve the theoretical bounds for numbers of linear regions.



Figure 2. Evolution of linear regions within a ReLU network for 2-dimensional input. Each neuron in the first layer defines a linear boundary that partitions the input space into two regions. Neurons in the second layer combine and split these linear boundaries into higher level patterns of regions, and so on. Ultimately, the input space is partitioned into a number of regions, on each of which the neural network is given by a (different) linear function. During training, both the partition into regions and the linear functions on them are learned.

Figure 1 also invites measures of complexity for piecewise linear networks beyond region counting. The boundary between two linear regions can be straight or can be bent in complex ways, for example, suggesting the volume of the boundary between linear regions as complexity measure for the resulting partition of input space. In the 2D example of Figure 1, this corresponds to computing perimeters of the linear pieces. As we detail below, this measure has another natural advantage: the volume of the boundary controls the typical distance from a random input to the boundary of its linear region (see §2.2). This measures the stability of the function computed by the network, and it is intuitively related to robustness under adversarial perturbation.

Our Contributions. In this paper, we provide mathematical tools for analyzing the complexity of linear regions of a



Figure 3. The sawtooth function on the left with 2^n teeth can be expressed succinctly by a ReLU network with only 3n + 4 neurons (construction from Telgarsky (2015)). However, slight perturbation of the weights and biases of the network (by Gaussian noise with standard deviation 0:1) greatly simplifies the linear regions captured by the network.

network with piecewise linear activations (such as ReLU) before, during, and after training. Our main contributions are as follows:

- For networks at initialization, the total surface area of the boundary between linear regions scales as the number of neurons times the number of breakpoints of the activation function. This is our main result, from which several corollaries follow (see Theorem 3, Corollary 4, and the discussion in §2).
- In particular, for any line segment through input space, the average number of regions intersecting it is linear in the number of neurons, far below the exponential number of regions that is theoretically attainable.
- Theorem 3 also allows us to conclude that, at initialization, the average distance from a sample point to the nearest region boundary is bounded below by a constant times the reciprocal of the number of neurons (see Corollary 5).
- We find empirically that both the number of regions and the distance to the nearest region boundary stay roughly constant during training and in particular are far from their theoretical maxima. That this should be the case is strongly suggested by Theorem 3, though not a direct consequence of it.

Overall, our results stress that practical expressivity lags significantly behind theoretical expressivity. Moreover, both our theoretical and empirical findings suggest that for certain measures of complexity, trained deep networks are remarkably similar to the same networks at initialization. In the next section, we informally state our theoretical and empirical results and explore the underlying intuitions. Detailed descriptions of our experiments are provided in §3. The precise theorem statements for ReLU networks can be found in §5. The exact formulations for general piecewise linear networks are in Appendix A, with proofs in the rest of the Supplementary Material. In particular, Appendix B contains intuition for how our proofs are shaped, while details are completed in §C-D.



Figure 4. Graph of function computed by a ReLU net with input and output dimension 1 at initialization. The weights of the network are He normal (i.i.d. normal with variance = 2=fan-in) and the biases are i.i.d. normal with variance 10⁻⁶.

2. Informal Overview of Results

This section gives an informal introduction to our results. We begin in §2.1 by describing the case of networks with input dimension 1: In §2.2, we consider networks with higher input dimension. For simplicity, we focus throughout this section on fully connected ReLU networks. We emphasize, however, that our results apply to any piecewise linear activation. Moreover, the upper bounds we present in Theorems 1, 2, and 3 (and hence in Corollaries 4 and 5) can also be generalized to hold for feed-forward networks with arbitrary connectivity, though we do not go into details in this work, for the sake of clarity of exposition.

2.1. Number of Regions in 1D

Consider the simple case of a ReLU net N with input and output dimensions equal to 1: Such a network computes a piecewise linear function (see Figure 4), and we are interested in understanding both at initialization and during training the number of distinct linear regions. There is a simple universal upper bound:

where the maximum is over all settings of weight and biases. This bound depends on the architecture of N only via the number of neurons. For more refined upper bounds which take into account the widths of the layers, see Theorem 1 in Raghu et al. (2017) and Theorem 1 in Serra et al. (2018). The constructions in Montufar et al. (2014); Telgarsky (2015); Raghu et al. (2017); Serra et al. (2018) indicate that the bound in (1) is very far from sharp for shallow and wide networks but that exponential growth in the number of regions can be achieved in deep, skinny networks for very special choices of weights and biases. This is a manifestation of the expressive power of depth, the idea that repeated compositions allow deep networks to capture complex hierarchical relations more efficiently per parameter than their shallow cousins. However, there is no non-trivial lower bound for the number of linear regions:

min #fregionsg = 1; 8N:

The minimum is attained by setting all weights and biases to 0: This raises the question of the behavior for the average number of regions when the weights and biases are chosen at random (e.g. at initialization). Intuitively, configurations of weights and biases that come close to saturating the exponential upper bound (1) are numerically unstable in the sense that a small random perturbation of the weights and biases drastically reduces the number of linear regions (see Figure 3 for an illustration). In this direction, we prove a somewhat surprising answer to the question of how many regions N has at initialization. We state the result for ReLU but note that it holds for any piecewise linear, continuous activation function (see Theorems 3 and 6).

Theorem 1 (informal). Let N be a network with piecewise linear activation with input and output dimensions of N both equal 1. Suppose the weights and biases are randomly initialized so that for each neuron z, its pre-activation z(x) has bounded mean gradient

$$E[krz(x)k] C;$$
 some $C > 0$: (2)

This holds, for example, for ReLU networks initialized with independent, zero-centered weights with variance 2=fan-in: Then, for each subset I R of inputs, the average number of linear regions inside I is proportional to the number of neurons times the length of I

```
E [#fregions in Ig] jIj T #fneuronsg;
```

where T is the number of breakpoints in the non-linearity of N (for ReLU nets, T = 1). The same result holds when computing the number of linear regions along any fixed 1-dimensional curve in a high-dimensional input space.

This theorem implies that the average number of regions along a one-dimensional curve in input space is proportional to the number of neurons, but independent of the arrangement of those neurons. In particular, a shallow network and a deep network will have the same complexity, by this measure, as long as they have the same total number of neurons. Of course, as j1j grows, the bounds in Theorem 1 become less sharp. We plan to extend our results to obtain bounds on the total number of regions on all of R in the future. In particular, we believe that at initialization the mean total number of linear regions N is proportional to the number of neurons (this is borne out in Figure 5, which computes the total number of regions on an infinite line).

Theorem 1 defies the common intuition that, on average, each layer in N multiplies the number of regions formed up to the previous layer by a constant larger than one. This would imply that the average number of regions is exponential in the depth. To provide intuition for why this is not true for random weights and biases, consider the effect of each neuron separately. Suppose the pre-activation z(x)of a neuron z satisfies $jz^{0}(x)j = (1)$, a hallmark of any reasonable initialization. Then, over a compact set of inputs, the piecewise linear function $x \mid z(x)$ cannot be highly oscillatory over a large portion of the range of z. Thus, if the bias b_z is not too concentrated on any interval, we expect the equation $z(x) = b_z$ to have O(1) solutions. On average, then, we expect that each neuron adds a constant number of new linear regions. Thus, the average total number of regions should scale roughly as the number of neurons.

Theorem 1 follows from a general result, Theorem 3, that holds for essentially any non-degenerate distribution of weights and biases and with any input dimension. If krz(x)k and the bias distribution $b_{\frac{1}{2}}$ are well-behaved, then throughout training, Theorem 3 suggests the number of linear regions along a 1-dimensional curve in input space scales like the number of neurons in N. Figures 5-6 show experiments that give empirical verification of this heuristic.

