

---

# Adversarial Examples in Multi-Layer Random ReLU Networks

---

**Peter L. Bartlett**

Department of Electrical Engineering and Computer Science  
Department of Statistics  
UC Berkeley

**Sébastien Bubeck**

Microsoft Research Redmond

**Yeshwanth Cherapanamjeri**

Department of Electrical Engineering and Computer Science  
UC Berkeley

## Abstract

We consider the phenomenon of adversarial examples in ReLU networks with independent Gaussian parameters. For networks of constant depth and with a large range of widths (for instance, it suffices if the width of each layer is polynomial in that of any other layer), small perturbations of input vectors lead to large changes of outputs. This generalizes results of Daniely and Schacham (2020) for networks of rapidly decreasing width and of Bubeck et al (2021) for two-layer networks. Our proof shows that adversarial examples arise in these networks because the functions they compute are *locally* very similar to random linear functions. Bottleneck layers play a key role: the minimal width up to some point in the network determines scales and sensitivities of mappings computed up to that point. The main result is for networks with constant depth, but we also show that some constraint on depth is necessary for a result of this kind, because there are suitably deep networks that, with constant probability, compute a function that is close to constant.

## 1 Introduction and Main Result

Since the phenomenon of adversarial examples was first observed in deep networks [SZS<sup>+</sup>14], there has been considerable interest in why this extreme sensitivity to small input perturbations arises in deep networks [GSS15, SSRD19, BLPR19, DS20, BCGdC21] and how it can be detected and avoided [CW17a, CW17b, FCSG17, MMS<sup>+</sup>18, QMG<sup>+</sup>19]. Building on the work of Shamir et al [SSRD19], Daniely and Schacham [DS20] prove that small perturbations (measured in the Euclidean norm) can be found for any fixed input and most Gaussian parameters in certain ReLU networks—those in which each layer has vanishing width relative to the previous layer—and conjectured the same result without this strong constraint on the architecture. Bubeck, Cherapanamjeri, Gidel and Tachet des Combes [BCGdC21] prove that the same phenomenon occurs in general two-layer ReLU networks, and give experimental evidence of its presence in deeper ReLU networks.

In this paper, we prove that adversarial examples also arise in deep ReLU networks with random weights for a wide variety of network architectures—those with constant depth and polynomially-related widths. The key fact underlying this phenomenon was already observed in [DS20]: a high-dimensional linear function  $f(x) = w^\top x$  with input  $x \neq 0$  and random parameter vector  $w$  with a uniformly chosen direction will satisfy  $\|\nabla f(x)\| \|x\| \gg |f(x)|$  with high probability. This implies

the existence of a nearby adversarial example for this linear function: a perturbation of  $x$  of size  $|f(x)|/\|\nabla f(x)\| \ll \|x\|$  in the direction  $-f(x)\nabla f(x)$  will flip the sign of  $f(x)$ . This observation can be extended to nonlinear functions that are locally almost linear. Indeed, it is easy to show that for all  $x, u \in \mathbb{R}^d$ ,

$$|f(x+u) - (f(x) + \langle u, \nabla f(x) \rangle)| \leq \|u\| \sup \{ \|\nabla f(x) - \nabla f(x+v)\| : v \in \mathbb{R}^d, \|v\| \leq \|u\| \},$$

and thus to demonstrate the existence of an adversarial example near  $x$  for a function  $f$ , it suffices to show the smoothness property:

$$\text{for all } v \in \mathbb{R}^d \text{ with } \|v\| \lesssim |f(x)|/\|\nabla f(x)\|, \quad \|\nabla f(x) - \nabla f(x+v)\| \ll \|\nabla f(x)\|. \quad (1)$$

We show that for a deep ReLU network with random parameters and a high-dimensional input vector  $x$ , there is a relatively large ball around  $x$  where the function computed by the network is very likely to satisfy this smoothness property. Thus, adversarial examples arise in deep ReLU networks with random weights because the functions that they compute are very close to linear in this sense.

It is important to notice that ReLU networks are not smooth in a classical sense, because the nondifferentiability of the ReLU nonlinearity implies that the gradient can change abruptly. But for the smoothness condition (1), it suffices to have  $\|\nabla f(x) - \nabla f(x+v)\| \leq \epsilon + \phi(\|v\|)$ , for some increasing function  $\phi : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ , provided that  $\epsilon + \phi(\|v\|) \ll \|\nabla f(x)\|$ . We prove an inequality like this for ReLU networks, where the  $\epsilon$  term decreases with width.

Consider a network with input dimension  $d$ ,  $\ell+1$  layers, a single real output, and complete connections between layers. Let  $d_1, \dots, d_\ell$  denote the dimensions of the layers. The network has independent random weight matrices  $W_i \in \mathbb{R}^{d_i \times d_{i-1}}$  for  $i \in [\ell+1]$ , where we set  $d_0 = d$  and  $d_{\ell+1} = 1$ . For input  $x \in \mathbb{R}^d$ , the network output is defined as follows:

$$f(x) = W_{\ell+1} \cdot \sigma(W_\ell \cdot \sigma(W_{\ell-1} \cdot \sigma(\dots \sigma(W_1 \cdot x) \dots))) \text{ where } \sigma(x)_i = \max\{x_i, 0\}$$

$$W_{\ell+1} \sim \mathcal{N}(0, I/d_\ell) \text{ and } \forall i \in [\ell], (W_i)_{j,k} \stackrel{i.i.d.}{\sim} \mathcal{N}(0, 1/d_{i-1}). \quad (\text{NN-DEF})$$

Note that the scale of the parameters is chosen so that all the real-valued signals that appear throughout the network have roughly the same scale. This is only for convenience: because the ReLU is positively homogeneous (that is, for  $\alpha > 0$ ,  $\sigma(\alpha x) = \alpha \sigma(x)$ ), the scaling is not important for our results; the  $1/d_{i-1}$  in (NN-DEF) could be replaced by any constant without affecting the ratio between the norm of an input vector and that of a perturbation required to change the sign of the corresponding output.

The following theorem is the main result of the paper.

**Theorem 1.1.** *Fix  $\ell \in \mathbb{N}$ . There are constants  $c_1, c_2, c_3$  that depend on  $\ell$  for which the following holds. Fix  $\delta \in (0, 1)$  and let  $f(\cdot)$  be an  $(\ell+1)$ -layer ReLU neural network defined by (NN-DEF) with input dimension  $d$  and intermediate layers of width  $\{d_i\}_{i=1}^\ell$ . Suppose that the widths satisfy*

$$d_{\min} \geq c_1 (\log d_{\max})^{c_2} \log 1/\delta \text{ where } d_{\min} = \min \{ \{d_i\}_{i=1}^\ell, d \}, d_{\max} = \max \{ \{d_i\}_{i=1}^\ell, d \}.$$

*Then for any fixed input  $x \neq 0$ , with probability at least  $1 - \delta$ ,*

$$|f(x + \eta \nabla f(x))| \geq |f(x)| \text{ and } \text{sign}(f(x + \eta \nabla f(x))) \neq \text{sign}(f(x)),$$

*for an  $\eta$  satisfying*

$$\frac{\|\eta \nabla f(x)\|}{\|x\|} \leq c_3 \sqrt{\frac{\log 1/\delta}{d}}.$$

*It suffices to choose  $c_1 = (C_1 \ell)^{c_2}$ ,  $c_2 = C_2 \ell$ ,  $c_3 = C_3^\ell$ , for some absolute constants  $C_1, C_2, C_3$ .*

This theorem concerns networks of fixed depth, and the constants in the size of the perturbation and in the requirement on the network width are larger for deeper networks. We also prove a converse result that illustrates the need for some constraint on the depth. Theorem 3.1 shows that when the depth is allowed to grow polynomially in the input dimension  $d$ , the function computed by a random ReLU network is essentially constant, which rules out the possibility of adversarial examples.

The heart of the proof of Theorem 1.1 is to show a smoothness property like (1). It exploits a decomposition of the change of gradient between two input vectors. Define  $H_i : \mathbb{R}^d \rightarrow \mathbb{R}^{d_i \times d_i}$

as  $H_i(x)_{jk} = \mathbf{1}\{j = k, v_i(x)_j \geq 0\}$  with  $v_i(x) = W_i \sigma(\dots \sigma(W_1 x))$ . For two input vectors  $x, y \in \mathbb{R}^d$ , we will see in Section 2.4 that we can decompose the change of gradient as

$$\nabla f(x) - \nabla f(y) = \sum_{j=1}^{\ell} W_{\ell+1} \left( \prod_{i=\ell}^{j+1} H_i(x) W_i \right) \cdot (H_j(x) - H_j(y)) W_j \cdot \left( \prod_{i=j-1}^1 H_i(y) W_i \right).$$

(Here and elsewhere, indices of products of matrices run backwards, so  $\prod_{i=j}^k M_i = I$  when  $j < k$ .) For the  $j$ th term in the decomposition, we need to control the scale of: the gradient of the mapping from the input to the output of layer  $j$ , the change in the layer  $j$  nonlinearity  $H_j(x) - H_j(y)$ , and the gradient from layer  $j$  to the output. It turns out that controlling these quantities depends crucially on the width of the narrowest layer before layer  $j$ —we call this the *bottleneck layer* for layer  $j$ . This width determines the dimension of the image at layer  $j$  of a ball in the input space. In proving bounds on gradients and function values that hold uniformly over pairs of nearby vectors  $x$  and  $y$ , this dimension—the width of the bottleneck layer—dictates the size of a discretization (an  $\epsilon$ -net) that is a crucial ingredient in the proof of these uniform properties. Our analysis involves working separately with the segments between these bottleneck layers. We show that for an input  $x \in \mathbb{R}^d$  satisfying  $\|x\| = \sqrt{d}$  and any  $y$  in a ball around  $x$ , with high probability  $\|\nabla f(x) - \nabla f(y)\| = o(1)$ , but  $|f(x)|$  is no more than a constant and  $\|\nabla f(x)\|$  is at least a constant. This implies the existence of a small ( $o(\|x\|)$ ) perturbation of  $x$  in the direction  $-f(x)\nabla f(x)$  that flips the sign of  $f(x)$ .

These results suggest several interesting directions for future work. First, our results show that for high-dimensional inputs, adversarial examples are inevitable in random ReLU networks with constant depth, and unlikely in networks with polynomial depth. Beyond this, we do not know how the sensitivity to input perturbations decreases with depth. Similarly, both results are restricted to networks with subexponential width, and it is not clear what happens for very wide networks. Finally, we show that networks with random weights suffer from adversarial examples because their behavior is very similar to that of random linear functions. It would be worthwhile to determine whether randomly initialized trained networks retain this nearly linear behavior, and hence suffer from adversarial examples for the same reason.

**Related Work:** Related theoretical work include the recent result of Bubeck and Sellke [BS21] who, following up on Bubeck, Li and Nagaraj [BLN21], show that only *mildly* over-parameterized networks when trained on random data have large Lipschitz constants. While these results apply to a broader class of networks including those potentially trained on data, this weaker property on its own does not suffice to explain the prevalence of adversarial examples and computational ease of finding them. Indeed, establishing this fact requires understanding the behavior of the local landscape of the function computed by the network which these approaches do not capture.

Closely related empirical works include work by Madry, Makelov, Schmidt, Tsipras and Vladu [MMS<sup>+</sup>18] and by Qin, Martens, Gowal, Krishnan, Dvijotham, Fawzi, De, Stanforth and Kohli [QMG<sup>+</sup>19]. We note that both these works emphasize the importance of the interplay between the local behavior of the function and the existence of adversarial perturbations. [MMS<sup>+</sup>18] identify local near-linearity as a cause of adversarial examples and propose a robust training procedure attempting to eliminate this property. On the other hand, [QMG<sup>+</sup>19] suggest an alternative robust training procedure that retains local near-linearity but eliminates adversarial examples by ensuring that the learnt network (despite being locally linear) is robust to single-step perturbations. Our work lends theoretical grounding to this phenomenon showing that local linearity arises naturally at initialization.

## 2 Proof of Main Theorem

In this section, we provide an outline of the proof of [Theorem 1.1](#). As described above, we will prove our result first by showing that the gradient at  $x$  has large norm and changes negligibly in a large ball around  $x$ . The first condition is established in [Subsection 2.1](#). The second step is more intricate. First, we prove the decomposition of the gradient differences in [Subsection 2.2](#). Then, in [Subsection 2.3](#), we track the scale of the ball around  $x$  as it propagates through the network. Finally, in [Subsection 2.4](#),

we use this result to bound the terms in the decomposition of the gradient differences to show that our network is locally linear. For the rest of the proof, unless otherwise stated, we consider a fixed  $x \in \mathbb{R}^d$ , and we assume:

$$d_{\min} \geq (C\ell \log d_{\max})^{240\ell} \log 1/\delta, \quad \|x\| = \sqrt{d} \text{ and } R := \frac{\sqrt{d_{\min}}}{(\ell \log d_{\max})^{80\ell}} = \Omega((\ell \log d_{\max})^{40\ell}).$$

Additionally, we randomize the activations of neurons whenever they receive an input of 0. This does not change the behavior of the neural network in terms of its output or the images of the input through the layers of the network but greatly simplifies our proof. For  $x \in \mathbb{R}^d$ , we let:

$$x_0 := x, \quad \tilde{f}_i(x) := W_i f_{i-1}(x), \quad (D_i(y))_{j,k} = \begin{cases} 1 & \text{w.p. } \frac{1}{2} \text{ if } j = k, y_j = 0, \\ 1 & \text{if } j = k, y_j > 0, \\ 0 & \text{otherwise,} \end{cases} \quad f_i(x) = D_i(\tilde{f}_i(x)) \tilde{f}_i(x).$$

Our first key observation is that the randomization in the activation units allows us the following distributional equivalences, proved in [Appendix A.1](#).

**Lemma 2.1.** *Let  $m \in \mathbb{N}$ ,  $\{d_i\}_{i=0}^m \subset \mathbb{R}^d$  and  $W_i \in \mathbb{R}^{d_i \times d_{i-1}}$  be distributed such that each entry of  $W_i$  is drawn iid from any symmetric distribution. Then, defining for  $x \in \mathbb{R}^{d_0}$ :*

$$\begin{aligned} h_0(x) &= x, \\ \tilde{h}_i(x) &= W_i h_{i-1}(x) \quad \text{where } (D_i(y))_{i,j} = \begin{cases} 1, & \text{if } j = k \text{ and } y_j > 0 \\ 1, & \text{with probability } \frac{1}{2} \text{ if } j = k, y_j = 0 \\ 0, & \text{otherwise} \end{cases} \\ h_i(x) &= D_i(\tilde{h}_i(x)) \tilde{h}_i(x) \end{aligned}$$

we have the distributional equivalences for any  $x \neq 0$  and fixed diagonal matrices  $B_1, \dots, B_m$ :

$$\begin{aligned} W_m \prod_{j=m-1}^1 (D_j(\tilde{h}_j(x)) + B_j) W_j &\stackrel{d}{=} W_m \prod_{j=m-1}^1 (D_j + B_j) W_j \\ \left\| \prod_{j=m}^1 (D_j(\tilde{h}_j(x)) + B_j) W_j \right\| &\stackrel{d}{=} \left\| \prod_{j=m}^1 (D_j + B_j) W_j \right\| \\ \text{where } (D_j)_{k,l} &= \begin{cases} 1, & \text{with probability } 1/2 \text{ if } k = l \\ 0, & \text{otherwise} \end{cases} \end{aligned}$$

## 2.1 Concentration of Function Value and Gradient at a Fixed Point

We first present a simple lemma that shows that the gradient at  $x$  is at least a constant and that its output value is bounded. The proof gives an illustration of how [Lemma 2.1](#) will be used through the more involved proofs in the paper.

