
Network size and weights size for memorization with two-layers neural networks

Sébastien Bubeck
Microsoft Research

Ronen Eldan*
Weizmann Institute

Yin Tat Lee
University of Washington
& Microsoft Research

Dan Mikulincer
Weizmann Institute

Abstract

In 1988, Eric B. Baum showed that two-layers neural networks with threshold activation function can perfectly memorize the binary labels of n points in general position in \mathbb{R}^d using only $\lceil n/d \rceil$ neurons. We observe that with ReLU networks, using four times as many neurons one can fit arbitrary real labels. Moreover, for approximate memorization up to error ε , the neural tangent kernel can also memorize with only $O\left(\frac{n}{d} \cdot \log(1/\varepsilon)\right)$ neurons (assuming that the data is well dispersed too). We show however that these constructions give rise to networks where the *magnitude* of the neurons' weights are far from optimal. In contrast we propose a new training procedure for ReLU networks, based on *complex* (as opposed to *real*) recombination of the neurons, for which we show approximate memorization with both $O\left(\frac{n}{d} \cdot \frac{\log(1/\varepsilon)}{\varepsilon}\right)$ neurons, as well as nearly-optimal size of the weights.

1 Introduction

We study two-layers neural networks in \mathbb{R}^d with k neurons and non-linearity $\psi : \mathbb{R} \rightarrow \mathbb{R}$. These are functions of the form:

$$x \mapsto \sum_{\ell=1}^k a_{\ell} \psi(w_{\ell} \cdot x + b_{\ell}), \quad (1)$$

with $a_{\ell}, b_{\ell} \in \mathbb{R}$ and $w_{\ell} \in \mathbb{R}^d$ for any $\ell \in [k]$. We are mostly concerned with the Rectified Linear Unit non-linearity, namely $\text{ReLU}(t) = \max(0, t)$, in which case wlog one can restrict the recombination weights (a_{ℓ}) to be in $\{-1, 1\}$ (this holds more generally for positively homogeneous non-linearities). We denote by $\mathcal{F}_k(\psi)$ the set of functions of the form (1). Under mild conditions on ψ (namely that it is not a polynomial), such neural networks are *universal*, in the sense that for k large enough they can approximate any continuous function [Cybenko, 1989, Leshno et al., 1993].

In this paper we are interested in approximating a target function on a *finite* data set. This is also called the *memorization* problem. Specifically, fix a data set $(x_i, y_i)_{i \in [n]} \in (\mathbb{R}^d \times \mathbb{R})^n$ and an approximation error $\varepsilon > 0$. We denote $\mathbf{y} = (y_1, \dots, y_n)$, and for a function $f : \mathbb{R}^d \rightarrow \mathbb{R}$ we write $\mathbf{f} = (f(x_1), \dots, f(x_n))$. The main question concerning the memorization capabilities of $\mathcal{F}_k(\psi)$ is as follows: How large should be k so that there exists $f \in \mathcal{F}_k(\psi)$ such that $\|\mathbf{f} - \mathbf{y}\|^2 \leq \varepsilon \|\mathbf{y}\|^2$ (where $\|\cdot\|$ denotes the Euclidean norm)? A simple consequence of universality of neural networks is that $k \geq n$ is sufficient (see Proposition 2). In fact (as was already observed by Baum [1988] for threshold ψ and binary labels, see Proposition 3) much more compact representations can be achieved by leveraging the high-dimensionality of the data. Namely we prove that for $\psi = \text{ReLU}$ and a data set in general position (i.e., any hyperplane contains at most d points), one only needs $k \geq 4 \cdot \lceil \frac{n}{d} \rceil$ to memorize the data perfectly, see Proposition 4. The size $k \approx n/d$ is clearly optimal,

*This work was partly done while R. Eldan and D. Mikulincer were visiting Microsoft Research.

by a simple parameter counting argument. We call the construction given in Proposition 4 a *Baum network*, and as we shall see it is of a certain combinatorial flavor. In addition we also prove that such memorization can in fact essentially be achieved in a kernel regime (with a bit more assumptions on the data): we prove in Theorem 2 that for $k = \Omega\left(\frac{n}{d} \log(1/\varepsilon)\right)$ one can obtain approximate memorization with the Neural Tangent Kernel [Jacot et al., 2018], and we call the corresponding construction the *NTK network*. Specifically, the kernel we consider is,

$$\mathbb{E} [\nabla_w \psi(w \cdot x) \cdot \nabla_w \psi(w \cdot y)] = \mathbb{E} [(x \cdot y) \psi'(w \cdot x) \psi'(w \cdot y)],$$

where ∇_w is the gradient with respect to the w variable and the expectation is taken over a random initialization of w .

Measuring regularity via total weight. One is often interested in fitting the data using functions which satisfy certain regularity properties. The main notion of regularity in which we are interested is the *total weight*, defined as follows: For a function $f : \mathbb{R}^d \rightarrow \mathbb{R}$ of the form (1), we define

$$\mathbf{W}(f) := \sum_{\ell=1}^k |a_\ell| \sqrt{\|w_\ell\|^2 + b_\ell^2}.$$

This definition is widely used in the literature, see Section 2 for a discussion and references. Notably, it was shown in Bartlett [1998] that this measure of complexity is better associated with the network's generalization ability compared to the size of the network. We will be interested in constructions which have both a small number of neurons and a small total weight.

Our main contribution: The complex network. As we will see below, both the Baum network and the NTK networks have sub-optimal total weight. The main technical contribution of our paper is a third type of construction, which we call the *harmonic network*, that under the same assumptions on the data as for the NTK network, has both near-optimal memorization size and near-optimal total weight:

Theorem 1 (Informal). *Suppose that $n \leq \text{poly}(d)$. Let $x_1, \dots, x_N \in \mathbb{S}^{d-1}$ such that*

$$|x_i \cdot x_j| = \tilde{O}\left(\frac{1}{\sqrt{d}}\right).$$

For every $\varepsilon > 0$ and every choice of labels $(y_i)_{i=1}^n$ such that $|y_i| = O(1)$ for all i , there exist $k = \tilde{O}\left(\frac{n}{d\varepsilon}\right)$ and $f \in \mathcal{F}_k(\psi)$ such that

$$\frac{1}{n} \sum_{i=1}^n \min\left(\left(y_i - f(x_i)\right)^2, 1\right) \leq \varepsilon$$

and such that $\mathbf{W}(f) = \tilde{O}(\sqrt{n})$.

We show below in Proposition 1 that for random data one necessarily has $\mathbf{W}(f) = \tilde{\Omega}(\sqrt{n})$, thus proving that the harmonic network has near-optimal total weight. Moreover we also argue in the corresponding sections that the Baum and NTK networks have total weight at least $n\sqrt{n}$ on random data, thus being far from optimal.

An iterative construction. Both the NTK network and the harmonic network will be built by iteratively adding up small numbers of neurons. This procedure, akin to boosting, is justified by the following lemma. It shows that to build a large memorizing network it suffices to be able to build a small network f whose scalar product with the data $\mathbf{f} \cdot \mathbf{y}$ is comparable to its variance $\|\mathbf{f}\|^2$:

Lemma 1 *Fix $(x_i)_{i=1}^n$. Suppose that there are $m \in \mathbb{N}$ and $\alpha, \beta > 0$ such that the following holds: For any choice of $(y_i)_{i=1}^n$, there exists $f \in \mathcal{F}_m(\psi)$ with $\mathbf{y} \cdot \mathbf{f} \geq \alpha \|\mathbf{y}\|^2$ and $\|\mathbf{f}\|^2 \leq \beta \|\mathbf{y}\|^2$. Then for all $\varepsilon > 0$, there exists $g \in \mathcal{F}_{mk}(\psi)$ such that*

$$\|\mathbf{g} - \mathbf{y}\|^2 \leq \varepsilon \|\mathbf{y}\|^2$$

with

$$k \leq \frac{\beta}{\alpha^2} \log(1/\varepsilon).$$

Moreover, if the above holds with $\mathbf{W}(f) \leq \omega$, then $\mathbf{W}(g) \leq \frac{\omega}{\alpha} \log(1/\varepsilon)$.

Proof. Denote $\eta = \frac{\alpha}{\beta}$ and $\mathbf{r}_1 = \mathbf{y}$. Then, there exists $f_1 \in \mathcal{F}_m(\psi)$, such that

$$\begin{aligned} \|\eta \mathbf{f}_1 - \mathbf{r}_1\|^2 &= \|\mathbf{r}_1\|^2 - 2\eta \mathbf{y} \cdot \mathbf{f}_1 + \eta^2 \|\mathbf{f}_1\|^2 \leq \|\mathbf{r}_1\|^2 \left(1 - 2\frac{\alpha^2}{\beta} + \frac{\alpha^2}{\beta}\right) \\ &\leq \|\mathbf{r}_1\|^2 \left(1 - \frac{\alpha^2}{\beta}\right) = \|\mathbf{y}\|^2 \left(1 - \frac{\alpha^2}{\beta}\right) \end{aligned}$$

The result is obtained by iterating the above inequality with $\mathbf{r}_i = \mathbf{y} - \eta \sum_{j=1}^{i-1} \mathbf{f}_j$ taken as the residuals. By induction, if we set $g = \eta \sum_{j=1}^k f_j$, we get

$$\|\mathbf{g} - \mathbf{y}\| = \|\eta \mathbf{f}_k - \mathbf{r}_k\| \leq \|\mathbf{r}_k\|^2 \left(1 - \frac{\alpha^2}{\beta}\right) = \|\mathbf{y}\|^2 \left(1 - \frac{\alpha^2}{\beta}\right)^k.$$

□

In both the NTK and harmonic constructions, the function f will have the largest possible correlation with the data set attainable for a network of constant size. However, the harmonic network will have the extra advantage that the function f will be composed of a single neuron whose weight is the smallest one attainable. Thus, the harmonic network will enjoy both the smallest possible number of neurons and smallest possible total weight (up to logarithmic factors). Note however that the dependency on ε is worse for the harmonic network, which is technically due to a constant order term in the variance which we do not know how to remove.

We conclude the introduction by showing that a total weight of $\Omega(\sqrt{n})$ is necessary for approximate memorization. Just like for the upper bound, it turns out that it is sufficient to consider how well can one correlate a single neuron. Namely the proof boils down to showing that a single neuron cannot correlate well with random data sets.