2.2. Higher-Dimensional Regions

For networks with input dimension exceeding 1; there are several ways to generalize counting linear regions. A unitmatching heuristic applied to Theorem 1 suggests

Proving this statement is work in progress by the authors. Instead, we consider here a natural and, in our view, equally important generalization. Namely, for a bounded K $R^{n_{in}}$, we consider the $(n_{in}$ 1)-dimensional volume density

$$vol_{n_{in} 1}(B_N \setminus K) vol_{n_{in}}(K);$$
 (3)

where

 $B_N = fx j r N (x)$ is not continuous at xg (4)

is the boundary of the linear regions for N . When $n_{in} = 1$,

 $vol_0(B_N \setminus K) + 1 = \# fregions in Kg;$

and hence the volume density (3) truly generalizes to higher input dimension of the number of regions. One reason for

studying the volume density (3) is that it gives bounds from below for distance (x; B_N), which in turn provides insight into the nature of the computation performed by N : Indeed, the exact formula

distance (x; B_N) =
$$\min_{\substack{ne \\ urons}} jz(x) \quad b_z j krz(x)k$$
;

shows that distance (x; B_N) measures the sensitivity over neurons at a given input x. In this formula, z(x) denotes the pre-activation for a neuron z and b_z is its bias, so that ReLU(z(x) b_z) is the post-activation. Moreover, the distance from a typical point to B_N gives a heuristic lower bound for the typical distance to an adversarial example: two inputs closer than the typical distance to a linear region boundary likely fall into the same linear region, and hence are unlikely to be classified differently. Our next result generalizes Theorem 1.

Theorem 2 (informal). Let N be a network with a piecewise linear activation, input dimension n_{in} and output dimension 1: Suppose its weights and biases are randomly initialized as in (2). Then, for K R^{d_{in}} bounded, the average volume of the linear region boundaries in K satisfies:

$$E \frac{\operatorname{vol}_{n} (B \setminus K)}{\operatorname{vol}_{n} \frac{1}{k}} T \text{ #fneuronsg;}$$

where T is the number of breakpoints in the non-linearity of N (for ReLU nets, T = 1). Moreover, if x 2 $[0; 1]^{n_{in}}$ is uniformly distributed, then the average, over both x and the weights/biases of N, distance from x to B_N satisfies

$$E[distance(x; B_N)] \quad C(#fneuronsg)^{\perp}; \quad C > 0:$$

Experimentally, distance $(x; B_N)$ remains comparable to (#fneuronsg) ¹ throughout training (see Figure 6).

3. Experiments

We empirically verified our theorems and further examined how linear regions of a network change during training. All experiments below were performed with fully-connected networks, initialized with He normal weights (i.i.d. with variance 2=fan-in) and biases drawn i.i.d. normal with variance 10⁶ (to prevent collapse of regions at initialization, which occurs when all biases are uniquely zero). Training was performed on the vectorized MNIST (input dimension 784) using the Adam optimizer at learning rate 10³. All networks attain test accuracy in the range 95 98%.

3.1. Number of Regions Along a Line

We calculated the number of regions along lines through the origin and and a random selected training example in input space. For each setting of weights and biases within the network during training, the number of regions along each



Figure 5. We here show how the number of regions along 1D lines in input space changes during training. In accordance with Theorem 3, we scale the number of regions by the number of neurons. Plots show (a) early training, up through 0.5 epochs, and (b) later training, up through 20 epochs. Note that for all networks, number of regions is a fixed constant times the number of neurons at initialization, as predicted, and that the number decreases (slightly) early in training before rebounding. [n₁; n₂; n₃] in the legend corresponds to an architecture with layer widths 784 (input); n₁; n₂; n₃; 10 (output).

Figure 6. We here consider the average distance to the nearest boundary, as evaluated over 10000 randomly selected sample points. In (a) we show that this distance is essentially bounded between 0:4=#fneuronsg and 1:5=#fneuronsg. Accordingly, in the next plot, we normalize the distance to the nearest boundary by dividing by the number of neurons. We plot this quantity against (b) epoch and (c) test accuracy. Observe that, in keeping with the findings of Figure 5, the distance to the nearest boundary first increases quickly (as the number of regions decreases), then rebounds more slowly as the network completes training. [n₁; n₂; n₃] in the legend corresponds to an architecture with layer widths 784 (input); n₁; n₂; n₃; 10 (output).

line is calculated exactly by building up the network one layer at a time and calculating how each region is split by the next layer of neurons. Figure 5 represents the average over 5 independent runs, from each of which we sample 100 lines; variance across the different runs is not significant.

Figure 5 plots the average number of regions along a line, divided by the number of neurons in the network, as a function of epoch during training. We make several observations:

- As predicted by Theorem 3, all networks start out with the number of regions along a line equal to a constant times the number of neurons in the network (the constant in fact appears very close to 1 in this case).
- Throughout training, the number of regions does not deviate significantly from the number of neurons in the network, staying within a small constant of the value at initialization, in keeping with our intuitive understanding of Theorem 3 described informally around Theorem 1 above.

- 3. The number of regions actually decreases during the initial part of training, then increases again. We explore this behavior further in other experiments below.
- 3.2. Distance to the Nearest Region Boundary

We calculated the average distance to the nearest boundary for 10000 randomly selected input points, for various networks throughout training. Points were selected randomly from a normal distribution with mean and variance matching the componentwise mean and variance of MNIST training data. Results were averaged over 12 independent runs, but variance across runs is not significant. Rerunning these experiments with sample points selected randomly from (i) the training data or (ii) the test data yielded similar results to random sample points.

In keeping with our results in the preceding experiment, the distance to the nearest boundary first increases then decreases during training. As predicted by Theorem 2, we find that for all networks, the distance to the nearest boundary is well-predicted by 1=#fneuronsg. Throughout training, we find that it approximately varies between the curves 0:4=#fneuronsg and 1:5=#fneuronsg (Figure 6(a)). At initialization, as we predict, all networks have the same value for the product of number of neurons and distance to the nearest region boundary (Figure 6(b)); these prod-ucts then diverge (slightly) for different architectures, first increasing rapidly and then decreasing more slowly.

We find Figure 6(c) fascinating, though we do not completely understand it. It plots the product of number of neurons and distance to the nearest region boundary against the test accuracy. It suggests two phases of training: first regions expand, then they contract. This lines up with observations made in Arpit et al. (2017) that neural networks "learn patterns first" on which generalization is simple and then refine the fit to encompass memorization of individual samples. A generalization phase would suggest that regions are growing, while memorization would suggest smaller regions are fit to individual data points. This is, however, speculation and more experimental (and theoretical) exploration will be required to confirm or disprove this intuition. We found it instructive to consider the full distribution of

Figure 7. Distribution of log distances from random sample points to the nearest region boundary for a network of depth 4 and width 16, at initialization and after 1 and 20 epochs of training on MNIST.

distances from sample points to their nearest boundaries, rather than just the average. For a single network (depth 4, width 16), Figure 7 indicates that this distribution does not significantly change during training, although there appears to be a slight skew towards larger regions, in agreement with the findings in Novak et al. (2018). The histogram shows log-distances. Hence, distance to the nearest region boundary varies over many orders of magnitude. This is consistent with Figures 1 and 4, which lend credence to the intuition that small distances to the nearest region boundary are explained by the presence of many small regions. According to Theorem 3, this should correlate with a combination of regions in input space at which some neurons have a large gradient and neurons with highly peaked biases distributions. We hope to return to this in future work.

3.3. Regions Within a 2D Plane

We visualized the regions of a network through training. Specifically, following experiments in Novak et al. (2018), we plotted regions within a plane in the 784-dimensional input space (Figure 8) through three data points with different labels (0, 1, and 2, in our case) inside a square centered at the circumcenter of the three examples. The network shown has depth 3 and width 64. We observe that, as expected from our other plots, the regions expand initially during training and then contract again. We expect the number of regions within a 2-dimensional subspace to be on the order of the square of the number of neurons – that is, $(643)^2$ 410⁴, which we indeed find.

Our approach for calculating regions is simple. We start with a single region (in this case, the square), and subdivide it by adding neurons to the network one by one. For each new neuron, we calculate the linear function it defines on each region, and determine whether that region is split into two. This approach terminates within a reasonable amount of time precisely because our theorem holds: there are relatively few regions. Note that we exactly determine all regions within the given square by calculating all region boundaries; thus our counts are exact and do not miss any small regions, as might occur if we merely estimated regions by sampling points from input space.