**Lemma 2.2.** *For some universal constant  $c$ , with probability at least  $1 - \delta$  we have:*

$$|f(x)| \leq c2^\ell \sqrt{\log 1/\delta} \text{ and } \|\nabla f(x)\| \geq \frac{1}{2^{\ell+1}}.$$

*Proof.* Note that

$$\nabla f(x) = W_{\ell+1} \prod_{i=\ell}^1 D_i(\tilde{f}_i(x)) W_i \text{ and } f(x) = \nabla f(x)x.$$

And we have from [Lemma 2.1](#),  $\nabla f(x) \stackrel{d}{=} \tilde{W}_{\ell+1} \prod_{i=\ell}^1 D_i \tilde{W}_i$ , where

$$(D_i)_{j,k} = \begin{cases} 1 & \text{with probability } 1/2 \text{ if } j = k, \\ 0 & \text{otherwise,} \end{cases} \quad \{W_i\}_{i=1}^{\ell+1} \stackrel{d}{=} \{\tilde{W}_i\}_{i=1}^{\ell+1}.$$

Therefore, it suffices to analyze the random vector  $\tilde{W}_{\ell+1} \prod_{i=\ell}^1 D_i \tilde{W}_i$ . We first condition on a favorable event for the  $D_i$ . Note that we have by the union bound and an application of Hoeffding's inequality (e.g., [BLM13, Theorem 2.8]) that:

$$\forall i \in [\ell] : \text{Tr } D_i \geq \frac{d_i}{3} \text{ with probability at least } 1 - \delta/4,$$

since  $d_i \geq c \log(4\ell/\delta)$ . We now condition on the  $D_i$  and note that:

$$\left\| \widetilde{W}_{\ell+1} \prod_{i=\ell}^1 D_i \widetilde{W}_i \right\| \stackrel{d}{=} \left\| W_{\ell+1}^\dagger \prod_{i=\ell}^1 W_i^\dagger \right\| \text{ where}$$

$$\forall i \in \{2, \dots, \ell\} : W_i^\dagger \in \mathbb{R}^{\text{Tr } D_i \times \text{Tr } D_{i-1}}, W_{\ell+1}^\dagger \in \mathbb{R}^{1 \times \text{Tr } D_\ell}, W_1^\dagger \in \mathbb{R}^{\text{Tr } D_1 \times d}$$

$$(W_1^\dagger)_{j,:} \sim \mathcal{N}(0, I/d), W_{\ell+1}^\dagger \sim \mathcal{N}(0, I/d_\ell), \forall i \in \{2, \dots, \ell+1\} : (W_i^\dagger)_{j,:} \sim \mathcal{N}(0, I/d_{i-1}).$$

From the above display, we obtain from [Theorem C.3](#) and its corollary [C.4](#),

$$\begin{aligned} \left\| W_{\ell+1}^\dagger \prod_{i=\ell}^1 W_i^\dagger \right\| &\geq \frac{1}{2} \cdot \left\| W_{\ell+1}^\dagger \prod_{i=\ell}^2 W_i^\dagger \right\| && \text{with probability at least } 1 - \delta/(4\ell) \\ &\geq \frac{1}{2^{\ell+1}} && \text{with probability at least } 1 - \delta/4 \text{ by induction,} \end{aligned}$$

since  $\min_{0 \leq i \leq \ell} d_i \geq c \log(4\ell/\delta)$ . Through a similar argument, we obtain:

$$\left\| \widetilde{W}_{\ell+1} \prod_{i=\ell}^1 D_i \widetilde{W}_i \right\| \leq 2^{\ell+1} \tag{2}$$

with probability at least  $1 - \delta/4$ . We also have:

$$\left( \widetilde{W}_{\ell+1} \prod_{i=\ell}^1 D_i \widetilde{W}_i \left\| \widetilde{W}_{\ell+1} \prod_{i=\ell}^1 D_i \widetilde{W}_i \right\| = m \right) \stackrel{d}{=} \text{Unif}(m\mathbb{S}^{d-1}).$$

Combining this, [Eq. \(2\)](#) and [Lemma C.6](#), we get that:

$$|f(x)| \leq c2^\ell \sqrt{\log 1/\delta}$$

with probability at least  $1 - \delta/2$ , for some absolute constant  $c$ . A union bound over all the preceding events concludes the lemma.  $\square$

## 2.2 A Decomposition of Local Gradient Changes

The following decomposition allows us to reason about deviations layer by layer:

$$\begin{aligned} \nabla f(x) - \nabla f(y) &= W_{\ell+1} \left( \prod_{i=\ell}^1 D_i(\tilde{f}_i(x)) W_i - \prod_{i=\ell}^1 D_i(\tilde{f}_i(y)) W_i \right) \\ &= W_{\ell+1} \left( \prod_{i=\ell}^1 D_i(\tilde{f}_i(x)) W_i - (D_\ell(\tilde{f}_\ell(x)) + (D_\ell(\tilde{f}_\ell(y)) - D_\ell(\tilde{f}_\ell(x)))) W_\ell \prod_{i=\ell-1}^1 D_i(\tilde{f}_i(y)) W_i \right) \\ &= W_{\ell+1} \left( D_\ell(\tilde{f}_\ell(x)) W_\ell \left( \prod_{i=\ell-1}^1 D_i(\tilde{f}_i(x)) W_i - \prod_{i=\ell-1}^1 D_i(\tilde{f}_i(y)) W_i \right) \right. \\ &\quad \left. + (D_\ell(\tilde{f}_\ell(x)) - D_\ell(\tilde{f}_\ell(y))) W_\ell \prod_{i=\ell-1}^1 D_i(\tilde{f}_i(y)) W_i \right) \\ &= \underbrace{\sum_{j=1}^{\ell} W_{\ell+1} \left( \prod_{i=\ell}^{j+1} D_i(\tilde{f}_i(x)) W_i \right) \cdot (D_j(\tilde{f}_j(x)) - D_j(\tilde{f}_j(y))) W_j \cdot \left( \prod_{i=j-1}^1 D_i(\tilde{f}_i(y)) W_i \right)}_{\Delta_j}. \end{aligned} \tag{GD-DECOMP}$$

We use this decomposition to show that the gradient is locally constant. Concretely, consider a fixed term in the above decomposition and let  $i^* = \arg \min_{i < j} d_i$ . Letting  $M_{i^*} = \prod_{i=i^*}^1 D_i(\tilde{f}_i(y)) W_i$ , we can bound a single term,  $\Delta_j$ , as follows:

$$\|\Delta_j\| \leq \left\| W_{\ell+1} \left( \prod_{i=\ell}^{j+1} D_i(\tilde{f}_i(x)) W_i \right) (D_j(\tilde{f}_j(x)) - D_j(\tilde{f}_j(y))) W_j \left( \prod_{i=j-1}^{i^*+1} D_i(\tilde{f}_i(y)) W_i \right) \right\| \|M_{i^*}\|.$$

We bound the above by bounding both of the two factors on the right-hand-side. In the above expression, the length of  $\Delta_j$  is bounded by a product of the length of the corresponding factor in a truncated network starting at the output of layer  $i^*$  and a product of masked weight matrices up to layer  $i^*$  corresponding to the activation patterns of  $y$ . Intuitively, the first factor is expected to be small if the images of  $x$  and  $y$  at layer  $i^*$  are close and hence, we show that an image of a suitably small ball around  $x$  remains close to the image of  $x$  through all the layers of the network ([Lemma 2.3](#)). We then prove a spectral norm bound on  $M_{i^*}$  in [Lemma 2.6](#) and finally, establish a bound on  $\|\Delta_j\|$  in [Lemma 2.8](#) where we crucially rely on the scale preservation guarantees provided by [Lemma 2.3](#). Finally, combining these results with those of [Subsection 2.1](#) completes the proof of [Theorem 1.1](#).

### 2.3 Scale Preservation of Local Neighborhoods

In this section, we show that the image of a ball around  $x$  remains in a ball of suitable radius around the image of  $x$  projected through the various layers of the network. Here, we introduce additional notation used in the rest of the proof:

$$\forall j \in \{0\} \cup [\ell + 1] : f_{j,j}(x) = x, \forall i > j : \tilde{f}_{i,j}(x) = W_i f_{i-1,j}(x), f_{i,j}(x) = D_i(\tilde{f}_{i,j}(x)) \tilde{f}_{i,j}(x).$$

We now describe the decomposition of the neural network into segments, which are bounded by what we call bottleneck layers, and our analysis works separately with these segments. This decomposition is crucial for reducing the sizes of the  $\epsilon$ -nets that arise in our proofs. Intuitively, when we construct an  $\epsilon$ -net to prove that some property holds uniformly over a ball, it is crucial to work with the lowest-dimensional image of that ball, which appears in a bottleneck layer. These bottleneck layers are denoted by indices  $\{i_j\}_{j=1}^m$ , defined recursively from the output layer backwards with the convention that  $d_0 = d$ :

$$i_1 := \arg \min_{i \leq \ell} d_i, \quad \forall j > 1 \text{ s.t. } i_j \geq 1, \quad i_{j+1} := \arg \min_{j < i_j} d_j, \quad d_{\min}^i = \min_{j < i} d_j. \quad (\text{NN-DECOMP})$$

Note, that  $i_m = 0$ , for all  $j \in [m - 1]$ ,  $d_{i_j} < d_{i_{j+1}}$  and for all  $k \in \{i_{j+1}, \dots, i_j - 1\}$ ,  $d_k \geq d_{i_{j+1}}$ .

The following technical lemma bounds the scaling of the images at a layer in the network of an input  $x$  and of a ball around  $x$ . The crucial properties we will exploit are that these images avoid the origin, and that the radius of the image of the ball is not too large. The full proof is deferred to [Appendix A.2](#). In the proof sketch, we carry out the section of the proof where we transition between bottleneck layers in full as it is a simple illustration of how such ideas are used through the rest of the proof.

**Lemma 2.3.** *We have with probability at least  $1 - \delta$ ,*

$$\begin{aligned} \forall i \in [\ell] : \|f_i(x)\| &\geq \frac{1}{2^i} \cdot \sqrt{d_i}, \\ \forall \|x' - x\| \leq R : \|\tilde{f}_i(x) - \tilde{f}_i(x')\| &\leq (C \log d_{\max})^{i/2} \cdot \left(1 + \sqrt{\frac{d_i}{d_{\min}^i}}\right) \cdot R, \\ \forall \|x' - x\| \leq R : \|f_i(x) - f_i(x')\| &\leq (C \log d_{\max})^{i/2} \cdot \left(1 + \sqrt{\frac{d_i}{d_{\min}^i}}\right) \cdot R. \end{aligned}$$

*Proof Sketch.* The proof of the first claim is nearly identical to that of [Lemma 2.2](#).

For the second claim, we use a gridding argument with some subtleties. Concretely, we construct a new grid over the image of  $\mathbb{B}(x, R)$  whenever the number of units in a hidden layer drops below all the previous layers in the network starting from the input layer. These layers are precisely defined by the indices  $i_j$  in [NN-DECOMP](#). We now establish the following claim inductively where we adopt the convention  $i_0 = \ell + 1$ .

**Claim 2.4.** Suppose for  $j \geq 1$  and  $\tilde{R} \leq d_{\max} \cdot R$ :

$$\forall y \in \mathbb{B}(x, R) : \|f_{i_j}(x) - f_{i_j}(y)\| \leq \tilde{R}.$$

Then:

$$\forall i \in \{i_j + 1, \dots, i_{j-1} - 1\}, y \in \mathbb{B}(x, R) : \|\tilde{f}_i(x) - \tilde{f}_i(y)\| \leq (C\ell \log d_{\max})^{(i-i_j)/2} \cdot \sqrt{\frac{d_i}{d_{i_j}}} \cdot \tilde{R},$$

$$\forall y \in \mathbb{B}(x, R) : \|\tilde{f}_{i_{j-1}}(x) - \tilde{f}_{i_{j-1}}(y)\| \leq (C\ell \log d_{\max})^{(i_{j-1}-i_j)/2} \cdot \tilde{R}$$

with probability at least  $1 - \delta/8l$ .

*Proof Sketch.* We start by constructing an  $\varepsilon$ -net [Ver18, Definition 4.2.1],  $\mathcal{G}$ , of  $f_{i_j}(\mathbb{B}(x, R))$  with  $\varepsilon = 1/(d_{\max}^\ell)^{32}$ . Note that we may assume  $|\mathcal{G}| \leq (10\tilde{R}/\varepsilon)^{d_{i_j}}$ . We will prove the statement on the grid and extend to the rest of the space. For layer  $i + 1$ , defining  $\tilde{x} = f_{i_j}(x)$ , we have  $\forall \tilde{y} \in \mathcal{G}$ :

$$\begin{aligned} \|\tilde{f}_{i+1, i_j}(\tilde{x}) - \tilde{f}_{i+1, i_j}(\tilde{y})\| &= \|W_{i+1}(f_{i, i_j}(\tilde{x}) - f_{i, i_j}(\tilde{y}))\| \\ &\leq \|f_{i, i_j}(\tilde{x}) - f_{i, i_j}(\tilde{y})\| \cdot \sqrt{\frac{d_{i+1}}{d_i}} \cdot \left(1 + \sqrt{\frac{\log 1/\delta'}{d_{i+1}}}\right) \end{aligned}$$

with probability at least  $1 - \delta'$  as before by [Theorem C.3](#). By setting  $\delta' = \delta/(16|\mathcal{G}|^{\ell^2})$  and noting that  $d_i \geq d_{i_j}$ , the conclusion holds for layer  $i + 1 \leq i_j$  on  $\mathcal{G}$  with probability at least  $1 - \delta/(16\ell^2)$ . By induction and the union bound, we get:

$$\forall i \in \{i_j + 1, \dots, i_{j-1} - 1\}, \tilde{y} \in \mathcal{G} : \|\tilde{f}_{i, i_j}(\tilde{y}) - \tilde{f}_{i, i_j}(\tilde{x})\| \leq (C\ell \log d_{\max})^{(i-i_j)/2} \cdot \sqrt{\frac{d_i}{d_{i_j}}} \cdot \tilde{R}$$

$$\forall \tilde{y} \in \mathcal{G} : \|\tilde{f}_{i_{j-1}, i_j}(\tilde{y}) - \tilde{f}_{i_{j-1}, i_j}(\tilde{x})\| \leq (C\ell \log d_{\max})^{(i_{j-1}-i_j)/2} \cdot \tilde{R}$$

with probability at least  $1 - \delta/(16\ell^2)$ . To extend to all  $y \in f_{i_j}(\mathbb{B}(x, R))$ , we condition on the bound on  $\|W_i\|$  given by [Lemma C.1](#) for all  $i \leq i_{j-1}$  and note that  $\forall y \in f_{i_j}(\mathbb{B}(x, R))$ , for  $\tilde{y} = \arg \min_{z \in \mathcal{G}} \|z - y\|$ , and for  $i_j + 1 \leq i < i_{j-1}$ ,

$$\begin{aligned} \|\tilde{f}_{i, i_j}(\tilde{x}) - \tilde{f}_{i, i_j}(y)\| &\leq \|\tilde{f}_{i, i_j}(\tilde{x}) - \tilde{f}_{i, i_j}(\tilde{y})\| + \|\tilde{f}_{i, i_j}(y) - \tilde{f}_{i, i_j}(\tilde{y})\| \\ &\leq \|\tilde{f}_{i, i_j}(\tilde{x}) - \tilde{f}_{i, i_j}(\tilde{y})\| + \|y - \tilde{y}\| \prod_{k=i_j+1}^i \|W_k\| \\ &\leq \|\tilde{f}_{i, i_j}(\tilde{x}) - \tilde{f}_{i, i_j}(\tilde{y})\| + \varepsilon \prod_{k=i_j+1}^i \left(C \sqrt{\frac{d_k}{d_{k-1}}}\right) \\ &= \|\tilde{f}_{i, i_j}(\tilde{x}) - \tilde{f}_{i, i_j}(\tilde{y})\| + \varepsilon C^{i-i_j} \sqrt{\frac{d_i}{d_{i_j}}}, \end{aligned}$$

using that  $d_i \geq d_{i_j}$ . Similarly, for  $i = i_{j-1}$ , we have

$$\|\tilde{f}_{i_{j-1}, i_j}(\tilde{x}) - \tilde{f}_{i_{j-1}, i_j}(y)\| \leq \|\tilde{f}_{i_{j-1}, i_j}(\tilde{x}) - \tilde{f}_{i_{j-1}, i_j}(\tilde{y})\| + \varepsilon C^{i_{j-1}-i_j} \sqrt{\frac{d_{i_{j-1}-1}}{d_{i_j}}}.$$

Our setting of  $\varepsilon$  concludes the proof of the claim.  $\square$

An inductive application of [Claim 2.4](#), a union bound and the observation that:

$$\|f_i(x) - f_i(y)\| \leq \|\tilde{f}_i(x) - \tilde{f}_i(y)\|$$

concludes the proof of the lemma.  $\square$

The following lemma shows that few neurons have inputs of small magnitude for network input  $x$ .

**Lemma 2.5.** *With probability at least  $1 - \delta$ , for all  $i \in [\ell]$ , we have:*

$$\#\left\{j : |\langle (W_{i+1})_j, f_i(x) \rangle| \geq \alpha_i \frac{\|f_i(x)\|}{\sqrt{d_i}}\right\} \geq \left(1 - 2\sqrt{\frac{2}{\pi}} \alpha_i\right) d_{i+1}.$$

*Proof.* Follows from the fact that  $\Pr(|Z| \leq c) \leq c\sqrt{2/\pi}$  for  $Z \sim N(0, 1)$ , plus a simple application of Hoeffding's Inequality.  $\square$

## 2.4 Proving Local Gradient Smoothness

In this section, we show that the gradient is locally constant, and thus complete the proof of [Theorem 1.1](#). In our proof, we will bound each of the terms in the expansion of the gradient differences ([GD-DECOMP](#)). First, we prove a structural lemma on the spectral norm of the matrices appearing in the right-hand-side of ([GD-DECOMP](#)), allowing us to ignore the portion of the network till the last-encountered bottleneck layer. Define, for all  $i > j$ ,  $M_{i,j}(y) = \prod_{k=i}^{j+1} D_k(\tilde{f}_k(y))W_k$ .

**Lemma 2.6.** *With probability at least  $1 - \delta$  over the  $\{W_k\}$ , we have:*

$$\forall \|y-x\| \leq R, j \in [m-1] : \mathbb{P}_{\{D_1(\cdot), \dots, D_\ell(\cdot)\}} \left\{ \left\| M_{i_j, i_{j+1}}(y) \right\| \leq (C \cdot \ell \cdot \log d_{\max})^{(i_j - i_{j+1})/2} \right\} = 1,$$

where the probability is taken with respect to the random choices in the definition of the  $D_k(\cdot)$ .

We provide the first part of the proof in full as the simplest application of ideas that find further application in the subsequent result establishing bounds on terms in ([GD-DECOMP](#)); see [Appendix A.3](#).