Proposition 1 *There exists a data set $(x_i, y_i)_{i \in [n]} \in (\mathbb{S}^{d-1} \times \{-1, 1\})^n$ such that for every function f of the form (1) with ψ L -Lipschitz and which satisfies $\|\mathbf{f} - \mathbf{y}\|^2 \leq \frac{1}{2}\|\mathbf{y}\|^2$, it holds that $\mathbf{W}(f) \geq \frac{\sqrt{n}}{8L}$.*

Proof. We have

$$\frac{1}{2}\|\mathbf{y}\|^2 \geq \|\mathbf{f} - \mathbf{y}\|^2 \geq \|\mathbf{y}\|^2 - 2\mathbf{f} \cdot \mathbf{y} \Rightarrow \mathbf{f} \cdot \mathbf{y} \geq \frac{1}{4}\|\mathbf{y}\|^2,$$

that is

$$\sum_{\ell=1}^k \sum_{i=1}^n y_i a_{\ell} \psi(w_{\ell} \cdot x_i - b_{\ell}) \geq \frac{n}{4},$$

which implies:

$$\max_{w, b} \sum_{i=1}^n y_i \frac{\psi(w \cdot x_i - b)}{\sqrt{\|w\|^2 + b^2}} \geq \frac{n}{4\mathbf{W}(f)}.$$

Now let us assume that y_i are ± 1 uniformly at random (i.e., Rademacher random variables), and thus by Talagrand's contraction lemma for the Rademacher complexity (see [Lemma 26.9, [Shalev-Shwartz and Ben-David \[2014\]](#)]) we have:

$$\begin{aligned} \mathbb{E} \max_{w, b} \sum_{i=1}^n y_i \frac{\psi(w \cdot x_i - b)}{\sqrt{\|w\|^2 + b^2}} &\leq L \cdot \mathbb{E} \max_{w, b} \sum_{i=1}^n y_i \frac{w \cdot x_i - b}{\sqrt{\|w\|^2 + b^2}} \\ &\leq L \cdot \mathbb{E} \sqrt{\left\| \sum_{i=1}^n y_i x_i \right\|^2 + n} \leq 2L\sqrt{n}, \end{aligned}$$

and thus $\mathbf{W}(f) \geq \frac{\sqrt{n}}{8L}$.

□

2 Related works

Exact memorization. The observation that n neurons are sufficient for memorization with essentially arbitrary non-linearity was already made in [Bach, 2017] (using Carathéodory’s theorem), and before that a slightly weaker bound with $n + 1$ neurons was already observed in [Bengio et al., 2006] (or more recently $2n + d$ in [Zhang et al., 2017]). The contribution of Proposition 2 is to show that this statement of exactly n neurons follows in fact from elementary linear algebra.

As already mentioned above, Baum [1988] proved that for threshold non-linearity and binary labels one can obtain a much better bound of n/d neurons for memorization, as long as the data is in general position. This was generalized to the ReLU non-linearity (but still binary labels) in Yun et al. [2019] (we note that this paper also considers some questions around memorization capabilities of deeper networks). Our modest contribution here is to generalize this to arbitrary real labels, see Proposition 4.

Gradient-based memorization. A different line of works on memorization studies whether it can be achieved via gradient-based optimization on various neural network architectures. The literature here is very large, but early results with minimal assumptions include Li and Liang [2018], Soltanolkotabi et al. [2018] which were notably generalized in [Allen-Zhu et al., 2019, Du et al., 2019]. Crucially these works leverage very large overparametrization, i.e., the number of neurons is a large polynomial in the number of data points. For a critique of this large overparametrization regime see [Chizat et al., 2019, Ghorbani et al., 2019, Yehudai and Shamir, 2019], and for a different approach based on a certain scaling limit of stochastic gradient descent for sufficiently overparametrized networks see [Chizat and Bach, 2018, Mei et al., 2018]. More recently the amount of overparametrization needed was improved to a small polynomial dependency in n and d in [Kawaguchi and Huang, 2019, Oymak and Soltanolkotabi, 2019, Song and Yang, 2019]. In the *random features* regime, Bresler and Nagaraj [2020] have also considered an iterative construction procedure for memorization. This is somewhat different than our approach, in which the iterative procedure updates the w_j ’s, and a much smaller number of neurons is needed as a result. Finally, very recently Amit Daniely [Daniely, 2019, 2020] showed that gradient descent already works in the optimal regime of $k = \tilde{O}(n/d)$, at least for random data (and random labels). This result is closely related to our analysis of the NTK network in Section 4. Minor distinctions are that we allow for arbitrary labels, and we take a “boosting approach” where neurons are added one by one (although we do not believe that this is an essential difference).

Total weight complexity. It is well-known since Bartlett [1998] that the total weight of a two-layers neural network is a finer measure of complexity than the number of neurons to control its generalization (see Neyshabur et al. [2015] and Arora et al. [2019] for more recent discussions on this, as well as Bartlett et al. [2017] for other notions of norms for deeper networks). Of course the bound $\mathbf{W} = \tilde{O}(\sqrt{n})$ proved here leads to vacuous generalization performance, as is necessary since the Harmonic network can memorize completely random data (for which no generalization is possible). It would be interesting to see if the weight of the Harmonic network can be smaller for more structured data, particularly given the context raised by the work [Zhang et al., 2017] (where it was observed that SGD on deep networks will memorize arbitrary data, hence the question of where does the seeming generalization capabilities of those networks come from). We note the recent work [Ji and Telgarsky, 2020] which proves for example that polylogarithmic size network is possible for memorization under a certain margin condition. Finally we also note that the effect in function space of bounding \mathbf{W} has been recently studied in Ongie et al. [2020], Savarese et al. [2019].

Complex weights. It is quite natural to consider neural networks with complex weights. Indeed, as was already observed by Barron [Barron, 1993], the Fourier transform $f(x) = \int \hat{f}(\omega) \exp(i\omega \cdot x) d\omega$ exactly gives a representation of f as a two-layers neural network with the non-linearity $\psi(t) = \exp(it)$. More recently, it was noted in Andoni et al. [2014] that randomly perturbing a neuron with *complex weights* is potentially more beneficial than doing a mere real perturbation. We make a similar observation in Section 5 for the construction of the Harmonic network, where we show that complex perturbations allow to deal particularly easily with higher order terms in some key Taylor expansion. Moreover we also note that Andoni et al. [2014] considers non-linearity built from Hermite polynomials, which shall be a key step for us too in the construction of the Harmonic

network (the use of Hermite polynomials in the context of learning theory goes back to [Kalai et al., 2008]).

While orthogonal to our considerations here, we also note the work of Fefferman [Fefferman, 1994], where he used the analytical continuation of a (real) neural network to prove a certain uniqueness property (essentially that two networks with the same output must have the same weights up to some obvious symmetries and obvious counter-examples).

3 Elementary results on memorization

In this section we give a few examples of elementary conditions on k , ψ and the data set so that one can find $f \in \mathcal{F}_k(\psi)$ with $\mathbf{f} = \mathbf{y}$ (i.e., exact memorization). We prove three results: (i) $k \geq n$ suffices for any non-polynomial ψ , (ii) $k \geq \frac{n}{d} + 3$ with $\psi(t) = \mathbb{1}\{t \geq 0\}$ suffices for binary labels with data in general position (this is exactly Baum [1988]’s result), and (iii) $k \geq 4 \cdot \lceil \frac{n}{d} \rceil$ with $\psi = \text{ReLU}$ suffices for data in general position and arbitrary labels.

We start with the basic linear algebraic observation that having a number of neurons larger than the size of the data set is always sufficient for perfect memorization:

Proposition 2 *Assuming that ψ is not a polynomial, there exists $f \in \mathcal{F}_n(\psi)$ such that $\mathbf{f} = \mathbf{y}$.*

Proof. Note that the set of functions of the form (1) (with arbitrary k) corresponds to the vector space V spanned by the functions $\psi_{w,b} : x \mapsto \psi(w \cdot x + b)$. Consider the linear operator $\Psi : V \rightarrow \mathbb{R}^n$ that corresponds to the evaluation on the data points (x_i) (i.e., $\Psi(f) = (f(x_i))_{i \in [n]}$). Since ψ is not a polynomial, the image of Ψ is $\text{Im}(\Psi) = \mathbb{R}^n$. Moreover $\text{Im}(\Psi)$ is spanned by the set of vectors $\Psi(\psi_{w,b})$ for $w \in \mathbb{R}^d, b \in \mathbb{R}$. Now, since $\dim(\text{Im}(\Psi)) = n$, one can extract a subset of n such vectors with the same span, that is there exists $w_1, b_1, \dots, w_n, b_n$ such that

$$\text{span}(\Psi(\psi_{w_1, b_1}), \dots, \Psi(\psi_{w_n, b_n})) = \mathbb{R}^n,$$

which concludes the proof. \square

In [Baum, 1988] it is observed that one can dramatically reduce the number of neurons for high-dimensional data:

Proposition 3 *Fix $\psi(t) = \mathbb{1}\{t \geq 0\}$. Let $(x_i)_{i \in [n]}$ be in general position in \mathbb{R}^d (i.e., any hyperplane contains at most d points), and assume binary labels, i.e., $y_i \in \{0, 1\}$. Then there exists $f \in \mathcal{F}_{\frac{n}{d}+3}(\psi)$ such that $\mathbf{f} = \mathbf{y}$.*

Proof. Baum [1988] builds a network iteratively as follows. Pick d points with label 1, say x_1, \dots, x_d , and let $H = \{x : u \cdot x = b\}$ be a hyperplane containing those points and no other points in the data, i.e., $x_i \notin H$ for any $i > d$. With two neurons (i.e., $f \in \mathcal{F}_2(\psi)$) one can build the indicator of a small neighborhood of H , namely $f(x) = \psi(u \cdot x - (b - \tau)) - \psi(u \cdot x - (b + \tau))$ with τ small enough, so that $f(x_i) = 1$ for $i \leq d$ and $f(x_i) = 0$ for $i > d$. Assuming that the label 1 is the minority (which is without loss of generality up to one additional neuron), one thus needs at most $2^{\lceil \frac{n}{2d} \rceil}$ neurons to perfectly memorize the data. \square

We now extend Proposition 3 to the ReLU non-linearity and arbitrary real labels. To do so we introduce the *derivative neuron* of ψ defined by:

$$f_{\delta, u, v, b} : x \mapsto \frac{\psi((u + \delta v) \cdot x - b) - \psi(u \cdot x - b)}{\delta}, \quad (2)$$

with $\delta \in \mathbb{R}$ and $u, v \in \mathbb{R}^d$. As δ tends to 0, this function is equal to

$$f_{u, v, b}(x) = \psi'(u \cdot x - b)v \cdot x \quad (3)$$

for any x such that ψ is differentiable at $u \cdot x - b$. In fact, for the ReLU one has for any x such that $u \cdot x \neq b$ that $f_{\delta, u, v, b}(x) = f_{u, v, b}(x)$ for δ small enough (this is because the ReLU is piecewise

linear). We will always take δ small enough and u such that $f_{\delta,u,v,b}(x_i) = f_{u,v,b}(x_i)$ for any $i \in [n]$, for example by taking

$$\delta = \frac{1}{2} \min_{i \in [n]} \frac{|u \cdot x_i - b|}{|v \cdot x_i|}. \quad (4)$$