4. Related Work

There are a number of works that touch on the themes of this article: (i) the expressivity of depth; (ii) counting the number of regions in networks with piecewise linear activations; (iii) the behavior of linear regions through training; and (iv) the difference between expressivity and learnability. Related to (i), we refer the reader to Eldan & Shamir (2016); Telgarsky (2016) for examples of functions that can be efficiently represented by deep but not shallow ReLU nets. Next, still related to (i), for uniform approximation over classes of functions, again using deep ReLU nets, see Yarotsky (2017); Rolnick & Tegmark (2018); Yarotsky (2018); Petersen & Voigtlaender (2018). For interesting results on (ii) about counting the maximal possible number of linear regions in networks with piecewise linear activations see Bianchini & Scarselli (2014); Montufar et al. (2014); Poole et al. (2016); Arora et al. (2018); Raghu et al. (2017). Next, in the vein of (iii), for both a theoretical and empirical perspective on the number of regions computed by deep networks and specifically how the regions change during training, see Poole et al. (2016); Novak et al. (2018). In the direction of (iv), we refer the reader to Shalev-Shwartz et al. (2018); Hanin & Rolnick (2018); Hanin (2018). Finally, for general insights into learnability and expressivity in deep vs. shallow networks see Mhaskar & Poggio (2016); Mhaskar et al. (2016); Zhang et al. (2017); Lin et al. (2017); Poggio et al.

Epoch 0: 9744 regions Epoch 1: 4196 regions Epoch 20: 8541 regions

Figure 8. Here we show the linear regions that intersect a 2D plane through input space for a network of depth 3 and width 64 trained on MNIST. Black dots indicate the positions of the three MNIST training examples defining the plane. Note that we obtain qualitatively different pictures from Novak et al. (2018), which may result partially from our using ReLU activation instead of ReLU6.

(2017); Neyshabur et al. (2017).

5. Formal Statement of Results

To state our results precisely, we fix some notation. Let d; n_{in} ; n_1 ; :::; n_d 1 and consider a depth d fully connected ReLU net N with input dimension n_{in} , output dimension 1, and hidden layer widths n_j ; j = 1; :::; d 1: As explained in the introduction, a generic configuration of its weights and biases partitions the input space $R^{n_{in}}$ into a union of polytopes P_j with disjoint interiors. Restricted to each P_j ; N computes a linear function.

Our main mathematical result, Theorem 3, concerns the set B_N of points x 2 $R^{n_{in}}$ at which the gradient r N is discontinuous at x (see (4)). For each $k = 1; :::; n_{in}$; we define

$$B_{N:k} = \text{the} \setminus (n_{in} \ k) - \text{dimensional piece" of } B_N : (5)$$

More precisely, we set $B_{N;0} := ;$ and recursively define $B_{N;k}$ to be the set of points x 2 $B_N nfB_{N;0}$ [[$B_{N;k}$ 1g so that in a neighborhood of x the set $B_N nfB_{N;0}$ [[$B_{N;k}$ 1g coincides with a co-dimension k hyperplane.

For example, when $n_{in} = 2$; the linear regions P_j are polygons, the set $B_{N;1}$ is the union of the open line segments making up the boundaries of the P_j , and $B_{N;2}$ is the collection of vertices of the P_j : Theorem 3 provides a convenient formula for the average of the $(n_{in} \quad k)$ dimensional volume of $B_{N;k}$ inside any bounded, measurable set K $\mathbb{R}^{n_{in}}$. To state the result, for every neuron z in N we will write

$$z(x) := pre-activation at z;$$
 '(z) = layer index of z (6)

and b_z := bias at z: Thus, for a given input x 2 R^{n_0} , the post-activation of z is

$$Z(x) := ReLU(z(x)) = maxf0; z(x) b_zg:$$
 (7)

Theorem 3 holds under the following assumption on the distribution of weights and biases:

- A1: The conditional distribution of any collection of biases b_{z_1} ;:::; b_{z_k} , given all the other weights and biases, has a density b_{z_1} ;:::; b_z (b_1 ;:::; b_k) with respect to Lebesgue measure on R^k .
- A2: The joint distribution of all the weights has a density with respect to Lebesgue measure on R^{#weights}.

These assumptions hold in particular when the weights and biases of N are independent with marginal distributions that have a density relative to Lebesgue measure on R (i.e. at initialization). They hold much more generally, however, and can intuitively be viewed as a non-degeneracy assumption on the behavior of the weights and biases of N. Specifically, they are used in Proposition 10 to ensure that the set $B_{N\,;k}$ consists of inputs where exactly k neurons turn off/on. Assumption (A1) also allows us, in Proposition 11, to apply the co-area formula (29) to compute the expect volume of the set of inputs where a given collection of neurons turn on/off. Our main result is the following.

Theorem 3. Suppose N is a feed-forward ReLU net with input dimension n_0 ; output dimension 1, and random weights/biases. Assume that the distribution of weights/biases satisfies Assumptions A1 and A2 above. Then, with the notation (6), for any bounded measurable set K R^{n in} and any k = 1; :::; n_{in}; the average (n_{in} k)

dimensional volume $E[vol_{n_{in}} k(B_N; k \setminus K)]$ of $B_N; k$ inside K is

$$E\left[vol_{n_{in}} k(B_{N;k} \setminus K)\right]$$
(8)

of $B_{N;k}$ inside K is, in the notation (6),

where $Y_{z_1;...;z_k}(x)$ is

$$kJ_{z_1;...;z_k}(x)k_{b_{z_1};...;b_{z_k}}(z_1(x);:::;z_k(x))$$

times the indicator function of the event that z_j is good at x for each j = 1; :::; k. Here, $J_{z_1; :::; z_k}$ is the k n_{in} Jacobian of the map x ! $(z_1(x); :::; z_k(x));$

$$kJ_{z_1;...;z_k}(x)k := det J_{z_1;...;z_k}(x) (J_{z_1;...;z_k}(x))^{T^{1=2}};$$

the function $b_{z_1''';b_{z_k}}$ is the density of the joint distribution of the biases $b_{z_1}; :::; b_{z_k}$, and we say a neuron z is good at x if there exists a path of neurons from z to the output in the computational graph of N so that each neuron along this path is open at x).

To evaluate the expression in (8) requires information on the distribution of gradients rz(x), the pre-activations z(x), and the biases b_z : Exact information about these quantities is available at initialization (Hanin, 2018; Hanin & Rolnick, 2018; Hanin & Nica, 2018), yielding the following Corollary.

Corollary 4. With the notation and assumptions of Theorem 3, suppose the weights are independent are drawn from a fixed probability measure on R that is symmetric around 0 and then rescaled to have Var[weights] = 2=fan-in. Fix k 2 f1; :::; ning. Then there exists C > 0 for which

$$\frac{E \left[vol_{n_{in}} k \left(B_{N;k} \setminus K \right) \right]}{vol_{n_{in}}(K)}$$
(9)
#fneuronsg
k $\left(C_{grad} C_{bias} \right)^{k};$

where

$$C_{bias} = \sup_{z \ b2R} \sup_{b2R} (b)$$

and

$$C_{grad} = \sup_{z} \sup_{x \ge R \min} E^{h} krz(x)k^{2k} e^{i_{1}=k} C e^{C^{P}} e^{-i_{n_{j}} - i_{n_{j}}}$$

where C > 0 depends only on but not on the architecture of N and n_j is the width of the $j^{t\,h}$ hidden layer. Moreover, we also have similar lower bounds

#fneuronsg
$$E [vol_{n_{in}} k (B_{N;k} \setminus K)]$$

 $k c^{pias} -vol_{n_{in}}(K)$ (10)

where

$$c_{bias} = \inf_{\substack{j \in j \\ j \in j}} (b);$$

and

$$= \bigotimes^{0} \frac{\sup_{x \ge K} kxk^{2}}{n} + \bigotimes^{1}_{j=1} A_{b_{j}} e^{C^{0} P_{j=1} d_{n_{j}}}; \frac{1}{n}$$

with $C^0 > 0$ depending only on the distribution of the weights in N .