*Proof Sketch.* To start, consider a fixed  $j \in [m]$  and condition on the conclusion of [Lemma 2.3](#) up to level  $i_{j+1}$ . Now, consider an  $\varepsilon$ -net of  $f_{i_{j+1}}(\mathbb{B}(x, R))$ ,  $\mathcal{G}$ , with resolution  $\varepsilon = 1/(d_{\max}^\ell)^{32}$ . As before,  $|\mathcal{G}| \leq (Cd_{\max})^{48\ell d_{i_{j+1}}}$  for some constant  $C$ . We additionally will consider subsets

$$\mathcal{S} = \left\{ (S_k)_{k=i_j-1}^{i_{j+1}+1} : S_k \subseteq [d_k], |S_k| \leq 4d_{i_{j+1}} \right\}.$$

Note that  $|\mathcal{G}| \cdot |\mathcal{S}|^2 \leq (d_{\max})^{64\ell d_{i_{j+1}}}$ . For  $y \in \mathcal{G}$ ,  $S^1, S^2 \in \mathcal{S}$ , consider the following matrix:

$$M_{y, S^1, S^2}^{i_j, i_{j+1}} = \prod_{k=i_j}^{i_{j+1}+1} (D_k(\tilde{f}_{k, i_{j+1}}(y)) + (D_{S_k^1} - D_{S_k^2}))W_k \text{ where } (D_S)_{i,j} = \begin{cases} 1, & \text{if } i = j \text{ and } i \in S, \\ 0, & \text{otherwise.} \end{cases}$$

We will bound the spectral norm of  $M_{y, S^1, S^2}^{i_j, i_{j+1}}$ . First, note that:

$$\left\| M_{y, S^1, S^2}^{i_j, i_{j+1}} \right\| \leq 2 \left\| \widetilde{M}_{y, S^1, S^2}^{i_j, i_{j+1}} \right\| \text{ where } \widetilde{M}_{y, S^1, S^2}^{i_j, i_{j+1}} := W_{i_j} \prod_{k=i_j-1}^{i_{j+1}+1} (D_k(\tilde{f}_{k, i_{j+1}}(y)) + (D_{S_k^1} - D_{S_k^2}))W_k$$

and observe that from [Lemma 2.1](#):

$$\widetilde{M}_{y, S^1, S^2}^{i_j, i_{j+1}} \stackrel{d}{=} W_{i_j} \prod_{k=i_j-1}^{i_{j+1}+1} (D_k + (D_{S_k^1} - D_{S_k^2}))W_k \text{ where } (D_k)_{i,j} = \begin{cases} 1 & \text{w.p } \frac{1}{2} \text{ if } i = j, \\ 0 & \text{otherwise.} \end{cases}$$

To bound the spectral norm, let  $\mathcal{B}$  be a  $1/3$ -net of  $\mathbb{S}^{d_{i_j}-1}$  and  $v \in \mathcal{B}$ . Applying [Theorem C.3](#),

$$\begin{aligned} \left\| v^\top \widetilde{M}_{y, S^1, S^2}^{i_j, i_{j+1}} \right\| &\leq \left\| v^\top W_{i_j} \prod_{k=i_j-1}^{i_{j+1}+2} (D_k + (D_{S_k^1} - D_{S_k^2}))W_k \right\| \cdot \left( 1 + \sqrt{\frac{\log 1/\delta'}{d_{i_{j+1}}}} \right) \\ &\leq \prod_{k=i_j}^{i_{j+1}} \left( 1 + \sqrt{\frac{\log 1/\delta'}{d_k}} \right) \end{aligned}$$

with probability at least  $\ell\delta'$ . But setting  $\delta' = \delta/(16\ell^4 \cdot |\mathcal{G}| \cdot |\mathcal{S}|^2)$  yields with probability at least  $1 - \delta/16$ :

$$\forall y \in \mathcal{G}, S_1, S_2 \in \mathcal{S} : \left\| \widetilde{M}_{y, S^1, S^2}^{i_j, i_{j+1}} \right\| \leq (C \cdot \ell \cdot \log d_{\max})^{(i_j - i_{j+1})/2}.$$

On the event in the conclusion of [Lemma 2.3](#), we have that  $f_i(y) \neq 0$  for all  $i \in [l], y \in \mathcal{G}$  and therefore, we have by a union bound over the discrete set  $\mathcal{G}$ :

$$\forall y \in \mathcal{G}, k \in \{i_{j+1}, \dots, i_j\}, m \in [d_k] : (\tilde{f}_{k, i_j}(y))_m \neq 0$$

proving the lemma for  $y \in \mathcal{G}$  as the activations are deterministic. For  $y \notin \mathcal{G}$ , the following claim concludes the proof of the lemma. The claim is essentially a generalization of [[BCGdC21](#), Eq. (18)] and its proof is deferred to the appendix.



**Claim 2.7.** With probability at least  $1 - \delta'/\ell^2$  over the  $\{W_k\}$ , we have for all  $m \in \{i_{j+1} + 1, \dots, i_j\}$  and  $y \in f_{i_{j+1}}(\mathbb{B}(x, R))$ :

$$\mathbb{P}_{\{D_k(\tilde{f}_{k,i_{j+1}}(y)), D_k(\tilde{f}_{k,i_{j+1}}(\tilde{y}))\}} \left\{ \text{Tr} |D_m(\tilde{f}_{m,i_{j+1}}(y)) - D_m(\tilde{f}_{m,i_{j+1}}(\tilde{y}))| \leq 4d_{i_{j+1}} \right\} = 1$$

where  $\tilde{y} = \arg \min_{z \in \mathcal{G}} \|z - y\|$ , and the probability is taken with respect to the random choices in the definition of the  $D_k(\cdot)$ .

The previous claim along with our previously established bounds establish the lemma.  $\square$

Our final technical result establishes the near-linearity of  $f$  in a ball around  $x$ . The full proof is deferred to [Appendix A.4](#) but in our proof sketch we identify sections which involve considerations unique to this lemma.

**Lemma 2.8.** For some absolute constant  $C$ , with probability at least  $1 - \delta$  over the  $\{W_k\}$ :

$$\forall \|x - y\| \leq R, j \in [\ell] : \mathbb{P}_{\{D_k(\tilde{f}_k(x)), D_k(\tilde{f}_k(y))\}} \left\{ \|\nabla f(x) - \nabla f(y)\| \leq \left( \frac{C^\ell}{\log^\ell d_{\max}} \right) \right\} = 1,$$

where the probability is taken with respect to the random choices in the definition of the  $D_k(\cdot)$ .

*Proof Sketch.* Consider a fixed term from [\(GD-DECOMP\)](#); that is, consider:

$$\text{Diff}_j(y) := W_{\ell+1} \left( \prod_{i=\ell}^{j+1} D_i(\tilde{f}_i(x)) W_i \right) \cdot (D_j(\tilde{f}_j(y)) - D_j(\tilde{f}_j(x))) W_j \cdot \left( \prod_{i=j-1}^1 D_i(\tilde{f}_i(y)) W_i \right).$$

We will show with high probability that  $\text{Diff}_j(y)$  is small for all  $y \in \mathbb{B}(x, R)$ . This will then imply the lemma by a union bound and [\(GD-DECOMP\)](#). Let  $k$  be such that  $i_k = \arg \min_{m < j} d_m$ . We will condition on the weights of the network up to layer  $i_k$ . Specifically, we will assume the conclusions of [Lemmas 2.3](#) and [2.6](#) up to layer  $i_k$ . We may now focus our attention solely on the segment of the network beyond layer  $i_k$  as a consequence of the following observation and [Lemma 2.6](#):

$$\|\text{Diff}_j(y)\| \leq \|\text{Diff}_{j,k}(x, y)\| \cdot \|M_{i_k,0}(y)\| \text{ where} \quad (3)$$

$$\text{Diff}_{j,k}(x, y) := W_{\ell+1} \left( \prod_{i=\ell}^{j+1} D_i(\tilde{f}_i(x)) W_i \right) (D_j(\tilde{f}_j(y)) - D_j(\tilde{f}_j(x))) W_j \left( \prod_{i=j-1}^{i_k+1} D_i(\tilde{f}_i(y)) W_i \right)$$

We will show for all  $y$  such that  $\|y - x\| \leq R$ :

$$\mathbb{P}_{\{D_m(\tilde{f}_m(x)), D_m(\tilde{f}_m(y))\}} \left\{ \|\text{Diff}_{j,k}(x, y)\| \geq \frac{C^\ell}{(\ell \log d_{\max})^{3\ell}} \right\} = 0$$

with probability at least  $1 - \delta/(16\ell^2)$ .

We have from [Lemma 2.1](#) that the random vector  $H$  defined below is spherically symmetric and satisfies  $\|H\| \leq 2^\ell$  with probability at least  $1 - \delta/(16\ell^4)$ :

$$W_{\ell+1} \prod_{i=\ell}^{j+1} D_i(\tilde{f}_i(x)) W_i \stackrel{d}{=} \tilde{W}_{\ell+1} \prod_{i=\ell}^{j+1} D_i \tilde{W}_i =: H.$$

As in the proof of [Lemma 2.6](#), let  $\mathcal{G}$  be an  $\varepsilon$ -net of  $f_{i_k}(\mathbb{B}(x, R))$  with  $\varepsilon$  as in the proof of [Lemma 2.6](#). We now break into two cases depending on how  $d_j$  compares to  $d_{i_k}$  and handle them separately. At this point, we have effectively reduced the multi-layer proof to the problem of analyzing deviations of activations at a fixed layer. For the remaining proof, we generalize approaches in [\[DS20\]](#) for the small width case and [\[BCGdC21\]](#) for the large width case.

**Case 1:**  $d_j \leq d_{i_k} (\ell \log d_{\max})^{20\ell}$ . In this case, the key observation already made in [\[DS20\]](#) is that under [Lemma 2.3](#), the number of neurons that may actually differ at layer  $j$  is at most  $d_j / \text{poly}(\ell \log d_{\max})^\ell$  between  $y \in \mathbb{B}(x, R)$  and  $x$ . Since,  $H$  is spherically distributed,  $\|H(D_j(\tilde{f}_j(x)) - D_j(\tilde{f}_j(y)))\|$  is very small and consequently the whole term is small.

**Case 2:**  $d_j \geq d_{i_k} (\ell \log d_{\max})^{20\ell}$ . This case is analogous to [BCGdC21]. The proof is technically involved and requires careful analysis of the random vector  $H(D_j(\tilde{f}_j(x)) - D_j(\tilde{f}_j(y)))W_j$ , which is complicated in our setting due to the matrices preceding it in (GD-DECOMP). It involves the distributional equivalence in Lemma 2.1.

A union bound and an application of the triangle inequality now imply the lemma.  $\square$

*Proof of Theorem 1.1.* On the intersection of the events in the conclusions of Lemmas 2.2, 2.3 and 2.8, we have that  $\nabla f(x)$  is deterministic and furthermore, we have:

$$\|\nabla f(x)\| \geq \frac{1}{2^{\ell+1}}, \quad |f(x)| \leq c2^\ell \sqrt{\log 1/\delta}$$

$$\forall y \text{ s.t } \|y - x\| \leq R : \|\nabla f(x) - \nabla f(y)\| \leq \frac{1}{(\ell \log d_{\max})^\ell} = o(1).$$

Assume  $f(x) > 0$  (the alternative is similar) and let  $\eta = -\frac{2^\ell \log d \cdot \sqrt{\log 1/\delta}}{\|\nabla f(x)\|^2}$ , we have for the point  $x + \eta \nabla f(x)$ , defining the function  $g(t) = f(x + t\eta \nabla f(x))$ :

$$f(x + \eta \nabla f(x)) = f(x) + \int_0^1 g'(t) dt = f(x) + \int_0^1 (\eta \nabla f(x))^\top \nabla f(x + t\eta \nabla f(x)) dt$$

$$= f(x) - (1 - o(1))2^\ell \log d \sqrt{\log 1/\delta} \leq -f(x).$$

Our lower bounds on  $\nabla f(x)$  ensure  $\|\eta \nabla f(x)\| \leq R$ , concluding the proof of the theorem.  $\square$

### 3 The Impact of Depth

Theorem 1.1 relies on the depth being constant. In this section, we show that some constraint on the depth is necessary in order to ensure the existence of adversarial examples. In particular, the following result gives an example of a sufficiently deep network for which, with high probability, the output of the network will have the same sign for all input patterns.

**Theorem 3.1.** *Fix a sufficiently large  $d \in \mathbb{N}$ , an  $\ell \geq d^3$  and  $(\ell d)^{20} \leq k \leq \exp(\sqrt{\ell})$ , and consider the randomly initialized neural network (NN-DEF) with  $d_1 = \dots = d_\ell = k$ . There is a universal constant  $C$  such that with probability at least 0.9,*

$$\forall x, y \in \mathbb{S}^{d-1} : \frac{|f(x) - f(y)|}{|f(x)|} \leq C \sqrt{\frac{\log d}{d}}.$$

The proof involves showing that the image of  $\mathbb{S}^{d-1}$  is bounded away from 0, and that the inner products between images of two input vectors throughout the network converge. The following lemma shows how the expected inner products evolve through the network. The lemma follows from computing a double integral; see [CS09, Eq. (6)]. The proof of Theorem 3.1 is in Appendix B.

**Lemma 3.2.** *Let  $d \in \mathbb{N}$ . Fix  $x, y \in \mathbb{R}^d$  with  $x, y \neq 0$ . Then for  $g \sim \mathcal{N}(0, I)$ :*

$$\frac{\mathbb{E} [\max \{g^\top x, 0\} \max \{g^\top y, 0\}]}{\sqrt{\mathbb{E}(\max \{g^\top x, 0\})^2} \cdot \sqrt{\mathbb{E}(\max \{g^\top y, 0\})^2}} = \frac{\sin \theta}{\pi} + \left(1 - \frac{\theta}{\pi}\right) \cos \theta$$

where  $\theta = \arccos(x^\top y / (\|x\| \|y\|))$ .

### 4 Disclosure of Funding and Competing Interests

PB and YC gratefully acknowledge the support of the NSF through grants DMS-2023505 and DMS-2031883, the Simons Foundation through award #814639, and Microsoft through the BAIR Open Research Commons. We declare no competing interests.

## References

- [BCGdC21] Sébastien Bubeck, Yeshwanth Cherapanamjeri, Gauthier Gidel, and Rémi Tachet des Combes. A single gradient step finds adversarial examples on random two-layers neural networks. *arXiv*, 2104.03863, 2021. [1](#), [8](#), [9](#), [10](#)
- [BLM13] Stéphane Boucheron, Gábor Lugosi, and Pascal Massart. *Concentration inequalities*. Oxford University Press, Oxford, 2013. [4](#), [27](#)
- [BLN21] Sébastien Bubeck, Yuanzhi Li, and Dheeraj M. Nagaraj. A law of robustness for two-layers neural networks. In Mikhail Belkin and Samory Kpotufe, editors, *Conference on Learning Theory, COLT 2021, 15-19 August 2021, Boulder, Colorado, USA*, volume 134 of *Proceedings of Machine Learning Research*, pages 804–820. PMLR, 2021. [3](#)
- [BLPR19] Sebastien Bubeck, Yin Tat Lee, Eric Price, and Ilya Razenshteyn. Adversarial examples from computational constraints. In Kamalika Chaudhuri and Ruslan Salakhutdinov, editors, *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 831–840. PMLR, 2019. [1](#)
- [BS21] Sébastien Bubeck and Mark Sellke. A universal law of robustness via isoperimetry. *CoRR*, abs/2105.12806, 2021. [3](#)
- [CS09] Youngmin Cho and Lawrence Saul. Kernel methods for deep learning. In Y. Bengio, D. Schuurmans, J. Lafferty, C. Williams, and A. Culotta, editors, *Advances in Neural Information Processing Systems*, volume 22. Curran Associates, Inc., 2009. [10](#)
- [CW17a] Nicholas Carlini and David Wagner. Towards evaluating the robustness of neural networks. In *2017 IEEE Symposium on Security and Privacy (SP)*, pages 39–57, 2017. [1](#)
- [CW17b] Nicholas Carlini and David A. Wagner. Adversarial examples are not easily detected: Bypassing ten detection methods. In Bhavani M. Thuraisingham, Battista Biggio, David Mandell Freeman, Brad Miller, and Arunesh Sinha, editors, *Proceedings of the 10th ACM Workshop on Artificial Intelligence and Security, AISec@CCS 2017, Dallas, TX, USA, November 3, 2017*, pages 3–14. ACM, 2017. [1](#)
- [DS20] Amit Daniely and Hadas Shacham. Most ReLU networks suffer from  $\ell^2$  adversarial perturbations. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 6629–6636. Curran Associates, Inc., 2020. [1](#), [9](#)
- [FCSG17] Reuben Feinman, Ryan R. Curtin, Saurabh Shintre, and Andrew B. Gardner. Detecting adversarial samples from artifacts. *arXiv*, 1703.00410, 2017. [1](#)
- [GSS15] Ian J. Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. In Yoshua Bengio and Yann LeCun, editors, *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015. [1](#)
- [MMS<sup>+</sup>18] Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards deep learning models resistant to adversarial attacks. In *International Conference on Learning Representations*, 2018. [1](#), [3](#)
- [QMG<sup>+</sup>19] Chongli Qin, James Martens, Sven Gowal, Dilip Krishnan, Krishnamurthy Dvijotham, Alhussein Fawzi, Soham De, Robert Stanforth, and Pushmeet Kohli. Adversarial robustness through local linearization. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019. [1](#), [3](#)