Thus, as far as memorization is concerned, we can assume that $f_{u,v,b} \in \mathcal{F}_2(\text{ReLU})$. With this observation it is now trivial to prove the following extension of Baum’s result:

Proposition 4 *Let $(x_i)_{i \in [n]}$ be in general position in \mathbb{R}^d (i.e., any hyperplane contains at most d points). Then there exists $f \in \mathcal{F}_{4 \cdot \frac{n}{d} \gamma}(\text{ReLU})$ such that $\mathbf{f} = \mathbf{y}$.*

Proof. Pick an arbitrary set of d points, say $(x_i)_{i \leq d}$, and let $H = \{x : u \cdot x = b\}$ be a hyperplane containing those points and no other points in the data, i.e., $x_i \notin H$ for any $i > d$. With four neurons one can build the function $f = f_{u,v,b-\tau} - f_{u,v,b+\tau}$ with τ small enough so that $f(x_i) = x_i \cdot v$ for $i \leq d$ and $f(x_i) = 0$ for $i > d$. It only remains to pick v such that $v \cdot x_i = y_i$ for any $i \leq d$, which we can do since the matrix given by $(x_i)_{i \leq d}$ is full rank (by the general position assumption). \square

4 The NTK network

The constructions in Section 3 are based on a very careful set of weights that depend on the entire dataset. Here we show that essentially the same results can be obtained in the *neural tangent kernel* regime. That is, we take pair of neurons as given in (2) (which corresponds in fact to (3) since we will take δ to be small, we will also restrict to $b = 0$), and crucially we will also have that the “main weight” u will be chosen at random from a standard Gaussian, and only the “small perturbation” v will be chosen as a function of the dataset. The guarantee we obtain is slightly weaker than in Proposition 4: we have a $\log(1/\varepsilon)$ overhead in the number of neurons, and moreover we also need to assume that the data is “well-spread”. Specifically we consider the following notion of “generic data”:

Definition 1 *We say that $(x_i)_{i \in [n]}$ are (γ, ω) -generic (with $\gamma \in (\frac{1}{2n}, 1)$ and $\omega > 0$) if $\|x_i\| \geq 1$ for all $i \in [n]$, $\frac{1}{n} \sum_{i=1}^n x_i x_i^\top \preceq \frac{\omega}{d} \cdot \mathbf{I}_d$, and $|x_i \cdot x_j| \leq \gamma \cdot \|x_i\| \cdot \|x_j\|$ for all $i \neq j$.*

In the following we fix such a (γ, ω) -generic data set. Note that i.i.d. points on the sphere are $\left(O\left(\sqrt{\frac{\log(n)}{d}}\right), O(1)\right)$ -generic. We now formulate our main theorem concerning the NTK network.

Theorem 2 *There exists $f \in \mathcal{F}_k(\text{ReLU})$, produced in the NTK regime (see Theorem 3 below for more details) with $\mathbb{E}[\|\mathbf{f} - \mathbf{y}\|^2] \leq \varepsilon \|\mathbf{y}\|^2$ (the expectation is over the random initialization of the “main weights”) provided that*

$$k \cdot d \geq 20\omega \cdot n \log(1/\varepsilon) \cdot \frac{\log(2n)}{\log(1/\gamma)}. \quad (5)$$

In light of Lemma 1, it will be enough to produce a width-2 network, $f \in \mathcal{F}_2(\text{ReLU})$, whose correlation with the data set is large.

Theorem 3 *There exists $f \in \mathcal{F}_2(\text{ReLU})$ with*

$$\mathbf{y} \cdot \mathbf{f} \geq \frac{1}{10} \cdot \sqrt{\frac{\log(1/\gamma)}{\log(2n)}} \cdot \|\mathbf{y}\|^2, \quad (6)$$

and

$$\|\mathbf{f}\|^2 \leq \frac{\omega \cdot n}{d} \|\mathbf{y}\|^2. \quad (7)$$

In fact, one can take the construction (2) with:

$$u \sim \mathcal{N}(0, \mathbf{I}_d), \quad v = \sum_{i: u \cdot x_i \geq 0} y_i x_i, \quad \delta = \frac{1}{2} \frac{\min_{i \in [n]} |u \cdot x_i|}{|v \cdot x_i|}. \quad (8)$$

which produces $f \in \mathcal{F}_2(\text{ReLU})$ such that (6) holds in expectation and (7) holds almost surely.

To deduce Theorem 2 from Theorem 3, apply Lemma 1 with $\alpha = \frac{1}{10} \cdot \sqrt{\frac{\log(1/\gamma)}{\log(2n)}}$ and $\beta = \frac{\omega \cdot n}{d}$. See supplementary material for the proof of Theorem 3.

For $u \in \mathbb{R}^d$, set

$$f_u(x) = \psi'(u \cdot x)v \cdot x, \quad (9)$$

where v is defined as in (8). Observe that as long as $u \cdot x_i \neq 0, \forall i \in [n]$, a small enough choice of δ ensures the existence of $f \in \mathcal{F}_2(\text{ReLU})$ such that $\mathbf{f} = \mathbf{f}_u$.

To prove Theorem 3, it therefore remains to show that \mathbf{f}_u satisfies (6) and (7) with positive probability as $u \sim \mathcal{N}(0, \mathbf{I}_d)$. This will be carried out in two steps: First we show that the correlation $\mathbf{y} \cdot \mathbf{f}$ for a derivative neuron has a particularly nice form as a function of u , see Lemma 2. Then, in Lemma 3 we derive a lower bound for the expectation of the correlation under $u \sim \mathcal{N}(0, \mathbf{I}_d)$. Taken together these lemmas complete the proof of Theorem 3.

Lemma 2 Fix $u \in \mathbb{R}^d$, the function f_u defined in (9) satisfies $\sum_{i=1}^n y_i f_u(x_i) = \left\| \sum_{i: u \cdot x_i \geq 0} y_i x_i \right\|^2$, and furthermore $\sum_{i=1}^n f_u(x_i)^2 \leq \frac{\omega \cdot n}{d} \cdot \sum_{i=1}^n y_i f(x_i)$.

Proof. We may write

$$\sum_{i=1}^n f_u(x_i) y_i = \sum_{i=1}^n \psi'(u \cdot x_i) y_i x_i \cdot v.$$

To maximize this quantity we take $v = \sum_{i=1}^n \psi'(u \cdot x_i) y_i x_i$ so that the correlation is exactly equal to $\|v\|^2 = \left\| \sum_{i=1}^n \psi'(u \cdot x_i) y_i x_i \right\|^2$ (note also that $\psi'(t) = \mathbf{1}\{t \geq 0\}$ for the ReLU). Moreover we also have (recall that for ReLU, $|\psi'(t)| \leq 1$)

$$\sum_{i=1}^n f_u(x_i)^2 = \sum_{i=1}^n (\psi'(x_i \cdot u))^2 (x_i \cdot v)^2 \leq \lambda_{\max} \left(\sum_{i=1}^n x_i x_i^\top \right) \cdot \|v\|^2. \quad (10)$$

□

Lemma 3 One has:

$$\mathbb{E}_{u \sim \mathcal{N}(0, \mathbf{I}_n)} \left\| \sum_{i: u \cdot x_i \geq 0} y_i x_i \right\|^2 \geq \frac{1}{10} \cdot \sqrt{\frac{\log(1/\gamma)}{\log(2n)}} \cdot \sum_{i=1}^n y_i^2 \|x_i\|^2.$$

Proof. First note that $\mathbb{E} \left\| \sum_{i: u \cdot x_i \geq 0} y_i x_i \right\|^2 = \mathbf{y}^\top H \mathbf{y}$, where

$$H_{i,j} = \mathbb{E}[x_i \cdot x_j \mathbf{1}\{u \cdot x_i \geq 0\} \mathbf{1}\{u \cdot x_j \geq 0\}] = \frac{2}{\pi} x_i \cdot x_j \left(\frac{1}{4} + \arcsin \left(\frac{x_i}{\|x_i\|} \cdot \frac{x_j}{\|x_j\|} \right) \right).$$

Let us denote V the matrix with entries $V_{i,j} = \frac{x_i}{\|x_i\|} \cdot \frac{x_j}{\|x_j\|}$ and D the diagonal matrix with entries $\|x_i\|$. Note that $V \succeq 0$ and thus we have (recall also that $\arcsin(t) = \sum_{i=0}^{\infty} \frac{(2i)!}{(2^i i!)^2} \cdot \frac{t^{2i+1}}{2i+1}$):

$$D^{-1} H D^{-1} \succeq \frac{2}{\pi} \sum_{i=0}^{\infty} \frac{(2i)!}{(2^i i!)^2} \cdot \frac{V^{\circ 2(i+1)}}{2i+1}.$$

Now observe that for any i , by the Schur product theorem one has $V^{\circ i} \succeq 0$. Moreover $V^{\circ i}$ is equal to 1 on the diagonal, and off-diagonal it is smaller than γ^i , and thus for $i \geq \frac{\log(2n)}{\log(1/\gamma)}$ one has $V^{\circ i} \succeq \frac{1}{2} \mathbf{I}_n$. The conclusion now follows by an easy calculation. □

5 The complex network

We now wish to improve upon the NTK construction, by creating a network with similar memorization properties and which has almost no excess total weight.

5.1 Correlation of a perturbed neuron with random sign

Towards understanding our construction, let us first revisit the task of correlating a *single* neuron with the data, namely we want to maximize over w the ratio between $|\sum_{i=1}^n y_i \psi(w \cdot x_i)|$ and $\sqrt{\sum_{i=1}^n \psi(w \cdot x_i)^2}$. Note that depending on whether the sign of the correlation is positive or negative, one would eventually take either neuron $x \mapsto \psi(w \cdot x)$ or $x \mapsto -\psi(w \cdot x)$. Let us first revisit the NTK calculation from the previous section, emphasizing that one can take a random sign for the recombination weight a .