We prove Corollary 7 in Appendix D. Let us state one final corollary of Theorem 3

Corollary 5. Suppose N is as in Theorem 3 and satisfies the hypothesis (14) in Corollary 7. Then, for any compact set K $R^{n_{in}}$ let x be a uniform point in K: There exists c > 0 independent of K so that

$$E[distance(x; B_N)] = C \frac{c}{C_{bias}C_{grad} # fneuronsg}$$

We prove Corollary 8 in §E. The basic idea is simple. For every > 0; we have

E [distance(x; B_N)] P (distance(x; B_N) >);

with the probability on the right hand side scaling like

1
$$\operatorname{vol}_{n_{in}}(\mathsf{T}(\mathsf{B}_{\mathsf{N}}) \setminus \mathsf{K}) = \operatorname{vol}_{n_{in}}(\mathsf{K});$$

where $T(B_N)$ is the tube of radius around B_N : We expect that its volume like $vol_n_{in1}(B_N)$. Taking " = c=#fneuronsg yields the conclusion of Corollary 8.

6. Conclusions and Further Work

The question of why depth is powerful has been a persistent problem for deep learning theory, and one that recently has been answered by works giving enhanced expressivity as the ultimate explanation. However, our results suggest that such explanations may be misleading. While we do not speak to all notions of expressivity in this paper, we have both theoretically and empirically evaluated one common measure: the linear regions in the partition of input space defined by a network with piecewise linear activations. We found that the average size of the boundary of these linear regions depends only on the number of neurons and not on the network depth – both at initialization and during training. This strongly suggests that deeper networks do not learn more complex functions than shallow networks. We plan to test this interpretation further in future work - for example, with experiments on more complex tasks, as well as by investigating higher order statistics, such as the variance.

We do not propose a replacement theory for the success of deep learning; however, prior work has already hinted at how such a theory might proceed. Notably, Ba & Caruana (2014) show that, once deep networks are trained to perform a task successfully, their behavior can often be replicated by shallow networks, suggesting that the advantages of depth may be linked to easier learning.

References

Arora, R., Basu, A., Mianjy, P., and Mukherjee, A. Understanding deep neural networks with rectified linear units. In ICLR, 2018.

- Arpit, D., Jastrzębski, S., Ballas, N., Krueger, D., Bengio,
 E., Kanwal, M. S., Maharaj, T., Fischer, A., Courville, A.,
 Bengio, Y., et al. A closer look at memorization in deep networks. In ICML, 2017.
- Ba, J. and Caruana, R. Do deep nets really need to be deep? In NeurIPS, pp. 2654–2662, 2014.
- Bianchini, M. and Scarselli, F. 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.
- Eldan, R. and Shamir, O. The power of depth for feedforward neural networks. In COLT, pp. 907–940, 2016.
- Hanin, B. Which neural net architectures give rise to exploding and vanishing gradients? In NeurIPS, 2018.
- Hanin, B. and Nica, M. Products of many large random matrices and gradients in deep neural networks. Preprint arXiv:1812.05994, 2018.
- Hanin, B. and Rolnick, D. How to start training: The effect of initialization and architecture. In NeurIPS, 2018.
- Lin, H. W., Tegmark, M., and Rolnick, D. Why does deep and cheap learning work so well? Journal of Statistical Physics, 168(6):1223–1247, 2017.
- Mhaskar, H., Liao, Q., and Poggio, T. Learning functions: when is deep better than shallow. Preprint arXiv:1603.00988, 2016.
- Mhaskar, H. N. and Poggio, T. Deep vs. shallow networks: An approximation theory perspective. Analysis and Applications, 14(06):829–848, 2016.
- Montufar, G. F., Pascanu, R., Cho, K., and Bengio, Y. On the number of linear regions of deep neural networks. In NeurIPS, pp. 2924–2932, 2014.
- Neyshabur, B., Bhojanapalli, S., McAllester, D., and Srebro, N. Exploring generalization in deep learning. In NeurIPS, pp. 5947–5956, 2017.
- Novak, R., Bahri, Y., Abolafia, D. A., Pennington, J., and Sohl-Dickstein, J. Sensitivity and generalization in neural networks: an empirical study. In ICLR, 2018.
- Petersen, P. and Voigtlaender, F. Optimal approximation of piecewise smooth functions using deep ReLU neural networks. Neural Networks, 108:296–330, 2018.
- Poggio, T., Mhaskar, H., Rosasco, L., Miranda, B., and Liao, Q. Why and when can deep – but not shallow – networks avoid the curse of dimensionality: a review. International Journal of Automation and Computing, 14(5):503–519, 2017.

- Poole, B., Lahiri, S., Raghu, M., Sohl-Dickstein, J., and Ganguli, S. Exponential expressivity in deep neural networks through transient chaos. In NeurIPS, pp. 3360–3368, 2016.
- Raghu, M., Poole, B., Kleinberg, J., Ganguli, S., and Dickstein, J. S. On the expressive power of deep neural networks. In ICML, pp. 2847–2854, 2017.
- Rolnick, D. and Tegmark, M. The power of deeper networks for expressing natural functions. In ICLR, 2018.
- Serra, T., Tjandraatmadja, C., and Ramalingam, S. Bounding and counting linear regions of deep neural networks. In ICML, 2018.
- Shalev-Shwartz, S., Shamir, O., and Shammah, S. Failures of gradient-based deep learning. In ICML, 2018.
- Telgarsky, M. Representation benefits of deep feedforward networks. Preprint arXiv:1509.08101, 2015.
- Telgarsky, M. Benefits of depth in neural networks. In COLT, 2016.
- Yarotsky, D. Error bounds for approximations with deep ReLU networks. Neural Networks, 94:103–114, 2017.
- Yarotsky, D. Optimal approximation of continuous functions by very deep ReLU networks. In COLT, 2018.
- Zhang, C., Bengio, S., Hardt, M., Recht, B., and Vinyals, O. Understanding deep learning requires rethinking generalization. In ICLR, 2017.

A. Formal Statement of Results for General Piecewise Linear Activations

In §5, we stated our results in the case of ReLU activation, and now frame these results for a general piecewise linear non-linearity. We fix some notation. Let : R ! R be a continuous piecewise linear function with T breakpoints $_0$ = 1 < $_1$ < $_2$ < < $_T$ < $_{T+1}$ = 1: That is, there exist p_i; q_i 2 R so that

$$t 2 [j; j+1]$$
) $(t) = q_j t + p_j; q_j = q_{j+1}$: (11)

The analog of Theorem 3 for general is the following.

Theorem 6. Let : R ! R be a continuous piecewise linear function with T breakpoints $_1 < _T$ as in (11). Suppose N is a fully connected network with input dimension n_{in} ; output dimension 1, random weights and biases satisfying A1 and A2 above, and non-linearity.

Let $J_{z_1;\ldots;z_k}$ be the k n_{in} Jacobian of the map x ! $(z_1(x);\ldots;z_k(x));$

$$kJ_{z_1;...;z_k}(x)k := det J_{z_1;...;z_k}(x) (J_{z_1;...;z_k}(x))^{T^{1=2}}$$

and write $b_{z_1^{i_1i_2i_2}}$ for the density of the joint distribution of the biases b_{z_1} ; :::; b_{z_k} . We say a neuron z is good at x if there exists a path of neurons from z to the output in the computational graph of N so that each neuron **b** along this path is open at x (i.e. ${}^0(\mathbf{b}(x) \ \mathbf{b}_b) = 0$).