- [SSRD19] Adi Shamir, Itay Safran, Eyal Ronen, and Orr Dunkelman. A simple explanation for the existence of adversarial examples with small hamming distance. *arXiv*, 1901.10861, 2019. [1](#)
- [SZS<sup>+</sup>14] Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus. Intriguing properties of neural networks. In *International Conference on Learning Representations*, 2014. [1](#)
- [Ver18] Roman Vershynin. *High-dimensional probability*, volume 47 of *Cambridge Series in Statistical and Probabilistic Mathematics*. Cambridge University Press, Cambridge, 2018. [7](#), [14](#), [22](#)

## A Deferred Proofs from Section 2

### A.1 Proof of Lemma 2.1

For the first claim, we introduce the following diagonal random signed matrices,  $S_i$  for  $i \in [m]$  with:

$$(S_i)_{j,j} = \pm 1 \text{ with equal probability, } (S_i)_{j,k} = 0 \text{ if } j \neq k.$$

We prove the claim by induction on  $m$ . When  $m = 1$ , the claim is trivially true. For  $m > 1$ , we have:

$$\begin{aligned} & W_m \prod_{j=m-1}^1 (D_j(\tilde{h}_j(x)) + B_j)W_j \\ & \stackrel{d}{=} W_m \left( \prod_{j=m-1}^2 (D_j(\tilde{h}_j(x)) + B_j)W_j \right) (D_1(S_1\tilde{h}_1(x)) + B_1)S_1W_1 \quad (S_1W_1 \stackrel{d}{=} W_1) \\ & \stackrel{d}{=} W_m \left( \prod_{j=m-1}^2 (D_j + B_j)W_j \right) (D_1(S_1\tilde{h}_1(x)) + B_1)S_1W_1 \quad (\text{hypothesis for } m-1) \\ & \stackrel{d}{=} W_m \left( \prod_{j=m-1}^2 (D_j + B_j)W_j \right) S_1(D_1(S_1\tilde{h}_1(x)) + B_1)S_1W_1 \quad (W_2S_1 \stackrel{d}{=} W_2) \\ & = W_m \left( \prod_{j=m-1}^2 (D_j + B_j)W_j \right) (D_1(S_1\tilde{h}_1(x)) + B_1)W_1 \quad (S_1B_1S_1 = B_1) \\ & \stackrel{d}{=} W_m \left( \prod_{j=m-1}^2 (D_j + B_j)W_j \right) (D_1 + B_1)W_1 \\ & \stackrel{d}{=} W_m \prod_{j=m-1}^1 (D_j + B_j)W_j. \end{aligned}$$

Similarly, for the second claim, we have:

$$\begin{aligned} & \left\| \prod_{j=m}^1 (D_j(\tilde{h}_j(x)) + B_j)W_j \right\| \\ & \stackrel{d}{=} \left\| (D_m(S_m\tilde{h}_m(x)) + B_m)S_mW_m \prod_{j=m-1}^1 (D_j(\tilde{h}_j(x)) + B_j)W_j \right\| \\ & \stackrel{d}{=} \left\| S_m(D_m(S_m\tilde{h}_m(x)) + B_m)S_mW_m \prod_{j=m-1}^1 (D_j(\tilde{h}_j(x)) + B_j)W_j \right\| \\ & \stackrel{d}{=} \left\| (D_m(S_m\tilde{h}_m(x)) + B_m)W_m \prod_{j=m-1}^1 (D_j(\tilde{h}_j(x)) + B_j)W_j \right\| \\ & \stackrel{d}{=} \left\| (D_m + B_m)W_m \prod_{j=m-1}^1 (D_j(\tilde{h}_j(x)) + B_j)W_j \right\| \\ & \stackrel{d}{=} \left\| \prod_{j=m}^1 (D_j + B_j)W_j \right\| \end{aligned}$$

□

## A.2 Proof of Lemma 2.3

We start with the first claim of the lemma. As in Lemma 2.2, we note from Lemma 2.1:

$$f_i(x) = \left( \prod_{j=i}^1 D_j(\tilde{f}_j(x)) W_j \right) x$$

$$\left\| \prod_{j=i}^1 D_j(\tilde{f}_j(x)) W_j \right\| \stackrel{d}{=} \left\| \prod_{j=i}^1 D_j \tilde{W}_j \right\| \text{ where}$$

$$(D_i)_{j,k} = \begin{cases} 1 & \text{w.p } 1/2 \text{ if } j = k \\ 0 & \text{otherwise} \end{cases} \text{ and } \{W_i\}_{i=1}^{\ell+1} \stackrel{d}{=} \{\tilde{W}_i\}_{i=1}^{\ell+1}.$$

Therefore, it suffices to analyze the distribution of  $\prod_{j=i}^1 D_j \tilde{W}_j x$ . Again, we condition on the following event:

$$\forall i \in [\ell] : \text{Tr } D_i \geq \frac{d_i}{3}$$

which occurs with probability at least  $1 - \delta/8$ . As in the proof of Lemma 2.2, we observe the following distributional equivalence, when we condition on the  $D_j$ :

$$\forall i \in [\ell] : \left\| \prod_{j=i}^1 D_j \tilde{W}_j x \right\| \stackrel{d}{=} \left\| \prod_{j=i}^1 W_j^\dagger x \right\| \text{ where}$$

$$\forall j \in \{2, \dots, i\}, W_j^\dagger \in \mathbb{R}^{\text{Tr } D_j \times \text{Tr } D_{j-1}}, W_1^\dagger \in \mathbb{R}^{\text{Tr } D_1 \times d}$$

$$\forall j \in \{2, \dots, i\}, (W_j^\dagger)_{k,:} \sim \mathcal{N}(0, I/d_{j-1}), (W_1^\dagger)_{k,:} \sim \mathcal{N}(0, I/d).$$

A recursive application of Theorem C.3 as in the proof of Lemma 2.2 yields:

$$\left\| \prod_{j=i}^1 D_j \tilde{W}_j x \right\| \geq \|x\| \prod_{j=i}^1 \frac{1}{2} \sqrt{\frac{d_j}{d_{j-1}}}$$

with probability at least  $1 - \delta/16\ell^2$ . A union bound over the  $j$  layers concludes the proof of the first claim of the lemma.

We will establish the second claim through a gridding based argument with some subtleties. Namely, we construct a new grid over the image of  $\mathbb{B}(x, R)$  whenever the number of units in a hidden layer drops lower than all the previous layers in the network starting from the input layer. These layers are precisely defined by the indices  $i_j$  in NN-DECOMP. We now establish the following claim inductively where we adopt the convention  $i_0 = \ell + 1$ .

**Claim A.1.** Suppose for  $j \geq 1$  and  $\tilde{R} \leq d_{\max} \cdot R$ :

$$\forall y \in \mathbb{B}(x, R) : \|f_{i_j}(x) - f_{i_j}(y)\| \leq \tilde{R}.$$

Then:

$$\forall i \in \{i_j + 1, \dots, i_{j-1} - 1\}, y \in \mathbb{B}(x, R) : \|\tilde{f}_i(x) - \tilde{f}_i(y)\| \leq (C\ell \log d_{\max})^{(i-i_j)/2} \cdot \sqrt{\frac{d_i}{d_{i_j}}} \cdot \tilde{R},$$

$$\forall y \in \mathbb{B}(x, R) : \|\tilde{f}_{i_{j-1}}(x) - \tilde{f}_{i_{j-1}}(y)\| \leq (C\ell \log d_{\max})^{(i_{j-1}-i_j)/2} \cdot \tilde{R}$$

with probability at least  $1 - \delta/8l$ .

*Proof.* We start by constructing an  $\varepsilon$ -net [Ver18, Definition 4.2.1],  $\mathcal{G}$ , of  $f_{i_j}(\mathbb{B}(x, R))$  with  $\varepsilon = 1/(d_{\max}^\ell)^{32}$ . Note that we may assume  $|\mathcal{G}| \leq (10\tilde{R}/\varepsilon)^{d_{i_j}}$ . We will prove the statement on the grid and extend to the rest of the space. For layer  $i + 1$ , defining  $\tilde{x} = f_{i_j}(x)$ , we have  $\forall \tilde{y} \in \mathcal{G}$ :

$$\|\tilde{f}_{i+1, i_j}(\tilde{x}) - \tilde{f}_{i+1, i_j}(\tilde{y})\| = \|W_{i+1}(f_{i, i_j}(\tilde{x}) - f_{i, i_j}(\tilde{y}))\|$$

$$\leq \|f_{i,i_j}(\tilde{x}) - f_{i,i_j}(\tilde{y})\| \cdot \sqrt{\frac{d_{i+1}}{d_i}} \cdot \left(1 + \sqrt{\frac{\log 1/\delta'}{d_{i+1}}}\right)$$

with probability at least  $1 - \delta'$  as before by [Theorem C.3](#). By setting  $\delta' = \delta/(16|\mathcal{G}|\ell^2)$  and noting that  $d_i \geq d_{i_j}$ , the conclusion holds for layer  $i + 1 \leq i_j$  on  $\mathcal{G}$  with probability at least  $1 - \delta/(16\ell^2)$ . By induction and the union bound, we get:

$$\begin{aligned} \forall i \in \{i_j + 1, \dots, i_{j-1} - 1\}, \tilde{y} \in \mathcal{G} : \| \tilde{f}_{i,i_j}(\tilde{y}) - \tilde{f}_{i,i_j}(\tilde{x}) \| &\leq (C\ell \log d_{\max})^{(i-i_j)/2} \cdot \sqrt{\frac{d_i}{d_{i_j}}} \cdot \tilde{R} \\ \forall \tilde{y} \in \mathcal{G} : \| \tilde{f}_{i_{j-1},i_j}(\tilde{y}) - \tilde{f}_{i_{j-1},i_j}(\tilde{x}) \| &\leq (C\ell \log d_{\max})^{(i_{j-1}-i_j)/2} \cdot \tilde{R} \end{aligned}$$

with probability at least  $1 - \delta/(16\ell^2)$ . To extend to all  $y \in f_{i_j}(\mathbb{B}(x, R))$ , we condition on the bound on  $\|W_i\|$  given by [Lemma C.1](#) for all  $i \leq i_{j-1}$  and note that  $\forall y \in f_{i_j}(\mathbb{B}(x, R))$ , for  $\tilde{y} = \arg \min_{z \in \mathcal{G}} \|z - y\|$ , and for  $i_j + 1 \leq i < i_{j-1}$ ,

$$\begin{aligned} \| \tilde{f}_{i,i_j}(\tilde{x}) - \tilde{f}_{i,i_j}(y) \| &\leq \| \tilde{f}_{i,i_j}(\tilde{x}) - \tilde{f}_{i,i_j}(\tilde{y}) \| + \| \tilde{f}_{i,i_j}(y) - \tilde{f}_{i,i_j}(\tilde{y}) \| \\ &\leq \| \tilde{f}_{i,i_j}(\tilde{x}) - \tilde{f}_{i,i_j}(\tilde{y}) \| + \|y - \tilde{y}\| \prod_{k=i_j+1}^i \|W_k\| \\ &\leq \| \tilde{f}_{i,i_j}(\tilde{x}) - \tilde{f}_{i,i_j}(\tilde{y}) \| + \varepsilon \prod_{k=i_j+1}^i \left( C \sqrt{\frac{d_k}{d_{k-1}}} \right) \\ &= \| \tilde{f}_{i,i_j}(\tilde{x}) - \tilde{f}_{i,i_j}(\tilde{y}) \| + \varepsilon C^{i-i_j} \sqrt{\frac{d_i}{d_{i_j}}}, \end{aligned}$$

using that  $d_i \geq d_{i_j}$ . Similarly, for  $i = i_{j-1}$ , we have

$$\| \tilde{f}_{i_{j-1},i_j}(\tilde{x}) - \tilde{f}_{i_{j-1},i_j}(y) \| \leq \| \tilde{f}_{i_{j-1},i_j}(\tilde{x}) - \tilde{f}_{i_{j-1},i_j}(\tilde{y}) \| + \varepsilon C^{i_{j-1}-i_j} \sqrt{\frac{d_{i_{j-1}-1}}{d_{i_j}}}.$$

Our setting of  $\varepsilon$  concludes the proof of the claim.  $\square$

An inductive application of [Claim A.1](#), a union bound and the observation that:

$$\|f_i(x) - f_i(y)\| \leq \| \tilde{f}_i(x) - \tilde{f}_i(y) \|$$

concludes the proof of the lemma.  $\square$

### A.3 Proof of [Lemma 2.6](#)

To start, consider a fixed  $j \in [m]$  and condition on the conclusion of [Lemma 2.3](#) up to level  $i_{j+1}$ . Now, consider an  $\varepsilon$ -net of  $f_{i_{j+1}}(\mathbb{B}(x, R))$ ,  $\mathcal{G}$ , with resolution  $\varepsilon = 1/(d_{\max}^\ell)^{32}$ . As before,  $|\mathcal{G}| \leq (Cd_{\max})^{48\ell d_{i_{j+1}}}$  for some constant  $C$ . We additionally will consider subsets

$$\mathcal{S} = \left\{ (S_k)_{k=i_j-1}^{i_{j+1}+1} : S_k \subseteq [d_k], |S_k| \leq 4d_{i_{j+1}} \right\}.$$

Note that  $|\mathcal{G}| \cdot |\mathcal{S}|^2 \leq (Cd_{\max})^{64\ell d_{i_{j+1}}}$ . For  $y \in \mathcal{G}, S^1, S^2 \in \mathcal{S}$ , consider the following matrix:

$$\begin{aligned} M_{y, S^1, S^2}^{i_j, i_{j+1}} &= \prod_{k=i_j}^{i_{j+1}+1} (D_k(\tilde{f}_{k, i_{j+1}}(y))) + (D_{S_k^1} - D_{S_k^2})W_k \text{ where} \\ (D_S)_{i,j} &= \begin{cases} 1, & \text{if } i = j \text{ and } i \in S \\ 0, & \text{otherwise} \end{cases}. \end{aligned}$$

We will bound the spectral norm of  $M_{y,S^1,S^2}^{i_j,i_{j+1}}$ . First, note that:

$$\left\| M_{y,S^1,S^2}^{i_j,i_{j+1}} \right\| \leq 2 \left\| \widetilde{M}_{y,S^1,S^2}^{i_j,i_{j+1}} \right\| \text{ where } \widetilde{M}_{y,S^1,S^2}^{i_j,i_{j+1}} := W_{i_j} \prod_{k=i_j-1}^{i_{j+1}+1} (D_k(\widetilde{f}_{k,i_{j+1}}(y)) + (D_{S_k^1} - D_{S_k^2}))W_k$$

and observe that from [Lemma 2.1](#):

$$\begin{aligned} \widetilde{M}_{y,S^1,S^2}^{i_j,i_{j+1}} &\stackrel{d}{=} W_{i_j} \prod_{k=i_j-1}^{i_{j+1}+1} (D_k + (D_{S_k^1} - D_{S_k^2}))W_k \text{ where} \\ (D_k)_{i,j} &= \begin{cases} 1 & \text{with probability } \frac{1}{2} \text{ if } i = j, \\ 0 & \text{otherwise.} \end{cases} \end{aligned}$$

To bound the spectral norm, let  $\mathcal{B}$  be a  $1/3$ -net of  $\mathbb{S}^{d_{i_j}-1}$  and  $v \in \mathcal{B}$ . Applying [Theorem C.3](#),

$$\begin{aligned} \left\| v^\top \widetilde{M}_{y,S^1,S^2}^{i_j,i_{j+1}} \right\| &\leq \left\| v^\top W_{i_j} \prod_{k=i_j-1}^{i_{j+1}+2} (D_k + (D_{S_k^1} - D_{S_k^2}))W_k \right\| \cdot \left( 1 + \sqrt{\frac{\log 1/\delta'}{d_{i_{j+1}}}} \right) \\ &\leq \prod_{k=i_j}^{i_{j+1}} \left( 1 + \sqrt{\frac{\log 1/\delta'}{d_k}} \right) \end{aligned}$$

with probability at least  $\ell\delta'$ . But setting  $\delta' = \delta/(16\ell^4 \cdot |\mathcal{G}| \cdot |\mathcal{S}|^2)$  yields

$$\begin{aligned} \frac{\log 1/\delta'}{d_k} &\leq \frac{C\ell d_{i_{j+1}} \log d_{\max}}{d_k} + \frac{\log 1/\delta}{d_k} \\ &\leq C'\ell \log d_{\max} \end{aligned}$$

because  $d_k \geq d_{i_{j+1}}$  and by the assumption on  $\delta$ . Hence, with probability at least  $1 - \delta/16$ :

$$\forall y \in \mathcal{G}, S_1, S_2 \in \mathcal{S} : \left\| \widetilde{M}_{y,S^1,S^2}^{i_j,i_{j+1}} \right\| \leq (C \cdot \ell \cdot \log d_{\max})^{(i_j - i_{j+1})/2}.$$

On the event in the conclusion of [Lemma 2.3](#), we have that for all  $y \in \mathbb{B}(x, R)$  and  $i \in [\ell]$ ,  $f_i(y) \neq 0$ , and therefore we have by a union bound over the discrete set  $\mathcal{G}$  in an event of probability 1,

$$\forall y \in \mathcal{G}, k \in \{i_{j+1}, \dots, i_j\}, m \in [d_k] : (\widetilde{f}_{k,i_j}(y))_m \neq 0.$$

We now show a basic structural claim of how activation patterns differ across the various layers between a point and its closest approximation in  $\mathcal{G}$ :

**Claim A.2.** With probability at least  $1 - \delta'/\ell^2$  over the  $W_k$ , we have for all  $m \in \{i_{j+1} + 1, \dots, i_j\}$  and  $y \in f_{i_{j+1}}(\mathbb{B}(x, R))$ :

$$\mathbb{P}\{D_n(\widetilde{f}_{n,i_{j+1}}(y)), D_n(\widetilde{f}_{n,i_{j+1}}(\tilde{y}))\} \left\{ \text{Tr} |D_m(\widetilde{f}_{m,i_{j+1}}(y)) - D_m(\widetilde{f}_{m,i_{j+1}}(\tilde{y}))| \leq 4d_{i_{j+1}} \right\} = 1,$$

where  $\tilde{y} = \arg \min_{z \in \mathcal{G}} \|z - y\|$ .