The key NTK-like observation is that a single neuron perturbed around the parameter w_0 and with random sign can be interpreted as a linear model over a feature mapping that depends on w . More precisely (note that the random sign cancels the 0^{th} order term in the Taylor expansion):

$$\mathbb{E}_{a \sim \{-\delta, \delta\}} a^{-1} \psi((w + av) \cdot x) = \Phi_w(x) \cdot v + O(\delta), \text{ where } \Phi_w(x) = \psi'(w \cdot x)x. \quad (11)$$

In particular the correlation to the data of such a single random neuron is equal in expectation to $\sum_i y_i \Phi_w(x_i) \cdot v + O(\delta)$, and thus it is natural to take the perturbation vector v to be equal to $v_0 = \eta \sum_i y_i \Phi_w(x_i)$ (where η will be optimized to balance with the variance term), and we now find that:

$$\mathbb{E}_{a \sim \{-\delta, \delta\}} \sum_{i=1}^n y_i a^{-1} \psi((w + av_0) \cdot x_i) = \left\| \eta \sum_i y_i \Phi_w(x_i) \right\|^2 + O(\delta) = \eta y^\top H(w) y + O(\delta), \quad (12)$$

where $H(w)$ is the Gram matrix of the feature embedding, namely

$$H(w)_{i,j} = \Phi_w(x_i) \cdot \Phi_w(x_j).$$

Note that for $\psi = \text{ReLU}$, one has in fact that the term $O(\delta)$ in (11) disappears for δ small is enough, and thus the correlation to the data is simply $\eta y^\top H(w) y$ in that case.

As we did with the NTK network, we now also take the base parameter w at random from a standard Gaussian. As we just saw, understanding the expected correlation then reduces to lower bound (spectrally) the Gram matrix H defined by $H_{i,j} = \mathbb{E}_{w \sim \mathcal{N}(0, I_d)} [\psi'(w \cdot x_i) \psi'(w \cdot x_j) x_i \cdot x_j]$. This was exactly the content of Lemma 3 for $\psi = \text{ReLU}$.

5.2 Eliminating the higher derivatives with a complex trick

The main issue of the strategy described above is that it requires to take δ small, which in turn may significantly increase the total weights of the resulting network. Our next idea is based on the following observation: Taking a random sign in (11) eliminates all the even order term in the Taylor expansion since $\mathbb{E}_{a \sim \{-1, 1\}} [a^{-1} a^m] = 0$ for any even m (while it is $= 1$ for any odd m). However, taking a *complex* a , would rid us of *all* terms except the first order term. Namely, one has $\mathbb{E}_{a \in \mathbb{C}: |a|=1} [a^{-1} a^m] = 0$ for any $m \neq 1$. This suggests that it might make sense to consider neurons of the form

$$x \mapsto \text{Re} (a^{-1} \psi((w + av) \cdot x)),$$

where a is a complex number of unit norm.

The challenge is now to give sense to $\psi(z)$ for a complex z , so that the rest of the argument remains unchanged. This gives rise to two caveats:

- There is no holomorphic extension of the ReLU function.
- The holomorphic extension of the activation function, even if exists, is a function of two (real) variables. The expression $\psi((w + av) \cdot x)$ when $a \notin \mathbb{R}$ is not a valid neuron to be used in our construction since we're only allowed to use the original activation function as our non-linearity.

To overcome these caveats, the construction will be carried out in two steps, where in the first step we use *polynomial* activation functions, and in the second step, we replace these by the original activation function. It turns out that the calculation in Lemma 3 is particularly simple when the

derivative of the activation function is a Hermite polynomial (see supplementary material for definitions), which is in particular obviously well-defined on \mathbb{C} and in fact holomorphic.

The first step of our proof in the supplementary material will be to obtain a result analogous to Theorem 1 where the ReLU is replaced by a Hermite activation. Given such a result, the second step towards Theorem 1 is to replace the polynomial attained by the above lemma by a ReLU. This will be achieved by:

- Observing that any polynomial in two variables $p(x, y)$ can be written as a linear combination of polynomials which only depend on one direction, hence polynomials of the form $q(ax + by)$.
- Using the fact that any nice enough function of one variable can be written as a mixture of ReLUs, due to the fact that the second derivative of the ReLU is a Dirac function (this was observed before, see e.g., [Lemma A.4, Ji et al. [2020]]).
- The above implies that one can write the function $(x, y) \mapsto \varphi(x + iy)$, where φ is the Hermite activation, as the expectation of ReLUs such that the variance at points close to the origin is not too large.

These steps will be carried out in the supplementary material.

Broader impact. This work does not present any foreseeable societal consequence

References

- Zeyuan Allen-Zhu, Yuanzhi Li, and Zhao Song. A convergence theory for deep learning via over-parameterization. In *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 242–252. PMLR, 2019.
- Alexandr Andoni, Rina Panigrahy, Gregory Valiant, and Li Zhang. Learning polynomials with neural networks. In *International conference on machine learning*, pages 1908–1916, 2014.
- Sanjeev Arora, Simon Du, Wei Hu, Zhiyuan Li, and Ruosong Wang. Fine-grained analysis of optimization and generalization for overparameterized two-layer neural networks. In *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 322–332. PMLR, 2019.
- Francis Bach. Breaking the curse of dimensionality with convex neural networks. *The Journal of Machine Learning Research*, 18(1):629–681, 2017.
- Andrew R. Barron. Universal approximation bounds for superpositions of a sigmoidal function. *IEEE Transactions on Information Theory*, 39(3):930–945, 1993.
- Peter L Bartlett. The sample complexity of pattern classification with neural networks: the size of the weights is more important than the size of the network. *IEEE transactions on Information Theory*, 44(2):525–536, 1998.
- Peter L Bartlett, Dylan J Foster, and Matus J Telgarsky. Spectrally-normalized margin bounds for neural networks. In *Advances in Neural Information Processing Systems*, pages 6240–6249, 2017.
- Eric B Baum. On the capabilities of multilayer perceptrons. *Journal of complexity*, 4(3):193–215, 1988.
- Yoshua Bengio, Nicolas L Roux, Pascal Vincent, Olivier Delalleau, and Patrice Marcotte. Convex neural networks. In *Advances in neural information processing systems*, pages 123–130, 2006.
- Guy Bresler and Dheeraj Nagaraj. A corrective view of neural networks: Representation, memorization and learning. *arXiv preprint arXiv:2002.00274*, 2020.
- Lenaic Chizat and Francis Bach. On the global convergence of gradient descent for over-parameterized models using optimal transport. In *Advances in Neural Information Processing Systems 31*, pages 3036–3046. 2018.
- Lenaic Chizat, Edouard Oyallon, and Francis Bach. On lazy training in differentiable programming. In *Advances in Neural Information Processing Systems 32*, pages 2937–2947. 2019.
- George Cybenko. Approximation by superpositions of a sigmoidal function. *Mathematics of control, signals and systems*, 2(4):303–314, 1989.
- Amit Daniely. Neural networks learning and memorization with (almost) no over-parameterization. *arXiv preprint arXiv:1911.09873*, 2019.
- Amit Daniely. Memorizing gaussians with no over-parameterization via gradient descent on neural networks. *arXiv preprint arXiv:2003.12895*, 2020.

- Simon Du, Jason Lee, Haochuan Li, Liwei Wang, and Xiyu Zhai. Gradient descent finds global minima of deep neural networks. In *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 1675–1685. PMLR, 2019.
- Charles Fefferman. Reconstructing a neural net from its output. *Revista Matemática Iberoamericana*, 10(3):507–555, 1994.
- Behrooz Ghorbani, Song Mei, Theodor Misiakiewicz, and Andrea Montanari. Limitations of lazy training of two-layers neural network. In *Advances in Neural Information Processing Systems*, pages 9108–9118, 2019.
- Arthur Jacot, Franck Gabriel, and Clément Hongler. Neural tangent kernel: Convergence and generalization in neural networks. In *Advances in neural information processing systems*, pages 8571–8580, 2018.
- Svante Janson. *Gaussian Hilbert spaces*, volume 129. Cambridge university press, 1997.
- Ziwei Ji and Matus Telgarsky. Polylogarithmic width suffices for gradient descent to achieve arbitrarily small test error with shallow relu networks. In *International Conference on Learning Representations*, 2020.
- Ziwei Ji, Matus Telgarsky, and Ruicheng Xian. Neural tangent kernels, transportation mappings, and universal approximation. In *International Conference on Learning Representations*, 2020.
- Adam Tauman Kalai, Adam R Klivans, Yishay Mansour, and Rocco A Servedio. Agnostically learning halfspaces. *SIAM Journal on Computing*, 37(6):1777–1805, 2008.
- Kenji Kawaguchi and Jiaoyang Huang. Gradient descent finds global minima for generalizable deep neural networks of practical sizes. In *57th Allerton Conference on Communication, Control, and Computing (Allerton)*. IEEE, 2019.
- Moshe Leshno, Vladimir Ya Lin, Allan Pinkus, and Shimon Schocken. Multilayer feedforward networks with a nonpolynomial activation function can approximate any function. *Neural networks*, 6(6):861–867, 1993.
- Yuanzhi Li and Yingyu Liang. Learning overparameterized neural networks via stochastic gradient descent on structured data. In *Advances in Neural Information Processing Systems*, pages 8157–8166, 2018.
- Song Mei, Andrea Montanari, and Phan-Minh Nguyen. A mean field view of the landscape of two-layer neural networks. *Proceedings of the National Academy of Sciences*, 115(33):E7665–E7671, 2018.
- Behnam Neyshabur, Ryota Tomioka, and Nathan Srebro. Norm-based capacity control in neural networks. In *Conference on Learning Theory*, pages 1376–1401, 2015.
- Greg Ongie, Rebecca Willett, Daniel Soudry, and Nathan Srebro. A function space view of bounded norm infinite width relu nets: The multivariate case. In *International Conference on Learning Representations*, 2020.
- Samet Oymak and Mahdi Soltanolkotabi. Towards moderate overparameterization: global convergence guarantees for training shallow neural networks. *CoRR*, abs/1902.04674, 2019.
- Pedro Savarese, Itay Evron, Daniel Soudry, and Nathan Srebro. How do infinite width bounded norm networks look in function space? In *Proceedings of the Thirty-Second Conference on Learning Theory*, volume 99 of *Proceedings of Machine Learning Research*, pages 2667–2690. PMLR, 2019.
- Shai Shalev-Shwartz and Shai Ben-David. *Understanding machine learning: From theory to algorithms*. Cambridge university press, 2014.
- Mahdi Soltanolkotabi, Adel Javanmard, and Jason D Lee. Theoretical insights into the optimization landscape of over-parameterized shallow neural networks. *IEEE Transactions on Information Theory*, 65(2):742–769, 2018.
- Zhao Song and Xin Yang. Quadratic suffices for over-parametrization via matrix chernoff bound. *CoRR*, abs/1906.03593, 2019.
- Gilad Yehudai and Ohad Shamir. On the power and limitations of random features for understanding neural networks. In *Advances in Neural Information Processing Systems*, pages 6594–6604, 2019.
- Chulhee Yun, Suvrit Sra, and Ali Jadbabaie. Small relu networks are powerful memorizers: a tight analysis of memorization capacity. In *Advances in Neural Information Processing Systems*, pages 15532–15543, 2019.
- Chiyuan Zhang, Samy Bengio, Moritz Hardt, Benjamin Recht, and Oriol Vinyals. Understanding deep learning requires rethinking generalization. In *5th International Conference on Learning Representations, ICLR 2017*. OpenReview.net, 2017.