Then, for any bounded, measurable set K $R^{n_{in}}$ and any k = 1; : : : ; n_{in} ; the average $(n_{in} \quad k)$ -dimensional volume

$$E[vol_{n_{in}} k(B_{N;k} \setminus K)]$$

of $B_{N;k}$ inside K is, in the notation of (6),

 $\begin{array}{ccc} X & X^{T} & Z \\ & & EY \begin{pmatrix} \vdots, \vdots, \vdots, z_{k} \end{pmatrix} \\ \substack{distinct neurons \ i_{1}; \ldots; i_{k} = 1 \\ z_{1}; \ldots; z_{k} \end{pmatrix} \\ k & K \\ \end{array}$

where $Y_{z_{i}}^{(i_{1}, \dots, i_{k})}(x)$ equals

$$kJ_{z_1;...;z_k}(x)k_{b_{z_1};...;b_{z_k}}(z_1(x)_{i_1};...;z_k(x)_{i_k})$$

(13)

multiplied by the indicator function of the event that z_j is good at x for every j:

Note that if in the definition (11) of we have that the possible values ${}^0(t) 2 fq_0; :::; q_T g$ do not include 0, then we may ignore the event that z_j are good at x in the definition of $Y_{z_1;::;z_k}^{(i_1;i_2;i_3)}$.

Corollary 7. With the notation and assumptions of Theorem 6, suppose in addition that the weights and biases are independent. Fix k 2 f1; :::; n_{ing} and suppose that for every collection of distinct neurons z_1 ; :::; z_k , the average magnitude of the product of gradients is uniformly bounded:

Then we have the following upper bounds

where T is the number of breakpoints in the non-linearity of N (see (11)) and

$$C_{\text{bias}} = \sup_{z \text{ b2R}} \sup_{b \in \mathbb{R}} (b):$$

We prove Corollary 7 in §D and state a final corollary of Theorem 3:

Corollary 8. Suppose N is as in Theorem 3 and satisfies the hypothesis (14) in Corollary 7 with constants C_{bias} ; C_{grad} . Then, for any compact set K R^{n_{in}} let x be a uniform point in K: There exists c > 0 independent of K so that

$$E[distance(x; B_N)] = \frac{c T}{C_{bias}C_{grad}#fneuronsg'}$$

where, as before, T is the number of breakpoints in the non-linearity of N .

We prove Corollary 8 in §E. The basic idea is simple. For every > 0; we have

E [distance(x; B_N)] P (distance(x; B_N) >);

with the probability on the right hand side scaling like

1 $\operatorname{vol}_{n_{in}}(T(B_N) \setminus K) \operatorname{vol}_{n_{in}}(K);$

where $T(B_N)$ is the tube of radius around B_N : We expect that its volume like $vol_n_{in1}(B_N)$. Taking " = c=#fneuronsg yields the conclusion of Corollary 8.

B. Outline of Proof of Theorem 6

The purpose of this section is to give an intuitive explanation of the proof of Theorem 3. We fix a non-linearity : R ! R with breakpoints $_1 < _{T}$ (as in (11)) and consider a fully connected network N with input dimension n_{in} 1, output dimension 1, and non-linearity : For each neuron z in N, we write

$$(z) := layer index of z$$
 (16)

and set

$$S_z := fx \ 2 \ R^{n_{in}} \ j \ z(x) \qquad b_z \ 2 \ f_1; :::;_T \ gg:$$
 (17)

We further

$$\mathfrak{S}_{z} := S_{z} \setminus O; \tag{18}$$

where

Intuitively, the set S_z is the collection of inputs for which the neuron z turns from on to off. In contrast, the set O is the collection of inputs x 2 R^{n_{in}} for which N is open in the sense that there is a path from the input to the output of N so that all neurons along this path compute are not constant in a neighborhood x. Thus, S g is the set of inputs at which neuron z switches between its linear regions and at which the output of neuron z actually affects the function computed by N :

We remark here that O = ; if in the non-linearity there are no linear pieces at which the slopes on equals $O(i.e. q_j = 0 \text{ for all } j$ in the definition (11) of). If, for example, is ReLU, then O need not be empty.

The overall proof of Theorem 3 can be divided into several steps. The first gives the following representation of B_N :

Proposition 9. Under Assumptions A1 and A2 of Theorem 3, we have, with probability 1;

$$B_N = \mathcal{S}_z$$
:
neurons z

The precise proof of Proposition 9 can be found in §C.1 below. The basic idea is that if for all y near a fixed input x 2 $R^{n_{1n}}$; none of the pre-activations $z(y) = b_z$ cross the boundary of a linear region for , then x 2 B_N : Thus, $B_N = \sum_{z} S_z$: Moreover, if a neuron z satisfies $z(x) = b_z = S_1^z$ for some i but there are no open paths from z to the output of N for inputs near x, then z is dead at x and hence does not influence N at x: Thus, we expect the more refined inclusion $B_N = \sum_{z} S_{ze}$ Finally, if x 2 S_e for some z then x 2 B_N unless the contribution from other neurons to r N (y) for y near x exactly cancels the discontinuity in r z (x): This happens with probability 0.

The next step in proving Theorem 3 is to identify the portions of B_N of each dimension. To do this, we write for any distinct neurons $z_1; \ldots; z_k$,

$$\mathfrak{S}_{z_1;\ldots;z_k} := \bigvee_{j=1}^k \mathfrak{S}_{z_j}:$$

The set \mathbf{S}_{z_1} ;...; z_k is, intuitively, the collection of inputs at which $z_i(x) = b_{z_i}$ switches between linear regions for and

at which the output of N is affected by the post-activations of these neurons. Proposition 9 shows that we may represent B_N as a disjoint union

 $B_{N} = B_{N;k};$

where

In words, $B_{N;k}$ is the collection of inputs in O at which exactly k neurons turn from on to off. The following Proposition shows that $B_{N;k}$ is precisely the " $(n_{in} \ k)$ -dimensional piece of B_N " (see (5)).

Proposition 10. Fix $k = 1; :::; n_{in}$; and k distinct neurons $z_1; :::; z_k$ in N : Then, with probability 1; for every x 2 $B_{N;k}$ there exists a neighborhood in which $B_{N;k}$ coincides with a $(n_{in} \quad k)$ dimensional hyperplane.

We prove Proposition 10 in §C.2. The idea is that each $\Re_{z_1;:::;z_k}$ is piecewise linear and, with probability 1, at every point at which exactly the neurons $z_1;:::;z_k$ contribute to B_N , its co-dimension is the number of linear conditions needed to define it. Observe that with probability 1, the bias vector (b_z ;:::; b_z) for any collection z_1 ;:::; z_{k+1} of distinct neurons is a regular value for $x ! (z_1(x);:::;z_{k+1}(x))$. Hence,

$$vol_{n_{in}}$$

k $S_{z_1;...;z}e_{+1} = 0:$

Proposition 10 thus implies that, with probability 1;

$$vol_{n_{in}} k (B_{N;k}) = X vol_{n_{in}} k S_{z} e_{j \dots z_{k}}$$

The final step in the proof of Theorem 3 is therefore to prove the following result.

Proposition 11. Let z_1 ;:::; z_k be distinct neurons in N : Then, for any bounded, measurable K $R^{n_{in}}$,

$$\begin{array}{cccc} h & i \\ E & vol_{n_{in}} & k & \mathfrak{S}_{z_{1}; \dots; z_{k}} \\ & & Z_{T} & \chi & h & i \\ & & & E & Y & (s_{1}; \dots; s_{i_{k}}) \\ & & & i_{1}; \dots; i_{k} = 1 \end{array}$$

where $Y_{z_1; \vdots; z_k}^{\,(S_{\,i_1}; \vdots; S_{i_k})}$ is defined as in (13).

We provide a detailed proof of Proposition 11 in §C.3. The intuition is that the image of the volume element dx under x $|z(x) = S_i$ is the volume element

$$kJ_{z_1;\ldots;z_k}(x)k dx$$

from (13). The probability of an infinitesimal neighborhood dx of x belonging to a $(n_{in} \ k)$ -dimensional piece of B_N is therefore the probability

$$b_{z_1};...;b_{z_k}(z_1(x) = S_{i_1};...;z_k(x) = S_{i_k})$$

 $kJ_{z_1;...;z_k}(x)k dx$

that the vector of biases $(b_{z_j}; j = 1; :::; k)$ belongs to the image of dx under map $z_j(x) = S_{i_j}; j = 1; :::; k$ for some collection of breakpoints S_{i_j} : The formal argument uses the co-area formula (see (29) and (30)).