*Proof.* We start by conditioning on the conclusion of [Lemma C.1](#) (with  $A = \sqrt{d_{k-1}}W_k$ ) up to layer  $i_{j+1}$ . Then:

$$\forall m \in \{i_{j+1} + 1, \dots, i_j\}, y \in f_{i_{j+1}}(\mathbb{B}(x, R)) : \|\widetilde{f}_{m,i_{j+1}}(y) - \widetilde{f}_{m,i_{j+1}}(\tilde{y})\| \leq \varepsilon \cdot (Cd_{\max})^{(m - i_{j+1})/2}.$$

Now, let  $m \in \{i_{j+1} + 1, \dots, i_j\}$  and  $y \in \mathcal{G}$ . We now operate on the conclusion of [Lemma 2.3](#) up to layer  $m - 1$  and hence,  $(\widetilde{f}_{m,i_{j+1}}(y))_n \neq 0$  for all  $n$ . Now, we have:

$$\begin{aligned} &\max_{z \text{ s.t. } \|z - y\| \leq \varepsilon} \#\{i : \text{sign}(\langle (W_m)_{i,:}, f_{m-1,i_{j+1}}(y) \rangle) \neq \text{sign}(\langle (W_m)_{i,:}, f_{m-1,i_{j+1}}(z) \rangle)\} \\ &\leq \#\{i : \exists z \text{ s.t. } \|z - y\| \leq \varepsilon \text{ and } \text{sign}(\langle (W_m)_{i,:}, f_{m-1,i_{j+1}}(y) \rangle) \neq \text{sign}(\langle (W_m)_{i,:}, f_{m-1,i_{j+1}}(z) \rangle)\}. \end{aligned}$$



Defining  $w_i := (W_m)_{i,:}$ , and

$$W_i(y) := \mathbf{1}\{\exists z \text{ s.t } \|z-y\| \leq \varepsilon \text{ and } \text{sign}(\langle (W_m)_{i,:}, f_{m-1,i_{j+1}}(y) \rangle) \neq \text{sign}(\langle (W_m)_{i,:}, f_{m-1,i_{j+1}}(z) \rangle)\},$$

we get:

$$\begin{aligned} \mathbb{P}\{W_i(y) = 1\} &\leq \mathbb{P}\left\{(Cd_{\max})^{(m-1-i_{j+1})/2} \cdot \varepsilon \cdot \|w_i\| \geq |\langle w_i, y \rangle|\right\} \\ &\leq \mathbb{P}\{\|w_i\| \geq 2\} + \mathbb{P}\left\{2\varepsilon(Cd_{\max})^{(m-1-i_{j+1})/2} \geq |\langle w_i, y \rangle|\right\} \leq \varepsilon \cdot (Cd_{\max})^{(m-i_{j+1})/2}. \end{aligned}$$

Therefore, we have:

$$\mathbb{P}\left\{\sum_{i=1}^{d_m} W_i(y) \geq 4d_{i_{j+1}}\right\} \leq d_m^{4d_{i_{j+1}}} \cdot \left(\varepsilon(Cd_{\max})^{(m-i_{j+1})/2}\right)^{4d_{i_{j+1}}}$$

and hence by the union bound and our setting of  $\varepsilon$  and bound on  $\mathcal{G}$ :

$$\mathbb{P}\left\{\exists y \in \mathcal{G} : \sum_{i=1}^{d_m} W_i(y) \geq 4d_{i_{j+1}}\right\} \leq (Cd_{\max})^{60\ell d_{i_{j+1}}} \cdot \varepsilon^{4d_{i_{j+1}}} \leq \frac{\delta'}{16\ell^4}.$$

□

The claim concludes the proof of the lemma. □

#### A.4 Proof of Lemma 2.8

Consider a fixed term in the decomposition of the gradient difference from **GD-DECOMP**; that is, consider the random vector-valued function:

$$\text{Diff}_j(y) := W_{\ell+1} \left( \prod_{i=\ell}^{j+1} D_i(\tilde{f}_i(x)) W_i \right) \cdot (D_j(\tilde{f}_j(y)) - D_j(\tilde{f}_j(x))) W_j \cdot \left( \prod_{i=j-1}^1 D_i(\tilde{f}_i(y)) W_i \right).$$

We will show with high probability that  $\text{Diff}_j(y) = o(1)$  with high probability for all  $\|x - y\| \leq R$ . This will then imply the lemma by a union bound and **GD-DECOMP**. Let  $k$  be such that  $i_k = \arg \min_{m < j} d_m$ . We will condition on the weights of the network up to layer  $i_k$ . Specifically, we will assume the conclusions of **Lemmas 2.3** and **2.6** up to layer  $i_k$ . We may now focus our attention solely on the segment of the network beyond layer  $i_k$  as a consequence of the following observations:

$$\|\text{Diff}_j(y)\| \leq \|\text{Diff}_{j,k}(x, y)\| \cdot \|M_{i_k,0}(y)\| \text{ where} \quad (4)$$

$$\text{Diff}_{j,k}(x, y) := W_{\ell+1} \left( \prod_{i=\ell}^{j+1} D_i(\tilde{f}_i(x)) W_i \right) (D_j(\tilde{f}_j(y)) - D_j(\tilde{f}_j(x))) W_j \left( \prod_{i=j-1}^{i_k+1} D_i(\tilde{f}_i(y)) W_i \right)$$

We will show for all  $y$  such that  $\|y - x\| \leq R$ :

$$\mathbb{P}\{D_m(\tilde{f}_m(x)), D_m(\tilde{f}_m(y))\} \left\{ \|\text{Diff}_{j,k}(x, y)\| \geq \frac{C^\ell}{(\ell \log d_{\max})^{3\ell}} \right\} = 0 \quad (5)$$

with probability at least  $1 - \delta/(16\ell^2)$ .

Observe now that from **Lemma 2.1**:

$$\begin{aligned} W_{\ell+1} \prod_{i=\ell}^{j+1} D_i(\tilde{f}_i(x)) W_i &\stackrel{d}{=} \widetilde{W}_{\ell+1} \prod_{i=\ell}^{j+1} D_i \widetilde{W}_i =: H \text{ where} \\ (D_i)_{m,n} &= \begin{cases} 1, & \text{with probability } 1/2 \text{ if } m = n \\ 0, & \text{otherwise} \end{cases}, \{W_i\}_{i=j+1}^{\ell+1} \stackrel{d}{=} \{\widetilde{W}_i\}_{i=j+1}^{\ell+1} \end{aligned}$$

Therefore,  $H$  is a spherically-symmetric random vector and we condition on the following high probability bound on its length:

$$\|H\| \leq \left\| \widetilde{W}_{\ell+1} \prod_{i=\ell}^{j+2} D_i \widetilde{W}_i \right\| \cdot \left( 1 + 2\sqrt{\frac{\log \ell / \delta'}{d_j}} \right) \leq (2)^{\ell+1-j} \quad (6)$$

with probability at least  $1 - \delta'$  and setting  $\delta' = \delta / (64\ell^2)$ . Observe that the distribution of  $H$  remains spherically symmetric even after conditioning on the above event. As in the proof of [Lemma 2.6](#), let  $\mathcal{G}$  be an  $\varepsilon$ -net of  $f_{i_k}(\mathbb{B}(x, R))$  with  $\varepsilon$  as in the proof of [Lemma 2.6](#).

**Claim A.3.** We have for all  $m \in \{i_k + 1, \dots, i_{k-1}\}$ ,  $y \in f_{i_k}(\mathbb{B}(x, R))$ :

$$\mathbb{P}_{\{D_m(\widetilde{f}_{m,i_k}(y)), D_m(\widetilde{f}_{m,i_k}(\widetilde{y}))\}} \left\{ \text{Tr} |D_m(\widetilde{f}_{m,i_k}(y)) - D_m(\widetilde{f}_{m,i_k}(\widetilde{y}))| \leq 4d_{i_k} \right\} = 1$$

where  $\widetilde{y} = \arg \min_{z \in \mathcal{G}} \|z - y\|$  with probability at least  $1 - \delta' / \ell^2$ .

*Proof.* The proof is identical to the proof of [Claim A.2](#).  $\square$

We now break into two cases depending on how  $d_j$  compares to  $d_{i_k}$  and handle them separately.

**Case 1:**  $d_j \leq d_{i_k} (\ell \log d_{\max})^{20\ell}$ . In this case, define the sets  $\mathcal{S}, \mathcal{Q}$  as follows:

$$\begin{aligned} \mathcal{S} &:= \left\{ (S_m)_{m=i_k+1}^{j-1} : S_m \subseteq [d_m], |S_m| \leq 4d_{i_k} \right\}, \\ \mathcal{Q} &:= \left\{ Q : Q \subset [d_j], |Q| \leq \frac{d_j}{(\ell \log d_{\max})^{60\ell}} \right\}. \end{aligned}$$

Now, for  $Q^1, Q^2 \in \mathcal{Q}, y \in \mathcal{G}, S^1, S^2 \in \mathcal{S}$ , define the random vector:

$$V_{y, Q^1, Q^2, S^1, S^2}^{j,k} = H(D_{Q^1} - D_{Q^2}) W_j \prod_{m=j-1}^{i_k+1} (D_m(\widetilde{f}_{m,i_k}(y)) + (D_{S_m^1} - D_{S_m^2})) W_m.$$

Note that we have from [Lemma 2.1](#):

$$\begin{aligned} V_{y, Q^1, Q^2, S^1, S^2}^{j,k} &\stackrel{d}{=} H(D_{Q^1} - D_{Q^2}) W_j \prod_{m=j-1}^{i_k+1} (D_m + (D_{S_m^1} - D_{S_m^2})) W_m \text{ where} \\ (D_m)_{i,j} &= \begin{cases} 1, & \text{w.p } 1/2 \text{ if } i = j \\ 0, & \text{otherwise} \end{cases}. \end{aligned}$$

Note that  $|\mathcal{Q}|^2 \cdot |\mathcal{G}| \cdot |\mathcal{S}|^2 \leq (d_{\max})^{64\ell d_{i_k}}$ , because of the condition on  $d_j$  in this case. We now get:

$$\begin{aligned} &\left\| V_{y, Q^1, Q^2, S^1, S^2}^{j,k} \right\| \\ &\leq \left\| H(D_{Q^1} - D_{Q^2}) W_j \prod_{m=j-1}^{i_k+2} (D_m + (D_{S_m^1} - D_{S_m^2})) W_m \right\| \cdot \left( 1 + \sqrt{\frac{\log 1 / \delta^\dagger}{d_{i_k}}} \right) \\ &\leq \|H(D_{Q^1} - D_{Q^2})\| \cdot \prod_{m=i_k+1}^j \left( 1 + \sqrt{\frac{\log 1 / \delta^\dagger}{d_{m-1}}} \right) \\ &\leq \|H(D_{Q^1} - D_{Q^2})\| \cdot (C\ell \log d_{\max})^{(j-i_k)/2} \end{aligned}$$

where the final inequality follows from the fact that  $d_m \geq d_{i_k}$  and by setting  $\delta^\dagger = \delta / (32 \cdot \ell^4 \cdot |\mathcal{Q}|^2 \cdot |\mathcal{G}| \cdot |\mathcal{S}|^2)$ . A union bound now implies that the previous conclusion holds for all  $Q^1, Q^2 \in \mathcal{Q}, y \in \mathcal{G}, S^1, S^2 \in \mathcal{S}$ . For the first term, we use the trivial bound:

$$\|H(D_{Q^1} - D_{Q^2})\| \leq \sqrt{|Q^1| + |Q^2|} \cdot \|H\|_\infty$$

$$\leq 10 \cdot \sqrt{\frac{d_j}{(\ell \log d_{\max})^{60\ell}}} \cdot \sqrt{\frac{\|H\|^2}{d_j} \cdot (\log d_{\max} + \log \ell/\delta)}.$$

with probability at least  $1 - \delta/(32 \cdot \ell^4)$  from [Lemma C.6](#) with  $M$  set to the standard basis vectors.

To conclude the proof, we get from [Lemmas 2.3](#) and [2.5](#), for all  $y$  s.t  $\|y - x\| \leq R$ :

$$\begin{aligned} & \text{Tr}|D_j(\tilde{f}_j(x)) - D_j(\tilde{f}_j(y))| \\ & \leq \# \left\{ i : |\langle (W_j)_i, f_{j-1}(x) \rangle| \leq \frac{\|f_{j-1}(x)\|}{4\sqrt{d_{j-1}} \cdot (\ell \log d_{\max})^{75\ell}} \right\} \\ & \quad + \frac{\|\tilde{f}_j(x) - \tilde{f}_j(y)\|^2}{\|f_{j-1}(x)\|^2} \cdot 16 \cdot d_{j-1} \cdot (\ell \log d_{\max})^{150\ell} \\ & \leq \frac{d_j}{2(\ell \log d_{\max})^{75\ell}} + d_j \left( \frac{R^2}{d_{\min}} \right) \frac{(C\ell \log d_{\max})^\ell}{d_{j-1}} \cdot 16d_{j-1} (\ell \log d_{\max})^{150\ell} \leq \frac{d_j}{(\ell \log d_{\max})^{60\ell}}. \end{aligned}$$

where the first inequality follows from the fact that for all  $t > 0$ :

$$|T| \leq \frac{\|\tilde{f}_j(x) - \tilde{f}_j(y)\|^2}{t^2} \text{ where}$$

$$T := \{i : |\langle (W_j)_i, f_{j-1}(x) \rangle| \geq t \text{ and } \text{sign}(\langle (W_j)_i, f_{j-1}(x) \rangle) \neq \text{sign}(\langle (W_j)_i, f_{j-1}(y) \rangle)\}.$$

Hence, we get with probability at least  $1 - \delta/(16\ell^4)$ :

$$\forall \|y - x\| \leq R : \text{Tr}|D_j(\tilde{f}_j(y)) - D_j(\tilde{f}_j(x))| \leq \frac{d_j}{(\ell \log d_{\max})^{60\ell}},$$

$$\forall \|y - x\| \leq R, \forall m \in \{i_k + 1, \dots, j-1\} : \text{Tr}|D_m(\tilde{f}_m(y)) - D_m(\tilde{f}_{m,i_k}(\tilde{y}))| \leq 4d_{i_k},$$

where  $\tilde{y} = \arg \min_{z \in \mathcal{G}} \|z - f_{i_k}(y)\|$ .

To conclude [\(5\)](#), note that for all  $y \in \mathcal{G}$  we have that  $(\tilde{f}_{m,i_k}(y))_n \neq 0$  almost surely on [Lemma 2.3](#). Hence, all the  $D_m(\tilde{f}_{m,i_k}(y))$  are deterministic on [Lemma 2.3](#) and similarly for  $x$ . This immediately yields [\(5\)](#) for  $y \in \mathcal{G}$ . For  $y \notin \mathcal{G}$ , the conclusion follows from the previous discussion and [Claim A.3](#).

**Case 2:**  $d_j \geq d_{i_k}(\ell \log d_{\max})^{20\ell}$ . As in the previous case, we start by defining the sets  $\mathcal{S}$ :

$$\mathcal{S} = \left\{ (S_m)_{m=i_k+1}^{j-1} : S_m \subseteq [d_m], |S_m| \leq 4 \cdot d_{i_k} \right\}.$$

Now, for  $y \in \mathcal{G}$ ,  $S^1, S^2 \in \mathcal{S}$ , consider the random matrix  $M_{y,S^1,S^2}^{j,k}$  and an application of [Lemma 2.1](#):

$$\begin{aligned} M_{y,S^1,S^2}^{j,k} & := W_{j-1} \prod_{m=j-2}^{i_k+1} (D_m(\tilde{f}_m(y)) + (D_{S_m^1} - D_{S_m^2}))W_m \\ M_{y,S^1,S^2}^{j,k} & \stackrel{d}{=} \tilde{W}_{j-1} \prod_{m=j-2}^{i_k+1} (D_m + (D_{S_m^1} - D_{S_m^2}))\tilde{W}_m. \end{aligned}$$

We will bound the spectral norm of  $M_{y,S^1,S^2}^{j,k}$  for all  $y, S^1, S^2$  with high probability as follows. Let  $\mathcal{V}$  be a  $1/9$ -grid of  $\mathbb{S}^{d_{i_k}}$  and  $v \in \mathcal{V}$ . We have:

$$\begin{aligned} \left\| M_{y,S^1,S^2}^{j,k} v \right\| & \leq \sqrt{\frac{d_{j-1}}{d_{j-2}}} \cdot \left( 1 + \sqrt{\frac{\log 1/\delta^\dagger}{d_{j-1}}} \right) \cdot \left\| \prod_{m=j-2}^{i_k+1} (D_m + (D_{S_m^1} - D_{S_m^2}))\tilde{W}_m v \right\| \\ & \leq 2^{(j-i_k+1)} \sqrt{\frac{d_{j-1}}{d_{i_k}}} \cdot \prod_{m=i_k+1}^{j-1} \left( 1 + \sqrt{\frac{\log 1/\delta^\dagger}{d_m}} \right) \\ & \leq \sqrt{\frac{d_{j-1}}{d_{i_k}}} \cdot (C\ell \log d_{\max})^{(j-i_k+1)/2} \end{aligned} \tag{7}$$

where the final inequality follows by the fact that  $d_m \geq d_{i_k}$  and by setting  $\delta^\dagger = \delta/(32 \cdot \ell \cdot |\mathcal{G}| \cdot |\mathcal{S}|^2 \cdot \nu)$ . By a union bound, the bound on the spectral norms of  $M_{y, S^1, S^2}^{j, k}$  follows. We now condition on this event and the conclusion of [Lemma 2.3](#) up to layer  $j - 1$  for the rest of the proof, which as before implies that  $M_{y, S^1, S^2}^{j, k}$  and  $D_{j-1}(\tilde{f}_{j-1, i_k}(y))$  are no longer random.