A Baum network total weight

We sketch here the calculation of the Baum network's total weight in the case that the x_i 's are independent uniform points on \mathbb{S}^{d-1} and y_i are ± 1 -Bernoulli distributed. We will show that the total weight is at least n^2/\sqrt{d} , thus more than n times the optimal attainable weight given in Proposition 1.

Consider the matrix X whose rows are the vectors $(x_i)_{i \leq d}$. The vector v taken in the neuron corresponding to those points solves the equation $Xv = y$ and since the distribution of X is absolutely continuous, we have that X is invertible almost surely and therefore $v = X^{-1}y$, implying that $|v| \geq \|X\|_{OP}^{-1}\sqrt{d}$. It is well-known (and easy to show) that with overwhelming probability, $\|X\|_{OP} = O(1)$, and thus $\|v\| = \Omega(\sqrt{d})$.

Observe that by normalizing the parameter δ accordingly, we can assume that $\|u\| = 1$. By definition we have $u \cdot x_i = b$ for all $i = 1, \dots, d$. A calculation shows that with probability $\Omega(1)$ we have $b = \Theta(1/\sqrt{d})$.

Next, we claim that $|v \cdot u| \leq (1 - \rho)\|v\|$ for some $\rho = \Omega(1)$. Indeed, suppose otherwise. Denote $c = \frac{1}{d} \sum_{i \in [d]} x_i$. It is easy to check that with high probability, $\|c\| = O\left(\frac{1}{\sqrt{d}}\right)$. Note that $v \cdot c = \frac{1}{d} \sum_{i \in [d]} y_i = O(1/\sqrt{d})$. This implies that

$$b(|v \cdot u| - O(1)) \leq |v \cdot (bu - c)| \leq \sqrt{\|v\|^2 - (v \cdot u)^2} \|bu - c\| \leq \sqrt{2\rho} \frac{\|v\|}{\sqrt{d}},$$

where we used the fact that $(bu - c) \perp (v \cdot u)u$. Thus we have

$$\Omega(1 - 2\rho) = b(1 - 2\rho)\|v\| = O(\sqrt{\rho}).$$

leading to a contradiction. To summarize, we have $\|v\| = \Omega(\sqrt{d})$, $\|u\| = 1$, $|u \cdot v| \leq (1 - \rho)\|v\|$, $\rho = \Omega(1)$, and $b = O(1/\sqrt{d})$. Since spherical marginals are approximately Gaussian, if x is uniform in \mathbb{S}^{d-1} we have that the joint distribution of $(x \cdot u, x \cdot v)$ conditional on v and u is approximately $\mathcal{N}\left(0, \frac{1}{d} \begin{pmatrix} 1 & (1 - \rho)\beta \\ (1 - \rho)\beta & \beta \end{pmatrix}\right)$ with $\rho = \Omega(1)$ and $\beta = \Theta(d)$. Therefore, with probability $\Omega(1/n)$ we have $|x \cdot v| = \Omega(1)$ and $|x \cdot u - b| = O(1/(n\sqrt{d}))$.

We conclude that

$$\mathbb{P}\left(\exists i \geq d + 1 \text{ s.t. } \frac{|x_i \cdot u - b|}{|x_i \cdot v|} = O\left(\frac{1}{n\sqrt{d}}\right) \mid x_1, \dots, x_d\right) = \Omega(1).$$

Therefore, we get $\delta = O(1/n\sqrt{d})$ which implies that the weight of the neuron is of order at least $\frac{\|u\|}{\delta} = \Omega(n\sqrt{d})$. This happens with probability $\Omega(1)$ for every one of the first $n/(2d)$ neurons, implying that the total weight is of order n^2/\sqrt{d} .

B NTK network total weight

We sketch here the calculation of the NTK network's total weight. Recall that the neurons are of the form (9). According to Lemma 3, we have that for typical neurons, $\|v\| = \Omega(\sqrt{n})$. Moreover, with high probability we have $\|u\| = \Theta(\sqrt{d})$, and thus the weight of a single neuron is at least $\frac{\|u\|}{\delta} = \frac{\sqrt{d}}{\delta}$. Adding up the neurons, this shows that the total weight is of order $\frac{\sqrt{d}}{\delta}$ (since $k = \tilde{\Theta}(n/d)$ and the coefficient in front of the neurons is of order $\tilde{\Theta}(\frac{d}{n})$).

Now suppose that δ is taken according to (4). The main observation (we omit the details of proof) is that u and v have a mutual distribution of roughly independent Gaussian random vectors (without loss of generality we can assume that $\sum y_i = 0$ which implies $\mathbb{E}u \cdot v = 0$). In this case we have $\delta = \tilde{O}\left(\frac{\sqrt{d}}{n\sqrt{n}}\right)$. This implies a total weight of order at least $n\sqrt{n}$.

C Hermite polynomials

Define the m 'th Hermite polynomial by:

$$H_m(x) = \frac{(-1)^m}{\sqrt{m!}} \left(\frac{d^m}{dx^m} e^{-\frac{x^2}{2}} \right) e^{\frac{x^2}{2}}.$$

For ease of notion we also define $H_{-1} \equiv 0$. The Hermite polynomials may also be regarded as the power series associated to the function $F(t, x) = \exp(tx - \frac{t^2}{2})$. Indeed,

$$\begin{aligned} F(t, x) &= \exp\left(\frac{x^2}{2} - \frac{(x-t)^2}{2}\right) \\ &= e^{\frac{x^2}{2}} \sum_{\ell=0}^{\infty} \frac{t^\ell}{\ell!} \left(\frac{d^\ell}{dt^\ell} e^{-\frac{(x-t)^2}{2}} \right) \Big|_{t=0} \\ &= \sum_{m=0}^{\infty} \frac{t^m}{\sqrt{m!}} H_m(x). \end{aligned} \tag{13}$$

Observe that $\frac{d}{dx}F(t, x) = tF(t, x)$, so that, since $H_0 \equiv 1$,

$$\sum_{m=1}^{\infty} \frac{t^m}{\sqrt{(m-1)!}} H_{m-1}(x) = \sum_{m=1}^{\infty} \frac{t^m}{\sqrt{m!}} H'_m(x),$$

and we deduce

$$H'_m = \sqrt{m} H_{m-1}. \tag{14}$$

Also $\frac{d}{dt}F(t, x) = (x-t)F(t, x)$ and a similar argument shows that

$$\sqrt{\frac{m}{m-1}} H_m(x) = \frac{x}{\sqrt{m-1}} H_{m-1}(x) - H_{m-2}(x). \tag{15}$$

Furthermore, we show that the family $\{H_m\}$ satisfies the following orthogonality relation, which we shall freely use.

Lemma 4 *Let $X, Y \sim \mathcal{N}(0, 1)$ be jointly Gaussian with $\mathbb{E}[XY] = \rho$. Then*

$$\mathbb{E}[H_m(X)H_{m'}(Y)] = \delta_{m,m'} \rho^m.$$

Proof. Fix $s, t \in \mathbb{R}$. We have the following identity

$$\mathbb{E}[F(s, X)F(t, Y)] = \mathbb{E}[\exp(sX + tY)] \exp\left(-\frac{s^2 + t^2}{2}\right) = e^{st \cdot \rho},$$

where in the second equality we have used the formula for the moment generating functions of bi-variate Gaussians. In particular, we have

$$\frac{d^{m+m'}}{ds^m dt^{m'}} \mathbb{E}[F(s, X)F(t, Y)] \Big|_{t=0, s=0} = \frac{d^{m+m'}}{ds^m dt^{m'}} e^{st \cdot \rho} \Big|_{t=0, s=0}.$$

By (13), the left hand side equals $\mathbb{E}[H_\ell(X)H_{\ell'}(Y)]$, while the right hand side is $\delta_{m,m'} \rho^m$. The proof is complete. \square

D More general non-linearities

We now consider an arbitrary L -Lipschitz non-linearity ψ that is differentiable except at a finite number of points and such that $\mathbb{E}_{X \sim \mathcal{N}(0,1)}[(\psi'(X))^2] < +\infty$. In particular, with H_1, H_2, \dots being the Hermite polynomials (normalized such that it forms an orthonormal basis) we have that there exists a sequence of real numbers (a_ℓ) such that

$$\psi' = \sum_{\ell \geq 0} a_\ell H_\ell.$$

Our generalization of Theorem 2 now reads as follows:

Theorem 4 Under the above assumptions on ψ , there exists $f \in \mathcal{F}_k(\psi)$ with $\|\mathbf{f} - \mathbf{y}\|^2 \leq \varepsilon \|\mathbf{y}\|^2$ provided that

$$k \cdot d \geq \frac{16\omega \cdot L}{\sum_{\ell \geq \frac{\log(2n)}{2\log(1/\gamma)}} a_\ell^2} \cdot n \log(1/\varepsilon).$$

In fact there is an efficient procedure that produces a random $f \in \mathcal{F}_k(\psi)$ with $\mathbb{E}[\|\mathbf{f} - \mathbf{y}\|^2] \leq \varepsilon \|\mathbf{y}\|^2$ when (5) holds.

Proof. First we follow the proof of Lemma 2, with the only change being: (i) in (9) there is an additive $O(\delta)$ term (also now the condition on u is that $u \cdot x_i$ is not in the finite set of points where ψ is not differentiable), and (ii) in (10) we use that $|\psi'| \leq L$. We obtain that for $u \in \mathbb{R}^d$ there exists $f \in \mathcal{F}_2(\psi)$ such that

$$\sum_{i=1}^n y_i f(x_i) \geq \frac{1}{2} \left\| \sum_{i=1}^n \psi'(u \cdot x_i) y_i x_i \right\|^2, \quad (16)$$

and furthermore

$$\sum_{i=1}^n f(x_i)^2 \leq \frac{2\omega \cdot n \cdot L}{d} \cdot \sum_{i=1}^n y_i f(x_i), \quad (17)$$

where the added term L is due to modification (ii) above and the added 2 is due to (i).