C. Proof of Theorem 3

C.1. Proof of Proposition 9

Recall that the non-linearity : R ! R is continuous and piecewise linear with T breakpoints $_{1} < _{T}$; so that, with $_{0} = _{1}$; $_{T+1} = _{1}$, we have

$$t 2 (i; i+1)$$
) $(t) = q_i t + p_i$

with $q_i = q_{i+1}$: For each x 2 R^{n_{in}}; write

 $Z^+ := z z(x)$ $b_z 2 (i; i+1)$ and $q_i = 0$ for some i Z_x x := z z(x) $b_z 2 (i; i+1)$ and $q_i = 0$ for some i $Z_x^0 := z z(x)$ $b_z = i$ for some i

Intuitively, Z_x^+ are the neurons that, at the input x are open (i.e. contribute to the gradient of the output N (x)) but do not change their contribution in a neighborhood of x, Z_x are the neurons that are closed, and Z^0 , are the neurons that, at x, produce a discontinuity in the derivative of N : Thus, for example, if = ReLU; then

$$Z_x := fz j sgn(z(x) b_z) = g;$$
 2 f+; ;0g: We

begin by proving that $B_N \xrightarrow{S} e$ by checking the contrapositive

$$\begin{bmatrix} & & & \\$$

Fix x 2 $\sum_{z}^{S} g_{z}^{c}$. Note that Z_{x} are locally constant in the sense that there exists " > 0 so that for all y with ky xk < ", we have

$$Z_{x} Z_{y}; Z_{x}^{+} Z_{y}; Z_{y}^{+}; Z_{y}^{+} [Z_{y}^{0} Z_{x} T_{x}]^{0}$$
 (20)

Moreover, observe that if in the definition (11) of none of the slopes q_i equal 0, then $Z = _{\gamma}$; for every γ . To prove (19), consider any path from the input to the output in the computational graph of N : Such a path consists of d + 1 neurons, one in each layer:

$$= z^{0}(; :::; z^{(d)}; (z^{(j)}) = j:$$

To each path we may associate a sequence of weights:

$$w^{(j)}$$
 := weight connecting $z^{(j-1)}$ to $z^{(j)}$; $j = 1; ...; d$:

We will also define

$$q^{(j)}(x) := \sum_{i=0}^{X^T} q_i \mathbf{1}_{f^{(x)}} b_{z^{(j)}} \mathbf{2}_{(i;i+1)} g^{(x)}$$

For instance, if = ReLU, then

$$q^{(j)}(x) = 1_{fz^{(j)}(x) \ b_2 \ Og};$$

and in general only one term in the definition of $q^{(j)}(x)$ is non-zero for each z: We may write

$$N(y) = \begin{array}{c} \chi^{in} & X & Y^{d} \\ y_{i} & q^{(j)}(y) \\ i=1 & paths: i ! out j = 1 \end{array} q^{(j)}(y) \\ (21)$$

Note that if x 2 $S_z \xi_z$ ^c, then for any path through a neuron z 2 Z_x^0 , we have

9 j s.t.
$$z^{(j)}$$
 2 Z_x :

This is an open condition in light of (20), and hence for all y in a neighborhood of x and for any path through a neuron z 2 Z_x we⁰also have that

9 j s.t.
$$z^{(j)}$$
 2 Z_v :

Thus, since the summand in (21) vanishes identically if Z = 3; we find that for y in a neighborhood of any x 2 $S_z S_z c$ we may write

$$N(y) = \begin{array}{ccc} X^{in} & X & Y^{d} \\ y_{i} & q^{(j)}(y)w^{(j)} + \text{ constant:} \\ z^{+} & z^{+} \end{array}$$

(22)

But, again by (20), for any fixed x, all y in a neighborhood of x and each z 2 Z_x^+ ; we have z 2 Z_y^+ as well. Thus, in particular,

$$z(x) \quad b_z \ 2(i; i+1) \quad) \quad z(y) \quad b_z \ 2(i; i+1)$$

Thus, for y sufficiently close to x; we have for every path in the sum (22) that

$$q^{(j)}(y) = q^{(j)}(x)$$
:

Therefore, the partial derivatives $(@N = @y_i)(y)$ are independent of y in a neighborhood of x and hence continuous at x. This proves (19). Let us now prove the reverse inclusion:

Note that, with probability 1; we have

$$vol_{n_{in}}(S_{z_1} \setminus S_{z_2}) = 0$$

for any pair of distinct neurons z_1 ; z_2 : Note also that since $x \mid N(x)$ is continuous and piecewise linear, the set B_N is closed. Thus, it is enough to show the slightly weaker inclusion $0 \quad 1$

$$\begin{bmatrix} & & & \\ & & & \\ & & & \\ & z & & & \\ & & & b = z \end{bmatrix} B_N$$
(24)

since the closure of Se $S_{b=z}$ Sb equals S_z : \mathbf{F} ix a neuron z and suppose x 2 Sz e $S_{b=z}$ Sb. By definition, we have that for every neuron $\mathbf{b} = \mathbf{z}$; either

This has two consequences. First, by (20), the map y ! z(y) is linear in a neighborhood of x: Second, in a neighborhood of x; the set \mathfrak{S}_z coincides with S_z . Hence, combining these facts, near x the set \mathfrak{S}_z coincides with the hyperplane

$$f x j z(x) \quad b_z = ig;$$
 for some i: (25)

We may take two sequences of inputs $y_n^+;\,y_n^-$ on opposite sides of this hyperplane so that

$$\lim_{n \ ! \ 1} y_n^+ = \lim_{n \ ! \ 1} y_n = x$$

and

$${}^{0}(z(y_{r}^{+}) b_{z}) = q_{i}; {}^{0}(z(y_{r}^{+}) b_{z}) = q_{i-1}; 8n;$$

where the index i the same as the one that defines the hyperplane (25). Further, since B_N has co-dimension 1 (it is contained in the piecewise linear co-dimension 1 set S_z , for example), we may also assume that y_n^+ ; $y_n^- 2 \ B_N : \tilde{C}$ onsider any path from the input to the output of the computational graph of N passing through z (so that $z = z^{(j)} 2$). By construction, for every n, we have

$$q^{(j)}(y_n^+) = q^{(j)}(y_n);$$

and hence, after passing to a subsequence, we may assume that the symmetric difference

$$Z_{v_n}^+ Z_{v_n}^+ = ;$$
 (26)

of the paths that contribute to the representation (21) for y_n^+ ; y_n^- is fixed and non-empty (the latter since it always contains z). For any y 2 B_N; we may write, for each $i = 1; :::; n_{in}$

$$\frac{@}{@N_i}(y) = X \qquad Y^d \qquad q^{(j)}(y)w^{(j)}: \qquad (27)$$

$$paths: ilout j=1$$

$$Z^+ y$$

Substituting into this expression $y = \gamma$, we find that there exists a non-empty collection of paths from the input to the output of N so that

$$\frac{@}{@N}(\gamma_n^+) = X \quad \forall \\ c^{(j)}w^{(j)} = \frac{@}{i} (\gamma_n^-) = A_j$$

where

a_i 2 f 1;1g; c^(j) 2 fq₀;:::;q_Tg:

Note that the expression above is a polynomial in the weights of N . Note also that, by construction, this polynomial is not identically zero due to the condition (26). There are only finitely many such polynomials since both a_j and $c^{(j)}$ range over a finite alphabet. For each such non-zero polynomial, the set of weights at which it vanishes has co-dimension 1. Hence, with probability 1; the difference $\frac{@N}{@y_i}(y_n^+) = \frac{@N}{@y_i}(y_n)$ is non-zero. This shows that the partial derivatives $\frac{@N}{@y_i}$ are not continuous at x and hence that $x \ge B_N$:

C.2. Proof of Proposition 10

Fix distinct neurons z_1 ;:::; z_k and suppose x 2 g_{z_1 ;:::; z_k but not in g_z for any $z = z_1$;:::; z_k : After relabeling, we may assume that they are ordered by layer index:

Since x 2 O, we also have that x 2 S_z for any z = $z_1; :::; z_k$: Thus, there exists a neighborhood U of x so $S_z \setminus U = ;$ for every z = $z_1; :::; z_k$: Thus, there exists a neighborhood of x on which y ! $z_1(y)$ is linear.