Let  $\tilde{M}_{y, S^1, S^2}^{j, k} = D_{j-1}(\tilde{f}_{j-1, i_k}(y)) \cdot M_{y, S^1, S^2}^{j, k}$  and for  $y \in \mathcal{G}$ , let  $\tilde{y} = f_{j-1, i_k}(y)$  and  $\tilde{x} = f_{j-1}(x)$ . We have by an application of [Lemma A.4](#) and a union bound over all  $y \in \mathcal{G}, S^1, S^2 \in \mathcal{S}$ :

$$\forall y \in \mathcal{G}, S^1, S^2 \in \mathcal{S} : \left\| H^\top \left( D_j(\tilde{f}_{j, i_k}(y)) - D_j(\tilde{f}_j(x)) \right) W_j \tilde{M}_{y, S^1, S^2}^{j, k} \right\| \leq \left( \frac{C}{\ell \log d_{\max}} \right)^{3\ell}$$

with probability at least  $1 - \delta'/\ell^2$  by recalling that  $d_j \geq d_{i_k}(\ell \log d_{\max})^{20\ell}$  in this case.

We now additionally condition on  $\|H\|_\infty$ . We have as a consequence of [Lemma C.6](#) that:

$$\|H\|_\infty \leq 4 \cdot 2^{\ell+1-j} \cdot \sqrt{\frac{\log d_j + \log 1/\delta^\dagger}{d_j}}$$

with probability at least  $1 - \delta^\dagger$ . We condition on this event and proceed as follows. Let  $T \subset [d_j]$  such that  $|T| \leq 4 \cdot d_{i_k}$  and  $y \in \mathcal{G}, S^1, S^2 \in \mathcal{S}$  and we observe:

$$\left\| H^\top D_T \tilde{W}_j \tilde{M}_{y, S^1, S^2}^{j, k} \right\| \leq \left\| H^\top D_T \tilde{W}_j U_{y, S^1, S^2}^{j, k} \right\| \cdot \left\| \tilde{M}_{y, S^1, S^2}^{j, k} \right\|$$

where  $U_{y, S^1, S^2}^{j, k}$  are the left singular vectors of  $\tilde{M}_{y, S^1, S^2}^{j, k}$  and observe:

$$\begin{aligned} \mathbb{P} \left\{ \left\| H^\top D_T \tilde{W}_j U_{y, S^1, S^2}^{j, k} \right\| \geq t \right\} &= \mathbb{P}_{\tilde{Z} \sim \mathcal{N}(0, \|H D_T\|^2 \cdot \frac{t}{d_j})} \left\{ \|\tilde{Z}\| \geq t \right\} \\ &\leq \mathbb{P}_{Z \sim \mathcal{N}(0, \|H\|_\infty^2 \cdot |T| \cdot \frac{t}{d_j})} \left\{ \|Z\| \geq t \right\} \end{aligned}$$

to get that (conditioned on  $H$  and noting that  $\text{rank}(U_{y, S^1, S^2}^{j, k}) \leq d_{i_k}$ ), with probability  $1 - \delta^\dagger$ :

$$\left\| H^\top D_T U_{y, S^1, S^2}^{j, k} \right\| \leq \|H\|_\infty \cdot \sqrt{|T|} \cdot \sqrt{\frac{d_{i_k}}{d_j}} \cdot \left( 1 + \sqrt{\frac{\log 1/\delta^\dagger}{d_{i_k}}} \right)$$

By setting  $\delta^\ddagger = \delta/(128\ell^2 d_j^{4d_{i_k}} |\mathcal{S}|^2 |\mathcal{G}|)$  and our bounds on  $d_j, \|H\|_\infty$  yield:

$$\forall T \subset [d_j] \text{ s.t } |T| \leq 4d_{i_k}, y \in \mathcal{G}, S^1, S^2 \in \mathcal{S} : \left\| H^\top D_T U_{y, S^1, S^2}^{j, k} \right\| \leq \frac{1}{(\ell \log d_{\max})^{4\ell}}$$

with probability at least  $1 - \delta/(64\ell^2)$  again by recalling  $d_j \geq d_{i_k}(\ell \log d_{\max})^{20\ell}$ .

To conclude our proof of [\(5\)](#), we proceed similarly to the previous case. As before, for all  $y \in \mathcal{G}$  we have that  $(\tilde{f}_{m, i_k}(y))_n \neq 0$  almost surely on [Lemma 2.3](#) and similarly for  $x$ . Hence, all the  $D_m(\tilde{f}_{m, i_k}(y))$  are deterministic on [Lemma 2.3](#). This immediately yields [\(5\)](#) for  $y \in \mathcal{G}$ . For  $y \notin \mathcal{G}$ , the conclusion follows from the previous discussion and [Claim A.3](#).

A union bound over all  $j \in [\ell]$ , an application of the triangle inequality with [GD-DECOMP](#) and [\(4\)](#) with [Lemma 2.6](#) conclude the proof of the lemma.  $\square$

**Lemma A.4.** *Let  $L, R, r > 0, k, m, n \in \mathbb{N}, M \in \mathbb{R}^{m \times n}$  such that  $m \geq n$  and  $R > r$ . Furthermore, suppose  $W \in \mathbb{R}^{k \times m}$  and  $H$  are distributed as follows:*

$$W = \begin{bmatrix} w_1^\top \\ w_2^\top \\ \vdots \\ w_k^\top \end{bmatrix} \text{ with } w_i \stackrel{i.i.d.}{\sim} \mathcal{N}(0, I/m), \text{ and } H \sim \text{Unif}(L \cdot \mathbb{S}^{k-1}).$$

Furthermore, suppose  $x \in \mathbb{R}^m$  satisfies  $\|x\| \geq R$  and  $y$  be such that  $\|y - x\| \leq r$ . Then:

$$\mathbb{P} \left\{ \left\| H^\top D W M \right\| \geq 1024 L \|M\| \left( \sqrt{\frac{3r}{R} \log R/r} \left( \frac{n + \log 1/\delta^\dagger}{m} \right) + \left( \frac{n + \log 1/\delta^\dagger}{\sqrt{km}} \right) \right) \right\} \leq \delta^\ddagger$$

$$\text{where } D = D_x - D_y \text{ with } (D_x)_{i,j} = \begin{cases} 1, & \text{if } i = j \text{ and } w_i^\top x > 0 \\ 1, & \text{w.p } 1/2 \text{ if } i = j \text{ and } w_i^\top x = 0. \\ 0, & \text{otherwise} \end{cases}$$

*Proof.* We may discard the cases where  $w_i^\top x, w_i^\top y = 0$  as these form a measure 0 set. Now, it suffices to analyze the random variable:

$$\tilde{H}^\top DWM \text{ for } \tilde{H} \sim \mathcal{N}(0, 2L^2 \cdot I/k).$$

Let  $V$  denote the two-dimensional subspace of  $\mathbb{R}^m$  containing  $x, y$ . We now decompose the norm as follows:

$$\|\tilde{H}^\top DWM\| \leq \|\tilde{H}^\top DW\mathcal{P}_V M\| + \|\tilde{H}^\top DW\mathcal{P}_V^\perp M\| \leq \|\tilde{H}^\top DW\mathcal{P}_V\| \cdot \|M\| + \|\tilde{H}^\top DW\mathcal{P}_V^\perp M\|. \quad (8)$$

To apply Bernstein's inequality, we first expand on the first term:

$$\left(\tilde{H}^\top DW\mathcal{P}_V\right)^\top = \sum_{i=1}^k \tilde{H}_i D_{i,i} \mathcal{P}_V w_i.$$

Letting  $Z_i = \tilde{H}_i D_{i,i} \mathcal{P}_V w_i$ , we note that  $\mathbb{E}[Z_i] = 0$  and bound its even moments as follows, in an orthonormal basis  $\{v_1, v_2\}$  for  $V$ :

$$\begin{aligned} \mathbb{E}[\langle v_j, Z_i \rangle^2] &= \mathbb{E}\left[\tilde{H}_i^2 \cdot D_{i,i}^2 \cdot \langle v_j, w_i \rangle^2\right] \\ &\leq \frac{2L^2}{k} \cdot \mathbb{E}[\|\mathcal{P}_V w_i\|^2] \cdot \mathbb{P}\{D_{i,i} \neq 0\} = \frac{2L^2}{k} \cdot \frac{2}{m} \cdot \mathbb{P}\{D_{i,i} \neq 0\} \\ \mathbb{E}[\langle v_j, Z_i \rangle^\ell] &= \mathbb{E}\left[\tilde{H}_i^\ell \cdot D_{i,i}^\ell \cdot \langle v_j, w_i \rangle^\ell\right] \leq (\ell-1)!! \left(\frac{2L^2}{k}\right)^{\ell/2} \cdot \mathbb{E}[\|\mathcal{P}_V w_i\|^\ell] \cdot \mathbb{P}\{D_{i,i} \neq 0\} \\ &\leq (\ell-1)!! \left(\frac{2L^2}{k}\right)^{\ell/2} \cdot \left(2^{\ell/2} \cdot \mathbb{E}[\langle w_i, v_1 \rangle^\ell + \langle w_i, v_2 \rangle^\ell]\right) \cdot \mathbb{P}\{D_{i,i} \neq 0\} \\ &= 2 \cdot (\ell-1)!! \left(\frac{4L^2}{k}\right)^{\ell/2} \cdot \mathbb{E}[\langle w_i, v_1 \rangle^\ell] \cdot \mathbb{P}\{D_{i,i} \neq 0\} \\ &= 2 \cdot ((\ell-1)!!)^2 \left(\frac{4L^2}{km}\right)^{\ell/2} \cdot \mathbb{P}\{D_{i,i} \neq 0\} \leq 2 \cdot \ell! \cdot \left(\frac{8L^2}{km}\right)^{\ell/2} \cdot \mathbb{P}\{D_{i,i} \neq 0\}. \end{aligned}$$

For odd  $\ell \geq 3$  we have by similar manipulations:

$$\begin{aligned} \mathbb{E}[|\langle v_j, Z_i \rangle|^\ell] &= \mathbb{E}\left[|\tilde{H}_i^\ell \cdot D_{i,i}^\ell \cdot \langle v_j, w_i \rangle^\ell|\right] \leq \mathbb{P}\{D_{i,i} \neq 0\} \cdot \mathbb{E}\left[|\tilde{H}_i^\ell \cdot \|\mathcal{P}_V w_i\|^\ell\right] \\ &\leq \mathbb{P}\{D_{i,i} \neq 0\} \cdot \sqrt{\mathbb{E}\left[\left(\tilde{H}_i \cdot \|\mathcal{P}_V w_i\|\right)^{2\ell}\right]} \leq \mathbb{P}\{D_{i,i} \neq 0\} \cdot \sqrt{2 \cdot (2\ell)! \cdot \left(\frac{8L^2}{km}\right)^\ell} \\ &\leq \mathbb{P}\{D_{i,i} \neq 0\} \cdot \sqrt{2 \cdot (2\ell)!} \cdot \left(\frac{8L^2}{km}\right)^{\ell/2} \leq \mathbb{P}\{D_{i,i} \neq 0\} \cdot 2^\ell \ell! \cdot \left(\frac{8L^2}{km}\right)^{\ell/2} \\ &\leq \mathbb{P}\{D_{i,i} \neq 0\} \cdot \ell! \cdot \left(\frac{32L^2}{km}\right)^{\ell/2} \end{aligned}$$

A union bound and an application of [Theorem C.2](#) (with  $\nu = 4L^2 \mathbb{P}\{D_{i,i} \neq 0\} / m$  and  $c = 6L/\sqrt{km}$ ) gives:

$$\left\|\sum_{i=1}^k Z_i\right\| \leq 2 \max_{j \in \{1,2\}} \left|\sum_{i=1}^k \langle v_j, Z_i \rangle\right| \leq 256 \cdot L \cdot \left(\sqrt{\frac{1}{m} \cdot \mathbb{P}\{D_{i,i} \neq 0\} \cdot \log 1/\delta'} + \frac{\log 1/\delta'}{\sqrt{km}}\right)$$

with probability at least  $1 - \delta'$ .

For the other term in Eq. (8), we use  $U$  to denote the left singular subspace of  $M$  and  $Y$  to denote the span of the left singular vectors of  $\mathcal{P}_V^\perp \mathcal{P}_U$  with non-zero singular values. Noting  $\|\mathcal{P}_V^\perp \mathcal{P}_U\| \leq 1$ , we now have:

$$\|\tilde{H}^\top DW \mathcal{P}_V^\perp M\| \leq \|M\| \cdot \|\tilde{H}^\top DW \mathcal{P}_V^\perp \mathcal{P}_U\| \leq \|M\| \cdot \|\tilde{H}^\top DW \mathcal{P}_Y\|.$$

We expand the term on the right as follows:

$$\left(\tilde{H}^\top DW \mathcal{P}_Y\right)^\top = \sum_{i=1}^k \tilde{H}_i D_{i,i} \mathcal{P}_Y w_i = \sum_{i=1}^k D_{i,i} \cdot \left(\tilde{H}_i \mathcal{P}_Y w_i\right) =: \sum_{i=1}^k Y_i.$$

Note that  $Y$  is orthogonal to  $V$  as  $\mathcal{P}_V(\mathcal{P}_V^\perp \mathcal{P}_U u) = 0$  for all  $u$  and hence,  $\tilde{H}_i \mathcal{P}_Y w_i$  and  $D_{i,i}$  are independent random variables (due to  $\mathcal{P}_V w_i$  and  $\mathcal{P}_V^\perp w_i$  being independent). Now, fix  $y \in Y$  s.t.  $\|y\| = 1$  and we bound the directional even moments of  $Y_i$  with the aim of applying Bernstein's inequality as before:

$$\begin{aligned} \mathbb{E} [\langle y, Y_i \rangle^2] &= \mathbb{E} \left[ D_{i,i}^2 \cdot \tilde{H}_i^2 \cdot \langle y, w_i \rangle^2 \right] = \frac{2L^2}{k} \cdot \frac{1}{m} \cdot \mathbb{P} \{D_{i,i} \neq 0\} \\ \mathbb{E} [\langle y, Y_i \rangle^\ell] &= \mathbb{E} \left[ D_{i,i}^\ell \cdot \tilde{H}_i^\ell \cdot \langle y, w_i \rangle^\ell \right] \\ &= ((\ell - 1)!!)^2 \cdot \left(\frac{2L^2}{km}\right)^{\ell/2} \cdot \mathbb{P} \{D_{i,i} \neq 0\} \leq \ell! \cdot \left(\frac{2L^2}{km}\right)^{\ell/2} \cdot \mathbb{P} \{D_{i,i} \neq 0\}. \end{aligned}$$

Similarly, for odd  $\ell \geq 3$ , we have by similar manipulations:

$$\begin{aligned} \mathbb{E} [|\langle y, Y_i \rangle|^\ell] &= \mathbb{E} \left[ \left| \tilde{H}_i^\ell \cdot D_{i,i}^\ell \cdot \langle y, w_i \rangle^\ell \right| \right] \leq \mathbb{P} \{D_{i,i} \neq 0\} \cdot \mathbb{E} \left[ |\tilde{H}_i^\ell \cdot \langle y, w_i \rangle^\ell| \right] \\ &\leq \mathbb{P} \{D_{i,i} \neq 0\} \cdot \sqrt{\mathbb{E} \left[ |\tilde{H}_i^\ell \cdot \langle y, w_i \rangle^{2\ell}| \right]} \leq \mathbb{P} \{D_{i,i} \neq 0\} \cdot \sqrt{(2\ell)! \cdot \left(\frac{2L^2}{km}\right)^\ell} \\ &\leq \mathbb{P} \{D_{i,i} \neq 0\} \cdot \sqrt{(2\ell)!} \cdot \left(\frac{2L^2}{km}\right)^{\ell/2} \leq \mathbb{P} \{D_{i,i} \neq 0\} \cdot 2^\ell \ell! \cdot \left(\frac{2L^2}{km}\right)^{\ell/2} \\ &\leq \mathbb{P} \{D_{i,i} \neq 0\} \cdot \ell! \cdot \left(\frac{8L^2}{km}\right)^{\ell/2} \end{aligned}$$