Next we follow the proof of Lemma 3, noting that the matrix H is now defined by (recall Lemma 4) $H_{i,j} = \sum_{\ell \geq 0} a_\ell^2 (x_i \cdot x_j)^{\ell+1}$, to obtain:

$$\mathbb{E}_{u \sim \mathcal{N}(0, I_n)} \left\| \sum_{i=1}^n \psi'(u \cdot x_i) y_i x_i \right\|^2 \geq \frac{1}{2} \sum_{\ell \geq \frac{\log(2n)}{2\log(1/\gamma)}} a_\ell^2 \cdot \sum_{i=1}^n y_i^2. \quad (18)$$

In particular we obtain from (16) and (18) that (6) holds true with the term $\frac{1}{10} \cdot \sqrt{\frac{\log(1/\gamma)}{\log(2n)}}$ replaced by $\frac{1}{4} \sum_{\ell \geq \frac{\log(2n)}{2\log(1/\gamma)}} a_\ell^2$, and from (17) that (7) holds true with ω replaced by $2\omega \cdot L$. We can thus conclude as we concluded Theorem 2 from Theorem 3. \square

E More details on the complex network

We will work under the assumptions that

$$\|x_i\| = 1 \text{ for every } i \in [n], \text{ and, } |x_i \cdot x_j| \leq \gamma \text{ for } i \neq j. \quad (19)$$

In light of Lemma 1, it is enough to find a single neuron whose scalar product with the data set is large. Thus, the rest of this section is devoted to proving the following theorem.

Theorem 5 Assume that (19) holds, that m is large enough so that $n\gamma^{m-2} \leq \frac{1}{2}$ and that for all $i \in [n]$, $y_i^2 \leq n\gamma^2$ with $\|\mathbf{y}\|^2 \leq n$. Then, there exist $w \in \mathbb{R}^d$ and $b, \sigma \in \mathbb{R}$, with

$$\|w\|^2, |b|^2 \leq C_m d \log(n)^m, |\sigma| = 1,$$

such that for

$$f(x) = \sigma \cdot \text{ReLU}(w \cdot x + b),$$

we have

$$\mathbf{y} \cdot \mathbf{f} \geq \frac{c_m}{\log(n)^{m^2/2}} \frac{1}{\sqrt{n\gamma^2}} \|\mathbf{y}\|^2,$$

and

$$\|\mathbf{f}\|^2 \leq \frac{n}{c_m} \log(n)^m,$$

where $c_m, C_m > 0$ are constants which depends only on m .

Equipped with the above result one can prove the following formal version of Theorem 1. If $A \subset [n]$ and $v \in \mathbb{R}^n$ we denote below by v_A the projection of v unto the indices contained in A . With this notation our result is:

Theorem 6 Assume that (19) holds, that m is large enough so that $n\gamma^{m-2} \leq \frac{1}{2}$ and that $\|\mathbf{y}\|^2 = n$. There exists $f \in \mathcal{F}_k(\text{ReLU})$ and $A \subset [n]$, with

$$k = \left\lceil C_m \gamma^2 \frac{\log(1/\varepsilon)}{\varepsilon} n \log(n)^{(m^2+m)} \right\rceil,$$

such that

$$\mathbb{E}[\|\mathbf{f}_A - \mathbf{y}_A\|^2] \leq \varepsilon \|\mathbf{y}\|^2, \quad |A| \geq n - \frac{1}{\gamma^2}, \quad (20)$$

and

$$\mathbf{W}(f) = \tilde{O} \left(\frac{\log(1/\varepsilon)}{\varepsilon} \sqrt{n\gamma^2 d} \right), \quad (21)$$

where C_m is a constant which depends only on m .

Observe that if $(x_i)_{i \in [n]}$ are uniformly distributed in the \mathbb{S}^{d-1} then $\gamma = \tilde{O} \left(\frac{1}{\sqrt{d}} \right)$ and we get that $\mathbf{W}(f) = \tilde{O} \left(\frac{\log(1/\varepsilon)}{\varepsilon} \sqrt{n} \right)$, which is optimal up to the logarithmic factors and the dependence on ε .

The proof of Theorem 6 follows an iterative procedure similar to the one carried out in Lemma 1. The only caveat is the condition $y_i^2 \leq n\gamma^2$ which appears in Theorem 5. Due to this condition we need to consider a slightly smaller set of indices at each iteration, ignoring ones where the residue becomes too big.

Proof.[of Theorem 6] We build the network iteratively. Set $f_0 \equiv 0$, $A_0 = [n]$ and $r_{0,i} = y_i$. Now, for $\ell \in \mathbb{N}$, suppose that there exists $f_\ell \in \mathcal{F}_\ell(\text{ReLU})$ with

$$\|(\mathbf{f}_\ell)_{A_\ell} - \mathbf{y}_{A_\ell}\| \leq \left(1 - \frac{c_m^3}{\log(n)^{m^2+m}} \frac{\varepsilon}{n\gamma^2} \right) \|\mathbf{y}\|^2.$$

Set $r_{\ell,i} = y_i - f_\ell(x_i)$ and $A_\ell = \{i \in A_{\ell-1} \mid r_{\ell,i}^2 \leq n\gamma^2\}$. We now invoke Theorem 5 with the residuals $\{r_{\ell,i} \mid i \in A_\ell\}$ to obtain a neuron $f \in \mathcal{F}_1(\text{ReLU})$, which satisfies

$$(\mathbf{r}_\ell)_{A_\ell} \cdot \mathbf{f} \geq \frac{c_m}{\log(n)^{m^2/2}} \frac{1}{\sqrt{n\gamma^2}} \|\mathbf{r}_\ell\|^2,$$

and

$$\|\mathbf{f}_{A_\ell}\|^2 \leq \frac{n}{c_m} \log(n)^m.$$

Since we may assume $\|(\mathbf{r}_\ell)_{A_\ell}\|^2 \geq n\varepsilon$ (otherwise we are done), the second condition can be rewritten as

$$\|\mathbf{f}_{A_\ell}\|^2 \leq \frac{\log(n)^m}{c_m \varepsilon} \|(\mathbf{r}_\ell)_{A_\ell}\|^2.$$

In this case the calculation done in Lemma 1 with $\alpha = \frac{c_m}{\log(n)^{m^2/2}} \frac{1}{\sqrt{n\gamma^2}}$ and $\beta = \frac{\log(n)^m}{c_m \varepsilon}$ shows that for $\eta := \frac{c_m^2 \varepsilon}{\log(n)^{m^2/2+m}}$, one has

$$\|\eta \mathbf{f}_{A_\ell} - (\mathbf{r}_\ell)_{A_\ell}\|^2 \leq \left(1 - \frac{c_m^3}{\log(n)^{m^2+m}} \frac{\varepsilon}{n\gamma^2} \right) \|(\mathbf{r}_\ell)_{A_\ell}\|^2.$$

In other words, if we define $f_{\ell+1} \in \mathcal{F}_{\ell+1}(\text{ReLU})$ by $f_{\ell+1} = f_\ell + \eta f$,

$$\|(\mathbf{f}_{\ell+1})_{A_\ell} - \mathbf{y}_{A_\ell}\|^2 \leq \left(1 - \frac{c_m^3}{\log(n)^{m^2+m}} \frac{\varepsilon}{n\gamma^2} \right)^{\ell+1} \|\mathbf{y}\|^2.$$

The estimate (20) is now obtained with the appropriate choice of k . Let us also remark that for any ℓ ,

$$\|(\mathbf{r}_{\ell+1})_{A_\ell}\|^2 \leq \|(\mathbf{r}_\ell)_{A_\ell}\|^2 \leq \|(\mathbf{r}_\ell)_{A_{\ell-1}}\|^2 - n\gamma^2|A_{\ell-1} \setminus A_\ell|.$$

By induction

$$\|(\mathbf{r}_{\ell+1})_{A_\ell}\|^2 \leq \|\mathbf{y}\|^2 - n\gamma^2(n - |A_\ell|)$$

This shows that $|A_\ell| \geq n - \frac{1}{\gamma^2}$. The bound on $\mathbf{W}(f_k)$ a direct consequence of Lemma 1. \square

E.1 Constructing the complex neuron

In the sequel, we fix $m \in \mathbb{N}$ so that

$$n\gamma^{m-2} \leq \frac{1}{2}. \quad (22)$$

Define

$$\varphi(z) = \frac{1}{\sqrt{m}} H_m(z), \quad z \in \mathbb{C}$$

where H_m is the m -th Hermite polynomial. Note that we also have $\varphi' = H_{m-1}$.

We first prove the following analogous result to Theorem 5 where ψ is replaced by φ .

Lemma 5 *Assume that (19), (22) hold, and that for all $i \in [n]$, one has $y_i^2 \leq n\gamma^2$. There exist $\tilde{w}, \tilde{w}' \in \mathbb{R}^d$ and $z \in \mathbb{C}$, $|z| = 1$, such that for*

$$g(x) = \operatorname{Re}(z \cdot \varphi((\tilde{w} + \mathbf{i}\tilde{w}') \cdot x)), \quad (23)$$

we have,

$$\mathbf{y} \cdot \mathbf{g} \geq \frac{1}{2C_m \sqrt{n\gamma^2}} \|\mathbf{y}\|^2.$$

Moreover, its weights admit the bounds

$$\|\tilde{w}\|^2, \|\tilde{w}'\|^2 \leq d(4C_m \log(n))^m \quad (24)$$

and for all $i \in [n]$,

$$|\tilde{w} \cdot x_i|, |\tilde{w}' \cdot x_i| \leq (4C_m \log(n))^{\frac{m}{2}}.$$

Our approach to Lemma 5 will be to construct an appropriate distribution on neurons of type (23), and then show that the desirable properties are attained with positive probability. In what follows, let $w \sim \mathcal{N}(0, \mathbf{I}_d)$. Define

$$v(w) := \frac{1}{\sqrt{n\gamma^2}} \sum_{i=1}^n y_i \varphi'(w \cdot x_i) x_i.$$

Next, let a be uniformly distributed in the complex unit circle, and finally define

$$g(x) = \operatorname{Re}(a^{-1} \varphi((w + av(w)) \cdot x)). \quad (25)$$

We will prove the following two bounds.