Hence, as explained near (25) above, \mathfrak{S}_{z_1} is a hyperplane near x: We now restrict our inputs to this hyperplane and repeat this reasoning to see that, near x; the set $\mathfrak{S}_{z_1;z_2}$ is a hyperplane inside \mathfrak{S}_{z_1} and hence, near x, is the intersection of two hyperplanes in $\mathbb{R}^{n_{1n}}$. Continuing in this way shows that in a neighborhood of x; the set $\mathfrak{S}_{z_1; \ldots; z_k}$ is equal to the intersection of k hyperplanes in $\mathbb{R}^{n_{1n}}$: Thus, $\mathfrak{S}_{z_1; \ldots; z_k} \mathfrak{n} \sum_{z=z_1; \ldots; z_k}^{c} \mathfrak{S}_z^{c}$ is precisely the intersection of

k hyperplanes in a neighborhood of each of its points.

C.3. Proof of Proposition 11

Let z_1 ;:::; z_k be distinct neurons in N; and fix a compact set K Rⁿ_{in}. We seek to compute the mean of

$$vol_{n_{in} \ k} \ \mathfrak{S}_{z_1; \ldots; z_k} \setminus K$$
, which we may rewrite as Z

$$= \begin{array}{c} 1n_{z_{j} \text{ is good at } x \text{ o}} dvol_{p_{n}} k(x) \qquad (28) \\ x^{T} Z \\ = \begin{array}{c} X^{T} Z \\ i_{1}; \dots; i_{k} = 1 \\ s_{z_{1}}^{(-1)} \cdots s_{z_{k}}^{(-1)} (x_{k}) \\ K \end{array} \\ K \qquad 1n_{z_{j} \text{ is good at } x^{0}} dvol_{n_{in}} k(x); \end{array}$$

where we've set

$$S_{z_{2}^{(i_{1};\dots;i_{k})}}^{(i_{1};\dots;i_{k})} = fx j z_{j}(x) \quad b_{z_{j}} = i_{j}; j = 1;\dots;kg:$$

Note that the map x ! $(z_1(x); :::; z_k(x))$ is Lipschitz, and recall the co-area formula, which says that if 2 $L^1(\mathbb{R}^n)$ and g : \mathbb{R}^n ! \mathbb{R}^m with m n is Lipschitz, then

equals

where J g is the m n Jacobian of g and

Ζ

$$kJg(x)k = det (Jg(x))(Jg(x))^{T^{1=2}}$$
:

We assumed that the biases b_{z_1} ; : : : ; b_{z_j} have a joint conditional density

$$b_{z} = b_{z_{1}}; ...; b_{z_{k}}$$

given all other weights and biases. The mean of the term in (28) corresponding to a fixed = $(i_1; :::; i_k)$ over the conditional distribution of $b_{z_1}; :::; b_{z_j}$ is therefore Z Z

$$\frac{db_b}{k} \begin{pmatrix} b \end{pmatrix} \qquad 1 \operatorname{n}_{z_j \text{ is good at } x \text{ o } } \frac{dvol_{n}}{j_n}$$

$$f_z \quad b = g \setminus K \quad j = 1; \dots; k$$

where we've abbreviated $b = (b_1; ...; b_k)$ as well as $z(x) = (z_1(x); ...; z_k(x))$. This can rewritten as Z = Zdb $b = (z(x) - 1) a_k$ is greated by $a_k(x)$:

$$\begin{array}{ccc} \text{AD} & & & & b \quad \left(Z(X) \right) \text{In} \quad z_j \text{ is good at } x^O \text{AVOI}_n \quad _{d} (X) \text{:} \\ \text{R}^k & & & \text{fz} = b \quad g \setminus K & & & j = 1 \text{; ::: } \text{; } k \end{array}$$

Thus, applying the co-area formula (29), (30) shows that the average of (28) over the conditional distribution of b_{z_1} ; :::; b_{z_j} is precisely

Taking the average over the remaining weighs and biases, we may commute the expectation E [] with the dx integral since the integrand is non-negative. This completes the proof of Proposition 11.

D. Proof of Corollary 7

We begin by proving the upper bound in (15). By Theorem 3, E [vol ($B_{N;k} \setminus K$)] equals

$$\begin{array}{ccc} X & X^T & Z & h & i\\ & E & Y_{z_{\dot{z}\dot{j}}}^{(i_{\dot{z}}^{(i_{\dot{z}}^{(i_{j},\dots,i_{k})})}(x))}(x) & (x)dx; \end{array}$$

distinct neurons z₁;...;z_k i₁;...;i_k=1 K

where, as in (13), $Y_{z_1;...;z_k}^{(i_1;...;i_k)}(x)$ is

$$kJ_{z_1;...;z_k}(x)k_{b_{z_1};...;b_{z_k}}(z_1(x)_{i_1};...;z_k(z)_{i_k})$$

times the indicator function of the even that z_j is good at x for every j: When the weights and biases of N are independent, we may write b_{z_1} ;...; b_{z_k} (b_1 ; : : : ; b_k) as

$$\begin{array}{ccc} Y_k & & & k \\ & & (b_1) & & sup & sup & (b_2) = C^k & : \\ y_{j=1} & & b_{z_j} & j & & neurons z \ b 2 R & b_z & & b ias \end{array}$$

Hence,

$$Y_{z_1;...;z_k}(x) = C_{biaks} \det J_{z_1;...;z_k}(x) (J_{z_1;...;z_k}(x))^{T^{1=2}}$$
: (30)
Note that

$$J_{z_1;...;z_k}(x) (J_{z_1;...;z_k}(x))^T = Gram (rz_1(x); :::; rz_k(x));$$

where for any v_i 2 Rⁿ

$$Gram(v_1; :::; v_k)_{i;j} = hv_i; v_j i$$

is the associated Gram matrix. The Gram identity says that det $J_{z_1;...;z_k}(x) (J_{z_1;...;z_k}(x))^T$ equals

$$krz_{1}(x) ^{r} rz_{k}(x)k;$$

which is the the k-dimensional volume of the parallelopiped in $R^{n_{in}}$ spanned by $frz_i(x)$; j = 1; :::; kg: We thus have

det
$$J_{z_1;...;z_k}(x) (J_{z_1;...;z_k}(x))^T \overset{1=2}{\underset{j=1}{\overset{Y^k}{\sum}} kr z_j(x) k:$$

The estimate (14) proves the upper bound (15). For the special case of = ReLU we use the AM-GM inequality and Jensen's inequality to write

Therefore, by Theorem 1 of Hanin & Nica (2018), there exist C_1 ; $C_2 > 0$ so that

$$2 3
F 4 krz_j(x)k5 C_1 e^{C_2 P_{j=1} \frac{1}{n_j}}$$

This completes the proof of the upper bound in (15). To prove the power bound, lower bound in (15) we must argue in a different way. Namely, we will induct on k and use the following facts to prove the base case k = 1:

1. At initialization, for each fixed input x; the random variables $f_{1_{f_z(x)>b_zg}}$ g are independent Bernoulli random variables with parameter 1=2: This fact is proved in Proposition 2 of Hanin & Nica (2018). In particular, the event fz is good at xg, which occurs when there exists a layer j 2 '(z) + 1;:::; d in which z(x) b_z for every neuron, is independent of fz(x); b_zg and satisfies

P (z is good at x) 1
$$X^{d}$$
 2 n_{j} : (31)

2. At initialization, for each fixed input x, we have

$$\frac{1}{2}E \quad z(x)^2 = \frac{kxk^2}{n_{in}} + \frac{\dot{x}^{(z)}}{\sum_{j=1}^{2} \dot{B}_j}$$
(32)

where ${}^{2}_{b_{j}}$:= Var[biases at layer j]. This is Equation (11) in the proof of Theorem 5 from Hanin & Rolnick (2018).