Now, consider a 1/3-net of  $\mathcal{G} := Y \cap \mathbb{S}^{m-1}$ . We get for any  $z \in Y$ :

$$\|z\| = \max_{y \in \mathbb{S}^{m-1}} \langle y, z \rangle = \max_{y \in \mathbb{S}^{m-1}} (\langle y - \tilde{y}, z \rangle + \langle \tilde{y}, z \rangle) \leq \frac{\|z\|}{3} + \max_{y \in \mathcal{G}} \langle y, z \rangle \implies \|z\| \leq \frac{3}{2} \max_{y \in \mathcal{G}} \langle y, z \rangle,$$

where  $\tilde{y} = \arg \min_{z \in \mathcal{G}} \|y - z\|$ . Since the rank of  $Y$  is at most  $n$ , We may assume  $|\mathcal{G}| \leq (30)^n$  [Ver18, Corollary 4.2.13]. By a union bound and Theorem C.2:

$$\left\| \sum_{i=1}^n Y_i \right\| \leq 2 \max_{y \in \mathcal{G}} \left\langle y, \sum_{i=1}^n Y_i \right\rangle \leq 256L \left( \sqrt{\mathbb{P} \{D_{i,i} \neq 0\}} \cdot \left(\frac{n + \log 1/\delta'}{m}\right) + \left(\frac{n + \log 1/\delta'}{\sqrt{km}}\right) \right)$$

with probability at least  $1 - \delta'$ . The lemma follows from the previous discussion and Lemma C.5 by picking  $\delta' = \delta^\dagger/4$  and applying a union bound.  $\square$

## B Proof of Theorem 3.1

We will now construct an architecture that when randomly initialized approximately maps every input to a random constant times its Euclidean norm. We will adopt the notation from previous sections; the output of our neural network denoted by  $f$  with  $\ell$  hidden layers all of fixed width  $k$  is defined as follows:

$$f(x) = W_{\ell+1} \cdot \sigma(W_\ell \cdot \sigma(\dots \sigma(W_1 \cdot x))) \text{ where } \sigma(x)_i = \max\{x_i, 0\}$$

$$\begin{aligned}
f_i(x) &= \sigma(W_i \cdot \sigma(\cdots \sigma(W_1 \cdot x))) \\
\forall i \in [\ell] \setminus \{1\} : W_i &\in \mathbb{R}^{k \times k} \text{ with } (W_i)_{m,n} \stackrel{iid}{\sim} \mathcal{N}\left(0, \frac{2}{k}\right) \\
W_1 &\in \mathbb{R}^{k \times d} \text{ with } (W_1)_{m,n} \stackrel{iid}{\sim} \mathcal{N}\left(0, \frac{2}{d}\right) \quad W_{\ell+1} \sim \mathcal{N}(0, 2I/k) \quad (\text{Lower-Bound-Init})
\end{aligned}$$

The scaling for the intermediate layers is chosen such that it preserves the length of the input from the previous layer. We now prove the main result of the section which implies [Theorem 3.1](#) by a simple rescaling:

**Theorem B.1.** *Fix a sufficiently large  $d \in \mathbb{N}$ , an  $\ell \geq d^3$  and  $(\ell d)^{20} \leq k \leq \exp(\sqrt{\ell})$ , and consider the randomly initialized neural network ([Lower-Bound-Init](#)). There is a universal constant  $C$  such that with probability at least 0.9,*

$$\begin{aligned}
\forall x \in \mathbb{S}^{d-1} : |f(x)| &\geq 0.04 \\
\forall x, y \in \mathbb{S}^{d-1} : |f(x) - f(y)| &\leq C \sqrt{\frac{\log d}{d}}
\end{aligned}$$

*Proof.* We start by picking an  $\varepsilon$ -net of  $\mathbb{S}^{d-1}$ ,  $\mathcal{G}$  with  $\varepsilon = (\frac{1}{10 \cdot 2^\ell})^{10}$ . Note we may assume that  $|\mathcal{G}| \leq (\frac{10}{\varepsilon})^d$ . Now, for fixed  $x \in \mathcal{G}$  and defining  $f_0(x) := x$ , we have:

$$\begin{aligned}
\forall i \in [\ell] : \|\sigma(W_i f_{i-1}(x))\| - \|\sigma(W'_i f_{i-1}(x))\| &\leq \|\sigma(W_i f_{i-1}(x)) - \sigma(W'_i f_{i-1}(x))\| \\
&\leq \|W_i - W'_i\| \cdot \|f_{i-1}(x)\| \leq \|W_i - W'_i\|_F \cdot \|f_{i-1}(x)\|.
\end{aligned}$$

Hence,  $\|\sigma(W_i f_{i-1}(x))\|$  is a  $\|f_{i-1}(x)\|$ -Lipschitz function of  $W_i$  and we get by an application of [Theorem C.3](#) (note that  $(W_i)_{l,m} \sim \mathcal{N}(0, 2/k)$ ) and a union bound over the  $\ell$  layers:

$$\forall i \in [\ell] : \|\|f_i(x)\| - \mathbb{E}[\|f_i(x)\| \mid f_{i-1}(x)]\| \leq 8 \|f_{i-1}(x)\| \cdot \sqrt{\frac{\log \delta + \log \ell}{k}}$$

with probability at least  $1 - \delta$ . By setting  $\delta = \frac{1}{16 \cdot |\mathcal{G}| \cdot d^{10}}$  and a union bound over all  $x \in \mathcal{G}$ , we have with probability at least  $1 - 1/(16d^{10})$ :

$$\forall x \in \mathcal{G}, i \in [\ell] : \|\|f_i(x)\| - \mathbb{E}[\|f_i(x)\| \mid f_{i-1}(x)]\| \leq \|f_{i-1}(x)\| \cdot \frac{1}{2048 \cdot (\ell d)^3}. \quad (9)$$

We also have by Jensen's inequality:

$$\mathbb{E}[\|f_i(x)\| \mid f_{i-1}(x)] \leq \sqrt{\mathbb{E}[\|f_i(x)\|^2 \mid f_{i-1}(x)]} = \|f_{i-1}(x)\|$$

and by integrating the tail bound from [Theorem C.3](#):

$$\mathbb{E}\left[\left(\|f_i(x)\| - \mathbb{E}[\|f_i(x)\| \mid f_{i-1}(x)]\right)^2 \mid f_{i-1}(x)\right] \leq \frac{32}{k} \cdot \|f_{i-1}(x)\|^2.$$

Hence, we get:

$$\left(1 - \frac{32}{k}\right) \|f_{i-1}(x)\| \leq \mathbb{E}[\|f_i(x)\| \mid f_{i-1}(x)] \leq \|f_{i-1}(x)\|.$$

From the above display and [Eq. \(9\)](#) and noting  $(1+x) \leq e^x \leq (1+2x)$  for  $0 \leq x \leq 1$ , we get:

$$\|f_i(x)\| \leq \left(1 + \frac{1}{2048 \cdot (\ell d)^3}\right)^i \leq 1 + \frac{i}{1024 \ell^3 d^3}$$

and similarly for the lower bound:

$$\|f_i(x)\| \geq \left(1 - \frac{32}{k} - \frac{1}{2048 \cdot (\ell d)^3}\right)^i \geq 1 - \frac{i}{512 \ell^3 d^3}.$$

Putting these together, we obtain:

$$1 - \frac{i}{512 \ell^3 d^3} \leq \|f_i(x)\| \leq 1 + \frac{i}{1024 \ell^3 d^3} \quad (10)$$

Similarly to the previous discussion, we have for all  $x, y \in \mathcal{G}$ :

$$\begin{aligned} & \forall i \in [\ell] : \|\sigma(W_i f_{i-1}(x)) - \sigma(W_i f_{i-1}(y))\| - \|\sigma(W'_i f_{i-1}(x)) - \sigma(W'_i f_{i-1}(y))\| \\ & \leq \|\sigma(W_i f_{i-1}(x)) - \sigma(W_i f_{i-1}(y)) - \sigma(W'_i f_{i-1}(x)) + \sigma(W'_i f_{i-1}(y))\| \\ & \leq \|W_i - W'_i\| \cdot (\|f_{i-1}(x)\| + \|f_{i-1}(y)\|) \leq \|W_i - W'_i\|_F \cdot (\|f_{i-1}(x)\| + \|f_{i-1}(y)\|). \end{aligned}$$

Therefore, another application of [Theorem C.3](#) and a union bound over the  $\ell$  layers yields:

$$\begin{aligned} & \forall i \in [\ell] : \|\|f_i(x) - f_i(y)\| - \mathbb{E}[\|f_i(x) - f_i(y)\| \mid f_{i-1}(x), f_{i-1}(y)]\| \\ & \leq 8(\|f_{i-1}(x)\| + \|f_{i-1}(y)\|) \cdot \sqrt{\frac{\log \ell / \delta}{k}}. \end{aligned}$$

Setting  $\delta = 1/(16 \cdot |\mathcal{G}|^2 \cdot d^{10})$  and a union bound over all  $x, y \in \mathcal{G}$  yields with probability at least  $1 - 1/(16d^{10})$  for all  $x, y \in \mathcal{G}, i \in [\ell]$ :

$$\|\|f_i(x) - f_i(y)\| - \mathbb{E}[\|f_i(x) - f_i(y)\| \mid f_{i-1}(x), f_{i-1}(y)]\| \leq \|f_{i-1}(x) + f_{i-1}(y)\| \frac{1}{2048 \cdot (\ell d)^3}. \quad (11)$$

We only need an upper bound on  $\mathbb{E}[\|f_i(x) - f_i(y)\| \mid f_{i-1}(x), f_{i-1}(y)]$ . Before, we do we need the following simple fact:

**Fact B.2.** *We have for some  $c > 0$ :*

$$\forall x \in [0, \pi] : \sin(x) - x \cos(x) \geq \frac{(1 - \cos x)^{3/2}}{15}.$$

*Proof.* Let  $f(x) = \sin(x) - x \cos(x)$  we have:

$$\begin{aligned} & \forall x \in \left[0, \frac{\pi}{2}\right] : f'(x) = x \sin(x) \geq \frac{2}{\pi} \cdot x^2 \\ & \forall x \in [0, \pi] : f'(x) = x \sin(x) \geq 0. \end{aligned}$$

Therefore, we have:

$$\begin{aligned} & \forall x \in \left[0, \frac{\pi}{2}\right] : f(x) = \int_0^x f'(x) dx \geq \frac{2}{3\pi} \cdot x^3 \\ & \forall x \in \left[\frac{\pi}{2}, \pi\right] : f(x) = \int_0^x f'(x) dx \geq \int_0^{\pi/2} f'(x) dx \geq \frac{\pi^2}{12} \geq \frac{x^3}{36}. \end{aligned}$$

By noting that  $1 - \cos x \leq x^2/2$ , we get:

$$\forall x \in [0, \pi] : f(x) \geq \frac{(1 - \cos(x))^{3/2}}{15}.$$

□

Now, defining  $\theta_i = \arccos(\langle f_i(x), f_i(y) \rangle / (\|f_i(x)\| \|f_i(y)\|))$ , we have:

$$\begin{aligned} & \mathbb{E}[\|f_i(x) - f_i(y)\| \mid f_{i-1}(x), f_{i-1}(y)] \\ & \leq \sqrt{\mathbb{E}[\|f_i(x) - f_i(y)\|^2 \mid f_{i-1}(x), f_{i-1}(y)]} \\ & = \sqrt{\|f_{i-1}(x)\|^2 + \|f_{i-1}(y)\|^2 + 2\mathbb{E}[\langle f_i(x), f_i(y) \rangle \mid f_{i-1}(x), f_{i-1}(y)]} \\ & = \sqrt{\|f_{i-1}(x)\|^2 + \|f_{i-1}(y)\|^2 - 2\|f_{i-1}(x)\| \|f_{i-1}(y)\| \cdot \left(\frac{\sin \theta_{i-1}}{\pi} + \left(1 - \frac{\theta_{i-1}}{\pi}\right) \cos \theta_{i-1}\right)} \\ & = \sqrt{\|f_{i-1}(x) - f_{i-1}(y)\|^2 - 2\|f_{i-1}(x)\| \|f_{i-1}(y)\| \cdot \left(\frac{\sin \theta_{i-1} - \theta_{i-1} \cos \theta_{i-1}}{\pi}\right)} \\ & \leq \sqrt{\|f_{i-1}(x) - f_{i-1}(y)\|^2 - 2\|f_{i-1}(x)\| \|f_{i-1}(y)\| \cdot \left(\frac{\sin \theta_{i-1} - \theta_{i-1} \cos \theta_{i-1}}{\pi}\right)} \end{aligned}$$



$$\begin{aligned}
&\leq \|f_{i-1}(x) - f_{i-1}(y)\| \cdot \sqrt{1 - \frac{2\|f_{i-1}(x)\|\|f_{i-1}(y)\|}{\|f_{i-1}(x) - f_{i-1}(y)\|^2} \cdot \frac{(1 - \cos \theta_{i-1})^{3/2}}{15\pi}} \\
&\leq \|f_{i-1}(x) - f_{i-1}(y)\| \cdot \left(1 - \frac{\|f_{i-1}(x)\|\|f_{i-1}(y)\|}{\|f_{i-1}(x) - f_{i-1}(y)\|^2} \cdot \frac{(1 - \cos \theta_{i-1})^{3/2}}{15\pi}\right). \tag{12}
\end{aligned}$$

On the event in Eq. (10), defining  $\tilde{f}_i(x) = \frac{f_i(x)}{\|f_i(x)\|}$  and similarly for  $y$ , we have:

$$\begin{aligned}
&2(1 - \cos \theta_{i-1}) \\
&= \|\tilde{f}_{i-1}(x)\|^2 + \|\tilde{f}_{i-1}(y)\|^2 - 2\langle \tilde{f}_{i-1}(x), \tilde{f}_{i-1}(y) \rangle \\
&\geq \left(1 + \frac{1}{1024\ell^2 d^3}\right)^{-2} (\|f_{i-1}(x)\|^2 + \|f_{i-1}(y)\|^2) - 2\left(1 + \frac{\text{sgn} \cos \theta_{i-1}}{256\ell^2 d^3}\right) \langle f_{i-1}(x), f_{i-1}(y) \rangle \\
&\geq \|f_{i-1}(x) - f_{i-1}(y)\|^2 - \frac{1}{256\ell^2 d^3} (\|f_{i-1}(x)\| + \|f_{i-1}(y)\|)^2 \\
&\geq \|f_{i-1}(x) - f_{i-1}(y)\|^2 - \frac{1}{48\ell^2 d^3}.
\end{aligned}$$

Therefore, we get:

$$(1 - \cos \theta_{i-1}) \geq \begin{cases} \frac{1}{4} \cdot \|f_{i-1}(x) - f_{i-1}(y)\|^2, & \text{if } \|f_{i-1}(x) - f_{i-1}(y)\|^2 \geq \frac{1}{16\ell^2 d^3} \\ 0, & \text{otherwise} \end{cases}.$$

By substituting into Eq. (12), we have:

$$\begin{aligned}
&\mathbb{E} [\|f_i(x) - f_i(y)\| \mid f_{i-1}(x), f_{i-1}(y)] \\
&\leq \|f_{i-1}(x) - f_{i-1}(y)\| \cdot \left(1 - \frac{\|f_{i-1}(x)\|\|f_{i-1}(y)\|}{\|f_{i-1}(x) - f_{i-1}(y)\|^2} \cdot \frac{(1 - \cos \theta_{i-1})^{3/2}}{15\pi}\right) \\
&\leq \|f_{i-1}(x) - f_{i-1}(y)\| \cdot \begin{cases} \left(1 - \frac{\|f_{i-1}(x) - f_{i-1}(y)\|}{1000}\right), & \text{if } \|f_{i-1}(x) - f_{i-1}(y)\|^2 \geq \frac{1}{16\ell^2 d^3} \\ 1, & \text{otherwise} \end{cases}.
\end{aligned}$$

By further substituting this into Eq. (11), we have:

$$\begin{aligned}
\forall i \in [\ell] : \|f_i(x) - f_i(y)\| &\leq \frac{1}{2048 \cdot (\ell d)^3} \\
&+ \|f_{i-1}(x) - f_{i-1}(y)\| \cdot \begin{cases} \left(1 - \frac{\|f_{i-1}(x) - f_{i-1}(y)\|}{1000}\right), & \text{if } \|f_{i-1}(x) - f_{i-1}(y)\|^2 \geq \frac{1}{16\ell^2 d^3} \\ 1, & \text{otherwise} \end{cases}.
\end{aligned}$$

For the rest of the proof, we break into two cases:

**Case 1:**  $\|f_j(x) - f_j(y)\|^2 \leq \frac{1}{4\ell^2 d^3}$  for some  $j \in [\ell]$ . In this case, we simply show that:

$$\forall i \geq [j] : \|f_i(x) - f_i(y)\|^2 \leq \frac{1}{\ell^2 d^3}.$$

Suppose for the sake of contradiction, assume the contrary and let  $i^*$  be the least index greater than  $j$  such that the above condition was violated. We have:

$$\begin{aligned}
\|f_{i^*}(y) - f_{i^*}(x)\| &\leq \frac{1}{2048 \cdot (\ell d)^3} \\
&+ \|f_{i^*-1}(x) - f_{i^*-1}(y)\| \cdot \begin{cases} \left(1 - \frac{\|f_{i^*-1}(x) - f_{i^*-1}(y)\|}{1000}\right), & \text{if } \|f_{i^*-1}(x) - f_{i^*-1}(y)\|^2 \geq \frac{1}{16\ell^2 d^3} \\ 1, & \text{otherwise} \end{cases}.
\end{aligned}$$

Now, if  $\|f_{i^*-1}(x) - f_{i^*-1}(y)\|^2 \leq \frac{1}{4(\ell^2 d^3)}$ , we have:

$$\|f_{i^*}(y) - f_{i^*}(x)\| \leq \frac{1}{2048 \cdot (\ell d)^3} + \frac{1}{2\ell d^{3/2}}$$

yielding the contradiction. Alternatively, we have:

$$\begin{aligned} \|f_{i^*}(y) - f_{i^*}(x)\| &\leq \frac{1}{2048 \cdot (\ell d)^3} + \|f_{i^*-1}(x) - f_{i^*-1}(y)\| - \frac{\|f_{i^*-1}(x) - f_{i^*-1}(y)\|^2}{1000} \\ &\leq \|f_{i^*-1}(x) - f_{i^*-1}(y)\| \end{aligned}$$

yielding a contradiction in this case as well.