Lemma 6 *Under the assumptions (19) and (22), one has*

$$\mathbb{E}[\mathbf{y} \cdot \mathbf{g}] \geq \frac{1}{2\sqrt{n\gamma^2}} \|\mathbf{y}\|^2.$$

Lemma 7 *Suppose that the assumptions (19) and (22) hold. Assume also that for every i we have $y_i \leq n\gamma^2$. Then one has, for a constant $C_m > 0$ which depends only on m ,*

$$\mathbb{E}[\|\mathbf{g}\|^2] \leq C_m n.$$

Moreover, for every $i \in [n]$ and $s > s_0$, for some constant s_0 ,

$$\mathbb{P}(|\operatorname{Re}((w + v(w)) \cdot x_i)| > s), \mathbb{P}(|\operatorname{Im}((w + v(w))) \cdot x_i| > s) \leq \exp\left(-\frac{1}{C_m} s^{-2/m}\right). \quad (26)$$

Recall the definition of the Gram matrix H ,

$$H_{i,j} = \mathbb{E}_{w \sim \mathcal{N}(0, \mathbf{I}_d)} [\varphi'(w \cdot x_i) \varphi'(w \cdot x_j) x_i \cdot x_j].$$

As suggested in (12), we will need to bound H from below. We will need the following lemma.

Lemma 8 *Under the assumptions (19) and (22), one has $H \succeq \frac{1}{2} \mathbf{I}_n$.*

Proof. If X and Y are standard, jointly-normal random variables with $\mathbb{E}[XY] = \rho$, by Lemma 4 one has $\mathbb{E}[H_{m-1}(X)H_{m-1}(Y)] = \rho^{m-1}$ and thus here $H_{i,j} = (x_i \cdot x_j)^m$. In particular if $n \cdot \gamma^m \leq 1/2$ we obtain that for all $i \in [n]$ one has $1 = H_{i,i} \geq 2 \sum_{j \neq i} |H_{i,j}|$. By diagonal dominance we conclude that $H \succeq \frac{1}{2} \mathbf{I}_n$. \square

Proof. [Proof of Lemma 6] For any $\beta \in \mathbb{N}, \beta \neq 1$, we have that $\mathbb{E}[a^{-1+\beta}] = 0$. Thus, since φ is an entire function, by taking its Taylor expansion around the point w , we obtain the identity

$$\mathbb{E}_a [a^{-1} \varphi((w + av(w)) \cdot x)] = \sum_{\beta=0}^{\infty} \frac{1}{\beta!} \mathbb{E}_a [a^{-1+\beta} \varphi^{(\beta)}(w \cdot x_i) (v(w) \cdot x)^\beta] = \varphi'(w \cdot x) v(w) \cdot x.$$

So we can estimate

$$\begin{aligned} \mathbb{E}_{w,a} \left[\sum_{i=1}^n y_i \operatorname{Re} (a^{-1} \varphi((w + av(w)) \cdot x_i)) \right] &= \sum_{i=1}^n y_i \mathbb{E}_w [\varphi'(w \cdot x_i) v(w) \cdot x_i] \\ &= \frac{1}{\sqrt{n\gamma^2}} \sum_{i,j} y_i y_j \mathbb{E}_w [\varphi'(w \cdot x_i) \varphi'(w \cdot x_j) x_i \cdot x_j] \\ &= \frac{1}{\sqrt{n\gamma^2}} \mathbf{y}^\top H \mathbf{y} \geq \frac{1}{2\sqrt{n\gamma^2}} \|\mathbf{y}\|^2, \end{aligned}$$

where the last inequality follows from Lemma 8. \square

Proof. [Proof of Lemma 7] In what follows, the expression C_m will denote a constant depending only on m , whose value may change between different appearances. Our objective is to obtain an upper bound on

$$\|\mathbf{g}\|^2 = \sum_{i=1}^n |\operatorname{Re} (a^{-1} \varphi((w + av(w)) \cdot x_i))|^2.$$

Since φ is a polynomial of degree m we have

$$\|\mathbf{g}\|^2 \leq C_m \sum_{i=1}^n (1 + |w \cdot x_i|^{2m} + |v(w) \cdot x_i|^{2m}).$$

Moreover $w \cdot x_i$ is a standard Gaussian and thus $\mathbb{E}[|w \cdot x_i|^{2m}] \leq (2m)^m$. It therefore remains to control, for $x \in \{x_1, \dots, x_n\}$, the expression

$$|v(w) \cdot x|^{2m} = \frac{1}{(n\gamma^2)^m} \left| \sum_{i=1}^n y_i H_{m-1}(w \cdot x_i) x_i \cdot x \right|^{2m}.$$

From hypercontractivity and the fact that the Hermite polynomials are eigenfunctions of the Ornstein-Uhlenbeck operator we have (see [Janson, 1997, Theorem 5.8])

$$\mathbb{E}[|v(w) \cdot x|^{2m}] \leq (2m)^{2m^2} \mathbb{E}[|v(w) \cdot x|^2]^m.$$

Thus, it will be enough to show that $\mathbb{E}_w[|v(w) \cdot x_j|^2] = O(1)$. We calculate

$$\begin{aligned} \mathbb{E}_w[|v(w) \cdot x_j|^2] &= \frac{1}{n\gamma^2} \mathbb{E} \left| \sum_{i=1}^n y_i H_{m-1}(w \cdot x_i) x_i \cdot x_j \right|^2 \\ &= \frac{1}{n\gamma^2} \left(\mathbb{E} \sum_{i=1}^n y_i^2 \mathbb{E}[(H_{m-1}(w \cdot x_i))^2] |x_i \cdot x_j|^2 \right. \\ &\quad \left. + \sum_{i \neq i'} y_i y_{i'} \mathbb{E}[H_{m-1}(w \cdot x_i) H_{m-1}(w \cdot x_{i'})] (x_i \cdot x_j)(x_{i'} \cdot x_j) \right) \\ &\leq \frac{1}{n\gamma^2} \left(\sum_{i=1}^n y_i^2 |x_i \cdot x_j|^2 + \frac{\gamma^{m-1}}{n\gamma^2} \sum_{i \neq i'} |y_i y_{i'} (x_{i'} \cdot x_j)(x_i \cdot x_j)| \right), \end{aligned}$$

where we used that $\mathbb{E}[(H_{m-1}(w \cdot x_i))^2] = 1$ and

$$|\mathbb{E}[H_{m-1}(w \cdot x_i) H_{m-1}(w \cdot x_{i'})]| = |x_i \cdot x_{i'}|^{m-1} \leq \gamma^{m-1},$$

valid whenever $i \neq i'$. By using that $\|\mathbf{y}\|^2 = O(n)$, we get

$$\frac{1}{n\gamma^2} \sum_{i=1}^n y_i^2 |x_i \cdot x_j|^2 \leq \frac{y_j^2}{n\gamma^2} + \frac{\|\mathbf{y}\|^2}{n} = O(1).$$

To deal with the last term, observe that since $i \neq i'$ then $|(x_{i'} \cdot x_j)(x_i \cdot x_j)| \leq \gamma$, thus

$$\frac{\gamma^{m-1}}{n\gamma^2} \sum_{i \neq i'} |y_i y_{i'} (x_{i'} \cdot x_j)(x_i \cdot x_j)| \leq \frac{\gamma^{m-2}}{n} \left(\sum_{i=1}^n |y_i| \right)^2 \leq \gamma^{m-2} \|\mathbf{y}\|^2 = O(1),$$

where in the last inequality we've used $\gamma^{m-2} \leq \frac{1}{n}$. So, $\mathbb{E}_w[|v(w) \cdot x_i|^2] = O(1)$ as required.

Finally, to see (26) observe that both $\text{Re}(w + v(w))$ and $\text{Im}(w + v(w))$ are given by degree m polynomials of w , a standard Gaussian random vector. In [Janson, 1997, Theorem 6.7] it is shown that there exists a constant a_m depending only on m , such that if P is a polynomial of degree m and X is a standard normal random variable, then for every $t > 2$,

$$\mathbb{P} \left(|p(X)| > t \sqrt{\mathbb{E}[p(X)^2]} \right) \leq \exp \left(-a_m t^{2/m} \right)$$

Thus, since

$$\mathbb{E} [|\text{Re}(w + v(w)) \cdot x_i|^2], \mathbb{E} [|\text{Im}(w + v(w)) \cdot x_i|^2] \leq \mathbb{E} [1 + |w \cdot x_i|^{2m} + |v(w) \cdot x_i|^{2m}] \leq C_m,$$

the bound (26) follows. \square

We are finally ready to prove the existence of the complex neuron.

Proof.[Proof of Lemma 5] Consider the random variable

$$F = \mathbf{g} \cdot \mathbf{y} = \sum_{i=1}^n y_i g(x_i)$$

and set $W = \text{Re}(w + v(w))$ and $W' = \text{Im}(w + v(w))$. Lemma 6 gives

$$\mathbb{E} [F] \geq \frac{1}{2\sqrt{n\gamma^2}} \|\mathbf{y}\|^2.$$

Using Lemma 7 and Cauchy-Schwartz we may see that

$$\mathbb{E} [F^2] \leq \sum_{i=1}^n y_i^2 \mathbb{E}_{w,a} \left[\sum_{i=1}^n g(x_i)^2 \right] \leq C_m n \|\mathbf{y}\|^2.$$

Define $G = \mathbb{1}_{\{\exists i: |W \cdot x_i|, |W' \cdot x_i| \geq (4C_m \log(n))^{\frac{m}{2}}\}}$. A second application of Cauchy-Schwartz gives

$$\mathbb{E}[FG] \leq \sqrt{C_m n \|\mathbf{y}\|^2 \mathbb{E}[G]}.$$

Now, the estimate (26) and a union bound yields

$$\mathbb{E}[G] \leq n \exp(-4 \log(n)) \leq \frac{1}{n^3}.$$

Therefore,

$$\mathbb{E}[FG] \leq \frac{1}{n} C_m \|\mathbf{y}\|.$$

Combining this with the lower bound of $\mathbb{E}[F]$, we finally have

$$\mathbb{E}[F(1-G)] \geq \frac{1}{2\sqrt{n\gamma^2}} \|\mathbf{y}\|^2 - \frac{1}{n} C_m \|\mathbf{y}\| \geq \frac{1}{4\sqrt{n\gamma^2}} \|\mathbf{y}\|^2,$$

where the last inequality is valid as long as n is large enough. The claim now follows via taking a realization that exceeds the expectation. Since we might as well assume that the sample contains an orthonormal basis, (24) follows as well. \square

E.2 Approximating a complex neuron with ReLU activation

Our goal in this section is to prove the following lemma, showing that the complex polynomial can be essentially replaced by a ReLU. We write $\psi(t) = \text{ReLU}(t)$ and recall that $\varphi(t) = \frac{1}{\sqrt{m}} H_m(t)$.