 At initialization, for every neuron z and each input x; we have
 h
 i

$$E krz(x)k^2 = 2$$
: (33)

This follows easily from Theorem 1 of Hanin (2018).

4. At initialization, for each 1 j n_{in} and every x 2 $R^{n_{in}}$

$$E \log n_{in} \frac{@z}{@x_{j}}(x)^{2} = \frac{5}{2} \frac{X^{(z)}}{\sum_{j=1}^{j} n_{j}} (34)$$

plus $O^{P'(z)}_{n^2 \underline{1}}$ where n_j is the width of the jth hidden layer and the implied constant depends only on the 4th moment of the measure according to which weights are distributed. This estimate follows immediately by combining Corollary 26 and Proposition 28 in Hanin & Nica (2018).

We begin by proving the lower bound in (15) when k = 1: We use (31) to see that $E[vol_{n_{in}} \ 1(B_N \setminus K)]$ is bounded below by

$$1 \begin{array}{ccc} X & d & X \\ 2 & n_{j} & E[krz(x)k_{b}(z(x))] dx: \\ j=1 & neurons z \\ K \end{array}$$

Next, we bound the integrand. Fix x 2 $R^{n_{1n}}$ and a parameter > 0 to be chosen later. The integrand E [krz(x)k_{bz}(z(x))] is bounded below by

$$E krz(x)k_{b_{z}}(z(x))\mathbf{1}_{fjz(x)jg}_{\substack{b_{z}\\j\notin nf}}(b) E krz(x)k \mathbf{1}_{fjz(x)jg};$$

which is bounded below by

 $\inf_{b_z}(b)E[krz(x)k] = Ekrz(x)k \mathbf{1}_{fjz(x)jg>}$:

Uşing Cauchy-Schwarz, the term E krz(x)k $1_{fjz(x)jg>}$ is bounded above by

$$E[krz(x)k]^{2} P(jz(x)j >)^{1=2}$$

which using (33) and (32) together with Markov's inequality, is bounded above by

$$\frac{0}{\frac{2}{1=2}} @\frac{kxk^2}{n_{in}} + \frac{X^{z}}{\sum_{j=1}^{b_i} A} :$$

Next, using Jensen's inequality twice, we write

$$\log E[krz(x)k] = \frac{1}{2} E \int_{2}^{n} \log krz(x)k^{2} = \frac{1}{2} E 4 \log \frac{(m_{1}^{2} + m_{2}^{2})}{(m_{2}^{2} + m_{2}^{2})} = \frac{1}{2} E \frac{1}{2} \log \frac{(m_{1}^{2} + m_{2}^{2})}{(m_{2}^{2} + m_{2}^{2})} = \frac{1}{2} \frac{1}{2} \log \frac{(m_{1}^{2} + m_{2}^{2})}{(m_{2}^{2} + m_{2}^{2})} = \frac{1}{2} \frac{1}{2} \log \frac{(m_{1}^{2} + m_{2}^{2})}{(m_{2}^{2} + m_{2}^{2})} = \frac{1}{2} \frac{1}{2} \log \frac{(m_{1}^{2} + m_{2}^{2})}{(m_{2}^{2} + m_{2}^{2})} = \frac{1}{2} \log \frac{(m_{1}^{2} + m_{2}^{2})}{(m_{1}^{2} + m_{2}^{2})} = \frac{1}{2} \log \frac{(m_{1}^{2} + m_{2}^{2})$$

where in the last inequality we applied (34). Putting this all together, we find that exists c > 0 so that

$$E[krz(x)k_{b_{2}}(z(x))] c \inf_{jbj} b_{2}(b);$$

where

4
$$\binom{0}{n_{in}} + \frac{X^{d}}{j=1} e^{\sum_{b_{j}} P} e^{\sum_{a_{j}} P} j^{a_{j}} \frac{1}{1 - \frac{1}{n_{j}} + 0} j^{a_{j}} j^{a_{j}} j^{a_{j}}$$

In particular, we may take

$$= @ \frac{\sup_{x_{2K}} kxk^{2}}{n_{in}} + \frac{\frac{1}{x_{2K}} e^{C_{j=1}^{P_{j=1}^{d_{1}}}}}{\sum_{j=1}^{n_{j}}}$$

for C sufficiently large. This completes the proof of the lower bound in (15) when k = 1. To complete the proof of Corollary 7, suppose we have proved the lower bound in (15) for all ReLU networks N and all collections of k 1 distinct neurons. We may assume after relabeling that the neurons z_1 ; :::; z_k are ordered by layer index:

With probability 1; the set $S_z = R^{n_{in}}$ is piecewise linear, co-dimension 1 with finitely many pieces, which we denote

by P. We may therefore rewrite
$$vol_{n_{in}} \ k \ S_2 F_1;...;z_k \setminus K$$

as $\chi \ vol_{n_{in}} \ k \ S_{z_2};...;z_k \setminus P \setminus K :$

We now define a new neural network N, obtained by restricting N to P: The input dimension for N equals n_{in} 1; and the weights and biases of N satisfy all the assumptions of Corollary 7. We can now apply our inductive hypothesis to the k 1 neurons z_2 ; :::; z_k in N and to the set K \setminus P: This gives

$$E \times \operatorname{vol}_{n_{in} k} \operatorname{Sp}_{2}; \dots; z_{k} \setminus P \setminus K$$

$$\inf_{z \text{ jbj}} \inf_{b_{2}^{k}} b) \quad E[\operatorname{vol}_{n_{in} 1}(P \setminus K)]:$$

Ħ

Summing this lower bound over yields h

Applying the inductive hypothesis once more completes the proof.

E. Proof of Corollary 8

We will need the following observation.

Lemma 12. Fix a positive integer n 1, and let S Rⁿ be a compact continuous piecewise linear submanifold with finitely many pieces. Define $S_0 = ;$ and let S_k be the union of the interiors of all k-dimensional pieces of Sn(S₀ [[S_k 1). Denote by T_"(X) the " tubular neighborhood of any X Rⁿ: We have

where $!_d$:= volume of ball of radius 1 in R^d :

Proof. Define d to be the maximal dimension of the linear pieces in S: Let x 2 T_"(S): Suppose x 2 T_"(S_k) for all k = 0; :::; d 1: Then the intersection of the ball of radius " around s with S is a ball inside S_d = U R^d. Using the convexity of this ball, there exists a point y in S_d so that the vector x y is parallel to the normal vector to S_d at y. Hence, x belong to the normal "-ball bundle B_"(N(S_d)) (i.e. the union of the fiber-wise "-balls in the normal bundle to S_d). Therefore, we have

$$vol_n(T_{*}(S)) vol_n(B_{*}(N(S_d))) + vol_n(T_{*}(S_{d-1}));$$

where we abbreviated $S_{d-1} := S \xrightarrow[k=0]{d-1} \frac{d}{2} \frac{d}{2$

and repeating this argument d $\ \ 1$ times completes the proof. $\hfill \Box$

We are now ready to prove Corollary 2. Let x 2 K = $[0; 1]^{n_{in}}$ be uniformly chosen. Then, for any " > 0, using Markov's inequality and Lemma 12, we have

$$E [distance(x; B_N)]
"P (distance(x; B_N) > ")
= "(1 P (distance(x; B_N) ")) = "(1 E [vol_{n_{in}} (T" (B_N))])
" 1 $\frac{\chi^{i_{in}}}{\sum_{k=1}^{l_{n_{in}-k}} \sum_{k=1}^{l_{n_{in}-k}} vol_{n_{in-k}} (B_{N;k})}{\sum_{k=1}^{l_{in}} \sum_{k=1}^{l_{in}} (C_{grad}C_{bias}"#fneuronsg)^{k}}$

$$= "(1 C^{0}C_{grad}C_{bias}"#fneuronsg)$$$$

for some $C^0 > 0$: Taking " to be a small constant times $1=(C_{grad}#fneuronsg)$ completes the proof.