**Case 2:**  $\|f_j(x) - f_j(y)\|^2 \geq \frac{1}{4\ell^2 d^3}$  for all  $j \in [\ell]$ . In this case, we have:

$$\forall i \in [\ell] : \|f_i(x) - f_i(y)\| \leq \|f_{i-1}(x) - f_{i-1}(y)\| \cdot \left(1 - \frac{\|f_{i-1}(x) - f_{i-1}(y)\|}{2000}\right).$$

Here, we prove:

$$\|f_\ell(x) - f_\ell(y)\| \leq \frac{1}{\sqrt{\ell d}}.$$

Suppose again for the sake of contradiction that the above condition is violated, then we have:

$$\|f_\ell(x) - f_\ell(y)\| \leq \|x - y\| \cdot \left(1 - \frac{1}{2000\sqrt{\ell d}}\right)^\ell \leq 2 \cdot \exp\left\{-\frac{\sqrt{\ell}}{2000d}\right\}$$

thus yielding a contradiction.

Therefore, we may assume from [Lemma C.1](#):

$$\begin{aligned} \forall x, y \in \mathcal{G} : \|f_\ell(x) - f_\ell(y)\| &\leq \frac{1}{\sqrt{\ell d}} \\ \forall x \in \mathcal{G}, i \in [\ell] : 1 - \frac{i}{512\ell^3 d^3} &\leq \|f_i(x)\| \leq 1 + \frac{i}{1024\ell^3 d^3} \\ \forall i \in [\ell] : \|W_i\| &\leq 4. \end{aligned}$$

Furthermore, we have with probability  $1 - \delta$ :

$$|f(x) - f(y)| = |W_{\ell+1}(f_\ell(x) - f_\ell(y))| \leq 2\|f_\ell(x) - f_\ell(y)\| \sqrt{\log 1/\delta} \leq 2\sqrt{\frac{\log 1/\delta}{\ell d^2}}.$$

By picking  $\delta = 1/(64 \cdot d^{10} \cdot |\mathcal{G}|^2)$ , we get with probability at least  $1 - 1/(64d^{10})$ :

$$|f(x) - f(y)| \leq C\sqrt{\frac{\log d}{d}}.$$

For  $x \notin \mathcal{G}$ , let  $\tilde{x} = \arg \min_{y \in \mathcal{G}} \|x - y\|$  and we have:

$$\|f_\ell(x) - f_\ell(\tilde{x})\| \leq 4^\ell \varepsilon \leq \frac{1}{2^\ell}$$

and

$$\|f(x) - f(\tilde{x})\| \leq \frac{1}{2^\ell} \cdot \|W_{\ell+1}\|.$$

On the event,  $\|W_{\ell+1}\| \leq 2\sqrt{k}$ , this yields the conclusion:

$$\forall x, y \in \mathbb{S}^{d-1} : |f(x) - f(y)| \leq C\sqrt{\frac{\log d}{d}}$$

with probability at least  $1 - 1/d^{10}$ .

Finally, we have by the anti-concentration of Gaussians that with probability at least 0.95 for fixed  $x \in \mathcal{G}$ :

$$|f(x)| \geq 0.05.$$

A union bound and the above two displays concludes the proof.  $\square$

## C Miscellaneous Results

We restate standard results used in our analysis. We start with a fact on Gaussian random matrices.

**Lemma C.1.** *Let  $A$  be an  $m \times n$  random matrix with  $A_{i,j} \stackrel{i.i.d}{\sim} \mathcal{N}(0, 1)$ . Then, we have that:*

$$\|A\| \leq 3(\sqrt{m} + \sqrt{n} + \sqrt{\log 1/\delta})$$

with probability at least  $1 - \delta$ .

*Proof.* We may assume without loss of generality that  $m \leq n$ . Let  $\mathcal{G}$  be a  $1/3$ -grid of  $\mathbb{S}^{m-1}$  and we have:

$$\|A\| = \max_{u \in \mathbb{S}^{m-1}} \|u^\top A\| = \max_{u \in \mathbb{S}^{m-1}} \|(u - \tilde{u})^\top A + \tilde{u}^\top A\| \leq \frac{\|A\|}{3} + \max_{u \in \mathcal{G}} \|u^\top A\|$$

where  $\tilde{u} = \arg \min_{v \in \mathcal{G}} \|u - v\|$  which implies:

$$\|A\| \leq \frac{3}{2} \cdot \max_{u \in \mathcal{G}} \|u^\top A\|.$$

Note, we may assume that  $|\mathcal{G}| \leq (10)^m$  and  $u^\top A \sim \mathcal{N}(0, I)$  and we have from [Theorem C.3](#):

$$\forall \|u\| = 1 : \mathbb{P} \left\{ \|u^\top A\| \leq \sqrt{n} + \sqrt{\log 1/\delta} \right\} \geq 1 - \delta'.$$

Setting  $\delta' = \delta/\mathcal{G}$  and a union bound yields:

$$\max_{u \in \mathcal{G}} \|u^\top A\| \leq \sqrt{n} + \sqrt{m \log(10)} + \sqrt{\log 1/\delta} \leq 2(\sqrt{m} + \sqrt{n} + \sqrt{\log 1/\delta})$$

with probability at least  $1 - \delta$  which yields the lemma from the previous discussion.  $\square$

We also recall Bernstein's Inequality used frequently throughout our analysis.

**Theorem C.2.** *[BLM13, Theorem 2.1] Let  $X_1, \dots, X_n$  be  $n$  independent real-valued random variables. Assume there exist positive numbers  $\nu$  and  $c$  such that:*

$$\sum_{i=1}^n \mathbb{E}[X_i^2] \leq \nu \text{ and } \sum_{i=1}^n \mathbb{E}[|X_i|^q] \leq \frac{q!}{2} \nu c^{q-2} \text{ for all } q \geq 3.$$

Then, we have:

$$\mathbb{P} \left\{ \sum_{i=1}^n (X_i - \mathbb{E}[X_i]) \geq \sqrt{2\nu t} + ct \right\} \leq e^{-t}.$$

We additionally recall the Tsirelson-Ibragimov-Sudakov inequality.

**Theorem C.3.** *[BLM13, Theorem 5.6] Let  $X = (X_1, \dots, X_n)$  be a vector of  $n$  independent standard normal random variables. Let  $g : \mathbb{R}^n \rightarrow \mathbb{R}$  denote a  $L$ -Lipschitz function. Then, we have:*

$$\forall t \geq 0 : \mathbb{P} \{g(X) - \mathbb{E}g(X) \geq t\} \leq \exp \left\{ -\frac{t^2}{2L^2} \right\}.$$

We frequently use this inequality with the 1-Lipschitz functions  $g(X) = \pm \|X\|$ . Combining [Theorem C.3](#) with Gautschi's inequality, which implies  $\mathbb{E}\|X\|/\sqrt{n} \in (\sqrt{1-1/n}, \sqrt{1+1/n}) \subset (1-1/\sqrt{n}, 1+1/\sqrt{n})$ , gives, for  $\delta \in (0, 1)$ ,

$$\mathbb{P} \left\{ \|X\| \geq \sqrt{n} + \sqrt{2 \ln 1/\delta} + 1 \right\} \leq \delta, \quad \mathbb{P} \left\{ \|X\| \leq \sqrt{n} - \sqrt{2 \ln 1/\delta} - 1 \right\} \leq \delta.$$

This immediately implies the following corollary.

**Corollary C.4.** *For  $\delta \in (0, 1)$ , a matrix  $W_i \in \mathbb{R}^{d_i \times d_{i-1}}$  with independent  $\mathcal{N}(0, 1/d_{i-1})$  entries, and vectors  $u \in \mathbb{R}^{d_i}$  and  $v \in \mathbb{R}^{d_{i-1}}$ , both of the following events have probability at least  $1 - \delta$ :*

$$\left| \|u^\top W_i\| - \|u\| \right| \leq \|u\| \frac{\sqrt{2 \ln(2/\delta)} + 1}{\sqrt{d_{i-1}}}, \quad \left| \|W_i v\| - \|v\| \right| \leq \|v\| \frac{\sqrt{2 \ln(2/\delta)} + 1}{\sqrt{d_{i-1}}}.$$

We prove another simple lemma.

**Lemma C.5.** *Let  $x, y \in \mathbb{R}^d$ ,  $R, r \in \mathbb{R}_+$  be such that  $\|x\| \geq R > 0$  and  $\|x - y\| \leq r$  with  $R \geq r$ . Then, we have:*

$$\mathbb{P}_{w \sim \mathcal{N}(0, I)} \{ \text{sign}(w^\top x) \neq \text{sgn}(w^\top y) \} \leq \frac{3r}{R} \cdot \sqrt{\log R/r}.$$

*Proof.* We have that  $X := w^\top x \sim \mathcal{N}(0, \|x\|^2)$  and  $Z := w^\top (y - x) \sim \mathcal{N}(0, \|y - x\|^2)$ . Hence, we have:

$$\mathbb{P} \{ \text{sign}(w^\top x) \neq \text{sign}(w^\top y) \} \leq \mathbb{P} \{ |Z| \geq |X| \}.$$

We have for any  $\delta > 0$ :

$$\mathbb{P} \left\{ |Z| \geq 2r\sqrt{\log 1/\delta} \right\} \leq \delta \text{ and } \mathbb{P} \left\{ |X| \leq 2r\sqrt{\log 1/\delta} \right\} \leq \frac{4r\sqrt{\log 1/\delta}}{\sqrt{2\pi}R} \leq 2 \cdot \frac{r}{R} \cdot \sqrt{\log 1/\delta}.$$

By a union bound, we have:

$$\forall \delta > 0 : \mathbb{P} \{ |Z| \geq |X| \} \leq \delta + \frac{2r}{R} \cdot \sqrt{\log 1/\delta}.$$

By setting  $\delta = r/R$ , we get the conclusion of the lemma.  $\square$

**Lemma C.6.** *Let  $d \in \mathbb{N}$  with  $d \geq 40$ . Then, we have for all  $k \in \mathbb{N}$ ,  $r \geq 0$ ,  $t \geq 0$ ,  $M \in \mathbb{R}^{k \times d}$ :*

$$\mathbb{P}_{Z \sim \text{Unif}(r\mathbb{S}^{d-1})} \{ \|MZ\| \geq t \} \leq 2\mathbb{P}_{Y \sim \mathcal{N}(0, 2r^2 I/d)} \{ \|MY\| \geq t \}.$$

*Proof.* The proof follows from the following manipulations:

$$\begin{aligned} \mathbb{P}_{Y \sim \mathcal{N}(0, 2r^2 I/d)} \{ \|MY\| \geq t \} &= \mathbb{E}_{Y \sim \mathcal{N}(0, 2r^2 I/d)} [ \mathbf{1} \{ \|MY\| \geq t \} ] \\ &\geq \mathbb{E}_{Y \sim \mathcal{N}(0, 2r^2 I/d)} [ \mathbf{1} \{ \|MY\| \geq t \} \mathbf{1} \{ \|Y\| \geq r \} ] \\ &= \mathbb{E}_{Y \sim \mathcal{N}(0, 2r^2 I/d)} \left[ \mathbf{1} \left\{ \frac{\|MY\|}{\|Y\|} \geq \frac{t}{\|Y\|} \right\} \mathbf{1} \{ \|Y\| \geq r \} \right] \\ &\geq \mathbb{E}_{Y \sim \mathcal{N}(0, 2r^2 I/d)} \left[ \mathbf{1} \left\{ \frac{\|MY\|}{\|Y\|} \geq \frac{t}{r} \right\} \mathbf{1} \{ \|Y\| \geq r \} \right] \\ &\geq \mathbb{E}_{Y \sim \mathcal{N}(0, 2r^2 I/d)} [ \mathbb{P}_{Z \sim \text{Unif}(r\mathbb{S}^{d-1})} \{ \|MZ\| \geq t \} \mathbf{1} \{ \|Y\| \geq r \} ] \\ &\geq \mathbb{P}_{Z \sim \text{Unif}(r\mathbb{S}^{d-1})} \{ \|MZ\| \geq t \} \cdot \mathbb{P}_{Y \sim \mathcal{N}(0, 2r^2 I/d)} \{ \|Y\| \geq t \} \\ &\geq \frac{1}{2} \cdot \mathbb{P}_{Z \sim \text{Unif}(r\mathbb{S}^{d-1})} \{ \|MZ\| \geq t \} \end{aligned}$$

concluding the proof of the lemma.  $\square$

## D Empirical Evaluations

In this section, we provide some experimental evidence for the role played by bottleneck layers in restricting the complexity of intermediate representations and in determining the size of the neighborhood around  $x$  where the network retains its near-linear behavior. Here, we restrict ourselves to 1-hidden layer neural networks with the single hidden layer serving as a bottleneck layer. For a 1-hidden layer network with  $k$  hidden units parameterized as  $f(x) := w_2^\top \text{ReLU}(W_1 x)$  where  $w_2 \in \mathbb{R}^k$  and  $W_1 \in \mathbb{R}^{k \times d}$ , we compute the smallest radius,  $r > 0$ , such that for any possible value of  $\nabla f(y)$  for arbitrary  $y \in \mathbb{R}^d$ , there exists  $w \in \mathbb{R}^d$  with  $\|w - x\| \leq r$  and  $\nabla f(w) = \nabla f(y)$ :

$$r := \inf \{ t > 0 : \forall y \in \mathbb{R}^d, \exists w \in \mathbb{R}^d \text{ s.t. } \nabla f(w) = \nabla f(y) \text{ and } \|w - x\| \leq t \}$$

A straightforward upper bound on  $r$  is the objective value of the following optimization problem:

$$\begin{aligned} &\min_{\delta} \|\delta\| \\ &\text{s.t. } W_1(x - \delta) = 0. \end{aligned} \quad (\text{Min-Pert})$$

whose solution is  $\|W_1^\dagger W_1 x\|$  where  $W_1^\dagger$  denotes the pseudo-inverse of  $W_1$ . Noting the scale invariance of the ReLU activation function and following [Theorem 1.1](#), we pick  $(w_2)_i, (W_1)_{i,j} \stackrel{iid}{\sim} \mathcal{N}(0, 1)$ . In this setting,  $\|W_1^\dagger W_1 x\|$  can be shown to concentrate around  $\|x\| \cdot \sqrt{k/d}$  with high probability. Hence, as the width the hidden layer increases, so does complexity of the intermediate representation and consequently, the radius within which the network approximates a linear function around  $x$  as the perturbation must “undo” a larger more complex representation. Note that despite the fact that the constraints on  $r$  are significantly stronger than the region of near linear behavior guaranteed by [Theorem 1.1](#), this bound is tight up to logarithmic factors. Indeed, [Theorem 1.1](#) guarantees near-linear behavior within a radius of  $\|x\| \cdot \tilde{\Omega}(\sqrt{k/d})$  of  $x$ .

In [Fig. 1](#), we track the optimal length of the perturbation  $\delta^*$  relative to the length of the input for a network with input dimension 10000; formally, we plot the value of  $\|\delta^*\|/\|x\|$  averaged over 10 runs for the number of hidden units ranging over  $\{10, 50, 100, 500, 1000, 5000, 10000\}$ . The figure shows that as the number of hidden units are increased, the size of the perturbation required to completely distort the gradient increases. The size of this perturbation is closely related to the complexity of the intermediate representation of the input at the bottleneck layer consistent with [Theorem 1.1](#).

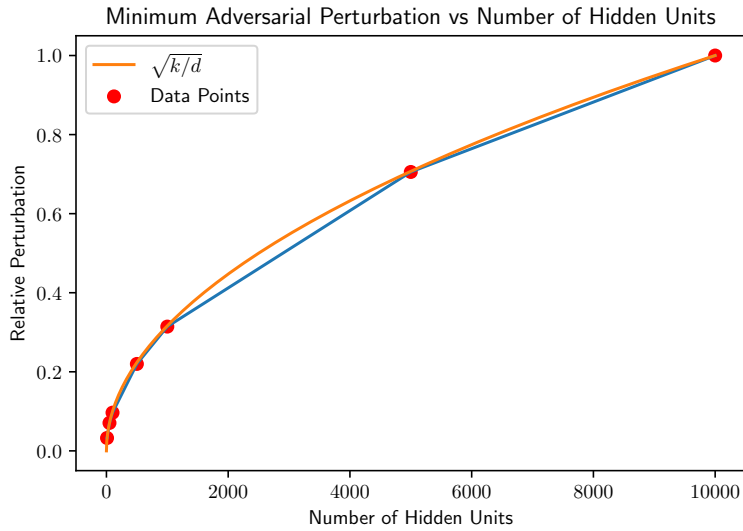


Figure 1: The average value of the optimal solution to [Min-Pert](#) as a function of the number of hidden units of a random one-hidden layer neural network with 10000 hidden units averaged over 10 runs. Observe that the data points are in close agreement with the theoretically predicted ratio of  $\sqrt{k/d}$ .