Lemma 9 *For any $w, w' \in \mathbb{R}^d$, $z \in \mathbb{C}$ with $|z| = 1$ and $M > 0$, there exist a pair of random variables S, B and a random vector $W \in \mathbb{R}^d$ such that for any $x \in \mathbb{S}^{d-1}$ with $m(|w \cdot x| + |w' \cdot x|) \leq M$,*

$$\mathbb{E}[S\psi(W \cdot x - B)] = \frac{c_{z,m}}{M^m} \text{Re}(z \cdot \varphi(w \cdot x + iw' \cdot x)),$$

where $c_{z,m}$ depends only on m and z and there exists another constant c_m , such that

$$\frac{1}{c_m} \geq c_{z,m} \geq c_m. \quad (27)$$

Moreover,

$$|S| = 1, |B| \leq M \text{ almost surely,}$$

and

$$W = w + j \cdot w' \text{ for some } j \in \{0, 1, \dots, m\}.$$

Let us first see how to complete the proof of Theorem 5 using the combination of the above with Lemma 5.

Proof.[Proof of Theorem 5] Invoke Lemma 5 to obtain a function

$$g(x) = \text{Re}(z \cdot \varphi(x \cdot \tilde{w} + ix \cdot \tilde{w}'))$$

such that

$$\mathbf{y} \cdot \mathbf{g} \geq \frac{1}{2C_m \sqrt{n\gamma^2}} \|\mathbf{y}\|^2,$$

and such that for every $i \in [n]$,

$$|\tilde{w} \cdot x_i|, |\tilde{w}' \cdot x_i| \leq C_m \log(n)^{\frac{m}{2}}.$$

Set $M = 2C_m m \log(n)^{\frac{m}{2}}$, so that $m(|\tilde{w} \cdot x_i| + |\tilde{w}' \cdot x_i|) \leq M$. By Lemma 9, we may find σ, w, b , such that

$$|b|^2 \leq M^2, \|w\|^2 \leq m^2(\|\tilde{w}\| + \|\tilde{w}'\|)^2 \leq 4C_m m^2 d \log(n)^m, \quad |\sigma| = 1,$$

for which we define $f(x) = \sigma\psi(w \cdot x - b)$. The lemma then implies,

$$\mathbf{y} \cdot \mathbf{f} \geq \frac{c_m}{M^m} \mathbf{y} \cdot \mathbf{g} \geq \frac{c'_m}{M^m \sqrt{n\gamma^2}} \|\mathbf{y}\|^2,$$

and

$$\begin{aligned} \|\mathbf{f}\|^2 &= \sum_{i=1}^n (\psi(w \cdot x_i - b))^2 \leq 2 \sum_{i=1}^n (|w \cdot x_i|^2 + b^2) \\ &\leq 2M^2n + 2 \sum_{i=1}^n |w \cdot x_i|^2. \end{aligned}$$

By Lemma 9, $w = \tilde{w} + j \cdot \tilde{w}'$ for some $j = 0, \dots, m$. Hence, $|w \cdot x_i|^2 \leq 2m(|\tilde{w} \cdot x_i|^2 + |\tilde{w}' \cdot x_i|^2)$ and

$$\|\mathbf{f}\|^2 \leq 2M^2n + 4m \sum_{i=1}^n (|\tilde{w} \cdot x_i|^2 + |\tilde{w}' \cdot x_i|^2) \leq 10mM^2n.$$

The proof is concluded by substituting M . \square

It remains to prove Lemma 9. This is done in the next subsections.

E.2.1 On homogeneous polynomials

Since our aim is to approximate a polynomial by ReLU, we first find an appropriate polynomial basis to work with.

Lemma 10 *Any polynomial of the form $(x, y) \rightarrow \text{Re}(z \cdot (x + \mathbf{i}y)^m)$ has the form,*

$$\sum_{j=0}^m a_j (x + j \cdot y)^m.$$

Proof. Define

$$\mathcal{H}_m = \{p(x, y) | p \text{ is a degree } m \text{ homogeneous polynomial}\},$$

and

$$A_m = \{(x + j \cdot y)^m | j = 0, \dots, m\}.$$

It will suffice to show that A_m forms a basis for \mathcal{H}_m . The result will follow since $\text{Re}(z \cdot (x + \mathbf{i}y)^m)$ is clearly homogeneous. For $0 \leq j \leq m$, set $p_j = (x + j \cdot y)^m$, so that $p_j \in A_m$ and

$$p_j = \sum_{k=0}^m \binom{m}{k} j^k y^k x^{m-k}.$$

Note that the set $\{\binom{m}{k} y^k x^{m-k} | k = 0, \dots, m\}$ forms a basis for \mathcal{H}_m and in that basis p_j has coordinates $(1, j, \dots, j^m)$. Taking the Vandermonde determinant of the matrix whose columns are $\{p_j : j = 0, \dots, m\}$, we see that it must also be a basis for \mathcal{H}_m . \square

Corollary 1 *Let $w, w' \in \mathbb{R}^d$ and $z \in \mathbb{C}$, then we have the following representation:*

$$\text{Re}(z \cdot \varphi(w \cdot x + \mathbf{i}w' \cdot x)) = \sum_{j=0}^m p_{z,j}((w + jw') \cdot x),$$

where each $p_{z,j}$ is a polynomial of degree m , which depends continuously on z .

Proof. The representation is immediate from the previous lemma. To address the point of continuity, we write

$$\begin{aligned} \text{Re}(z \cdot \varphi(w \cdot x + \mathbf{i}w' \cdot x)) &= \text{Re}(z) \text{Re}(\varphi(w \cdot x + \mathbf{i}w' \cdot x)) - \text{Im}(z) \text{Im}(\varphi(w \cdot x + \mathbf{i}w' \cdot x)) \\ &= \text{Re}(z) \text{Re}(\varphi(w \cdot x + \mathbf{i}w' \cdot x)) + \text{Im}(z) \text{Re}(\mathbf{i} \cdot \varphi(w \cdot x + \mathbf{i}w' \cdot x)) \\ &= \sum_{j=0}^m (\text{Re}(z) p_{1,j}((w + jw') \cdot x) + \text{Im}(z) p_{i,j}((w + jw') \cdot x)). \end{aligned}$$

So, $p_{z,j}$ is a linear combination of $p_{1,j}$ and $p_{i,j}$, with coefficients that vary continuously in z . \square

E.2.2 ReLUs as universal approximators

Next, we show how ReLU functions might be used to universally approximate compactly supported functions.

Proposition 5 *Let $f : \mathbb{R} \rightarrow \mathbb{R}$ be twice differentiable and compactly supported on $[-M, M]$. Then, there exists a pair of random variables S, B , such that, for every $x \in [-M, M]$,*

$$\mathbb{E}[S\psi(x - B)] = \frac{f(x)}{\int |f''|},$$

and such that, almost surely $|B| \leq M$ and $|S| = 1$.

Proof. Observe that, when considered as a distribution, $\psi(x)'' = \delta_0$. Therefore, there exists a linear function L such that

$$f(x) + L(x) = \int_{-M}^M \psi(x - y)f''(y)dy.$$

$f''(x)$ is the second derivative of a compactly supported function which implies that $f(x) + L(x)$ is compactly supported as well. Hence, $L(x) \equiv 0$. Let B be the random variable whose density is $\frac{|f''|}{\int_{-M}^M |f''|}$ and set $S = \text{sign}(f''(B))$. We now have

$$\mathbb{E}[S\psi(x - B)] = \frac{\int_{-M}^M \psi(x - y)f''(y)dy}{\int_{-M}^M |f''|} = \frac{f(x)}{\int |f''|}.$$

□

E.2.3 Completing the proof of Lemma 9

Set χ_M to be a bump function for the interval $[-M, M]$. That is,

- $\chi_M : \mathbb{R} \rightarrow \mathbb{R}$ is smooth.
- $0 \leq \chi_M \leq 1$.
- $\chi_M(x) = 1$ for $x \in [-M, M]$.
- $\chi_M(x) = 0$ for $|x| > 2M$.

By Corollary 1, for any $w, w' \in \mathbb{R}^d, z \in \mathbb{C}$ we have the representation

$$\text{Re}(z \cdot \varphi(w \cdot x + iw' \cdot x)) \chi_M(|w \cdot x| + m|w' \cdot x|) = \sum_{j=0}^m p_{z,j}((w + jw') \cdot x) \chi_M(|w \cdot x| + m|w' \cdot x|). \quad (28)$$

Proof.[Proof of Lemma 9] Define $X = \{x \in \mathbb{S}^{d-1}; m(|w \cdot x| + |w' \cdot x|) \leq M\}$. Observe that for all $x \in X$,

$$\text{Re}(\varphi(w \cdot x + iw' \cdot x)) = \text{Re}(\varphi(w \cdot x + iw' \cdot x)) \chi_M(m(|w \cdot x| + |w' \cdot x|)).$$

Moreover, if $j = 0, \dots, m$, then $\chi_M((w + jw') \cdot x) = 1$, as well. By invoking Proposition 5 we deduce that for every $j = 0, \dots, m$, there exists a pair of random variables S_j, B_j and a constant $c_{z,j} > 0$ depending only on j, m and z , such that

$$\mathbb{E}[S_j \psi((w + jw') \cdot x - B_j)] = \frac{c_{z,j}}{M^m} p_{z,j}((w + jw') \cdot x) \chi_M((w + jw') \cdot x), \quad \forall x \in X,$$

Here we have used the fact that if p_j is one of the degree m polynomials in the decomposition (28), then there exist some constants $C'_{z,j}, C_{z,j} > 0$, for which

$$C'_{z,j} M^m \leq \int_{-M}^M |p''_{z,j}| \leq \int_{-2M}^{2M} |p''_{z,j}| \leq C_{z,j} M^m.$$

We now set J to be a random index from the set $\{0, \dots, m\}$ with

$$\mathbb{P}(J = j) = \frac{c_{z,j}^{-1}}{\sum_{j'} c_{z,j'}^{-1}}.$$

If we set $c_{z,m} = \frac{1}{\sum_{j'} c_{z,j'}^{-1}}$, and $S := S_J, B = B_J, W = w + Jw'$ it follows from (28) that

$$\begin{aligned} \mathbb{E}[S\psi(W \cdot x - B)] &= \frac{c_{z,m}}{M^m} \sum_{j=0}^m p_{z,j}((w + jw') \cdot x) \chi_M((w + jw') \cdot x) \\ &= \frac{c_{z,m}}{M^m} \operatorname{Re}(z \cdot \varphi(w \cdot x + \mathbf{i}w' \cdot x)) \chi_M(m(|w \cdot x| + |w' \cdot x|)). \end{aligned}$$

Finally since, by Corollary 1, $c_{z,m}$ depends continuously on z , a compactness argument implies (27). \square