

The Connection Between Approximation, Depth Separation and Learnability in Neural Networks

Eran Malach*

School of Computer Science, Hebrew University

ERAN.MALACH@MAIL.HUJI.AC.IL

Gilad Yehudai*

Weizmann Institute of Science

GILAD.YEHUDAI@WEIZMANN.AC.IL

Shai Shalev-Shwartz

School of Computer Science, Hebrew University

SHAIS@CS.HUJI.AC.IL

Ohad Shamir

Weizmann Institute of Science

OHAD.SHAMIR@WEIZMANN.AC.IL

Editors: Mikhail Belkin and Samory Kpotufe

Abstract

Several recent works have shown separation results between deep neural networks, and hypothesis classes with inferior approximation capacity such as shallow networks or kernel classes. On the other hand, the fact that deep networks can efficiently express a target function does not mean that this target function can be learned efficiently by deep neural networks. In this work we study the intricate connection between learnability and approximation capacity. We show that learnability with deep networks of a target function depends on the ability of simpler classes to approximate the target. Specifically, we show that a necessary condition for a function to be learnable by gradient descent on deep neural networks is to be able to approximate the function, at least in a weak sense, with shallow neural networks. We also show that a class of functions can be learned by an efficient statistical query algorithm if and only if it can be approximated in a weak sense by some kernel class. We give several examples of functions which demonstrate depth separation, and conclude that they cannot be efficiently learned, even by a hypothesis class that can efficiently approximate them.

1. Introduction

The empirical success of deep networks has inspired a large number of theoretical works trying to understand what properties of deep neural networks make them so powerful. From a theoretical perspective, the success of deep networks is often attributed to their *approximation capacity*. Deep networks can efficiently implement arbitrary Boolean circuits, and thus can efficiently compute anything that can be efficiently computed by Turing machines. Therefore, in terms of expressive power, deep neural networks are the ultimate choice of hypothesis class.

In contrast, other hypothesis classes studied in the literature have inferior approximation capacity. For example, Kernel methods (i.e., linear functions in RKHS space) can approximate arbitrary functions, at the cost of having an exponential dimension or margin complexity (Rahimi and Recht, 2008; Sun et al., 2018). Shallow (two-layer) neural networks can also approximate almost any target function (e.g. Cybenko (1989); Leshno et al. (1993)), although possibly using an exponentially large number of neurons.

* equal contribution

Indeed, several theoretical studies have demonstrated *separation results*: explicit constructions of functions or function families that can be expressed using deep networks, but cannot be approximated using shallow networks or kernel predictors of reasonably bounded size. For example, [Eldan and Shamir \(2016\)](#); [Safran and Shamir \(2017\)](#); [Daniely \(2017\)](#) showed separation results between depth-2 and depth-3 neural networks while [Telgarsky \(2016\)](#) showed separation between depth n and depth $n^{1/3}$ neural networks when the input dimension is constant. The works of [Daniely and Malach \(2020\)](#); [Yehudai and Shamir \(2019\)](#); [Kamath et al. \(2020\)](#); [Allen-Zhu and Li \(2019\)](#) show separation results between neural networks and kernel methods.

With that said, all the above results only show an analysis of the approximation capacity of the hypothesis classes. This is unsatisfactory, since the fact that a certain hypothesis class can express some target class does not mean that there is an efficient algorithm that can learn it. Namely, given access to limited computational resources, we would hope to use hypothesis classes for which we have efficient algorithms that can recover the best hypothesis within the class.

In some sense, the point of depth separation results is to argue that depth is beneficial by showing function classes that cannot be efficiently expressed without depth. In this work we show, perhaps surprisingly, that essentially *there are no* target functions which are both efficiently learnable, and "truly deep" in the sense that shallow networks (or kernel classes) cannot even weakly approximate them. More concretely, we explore the intricate connection between approximation capacity and efficient learnability. We first define the notion of *weak approximation* (Definition 3), namely, a hypothesis class weakly approximates a target class if for every target function there exists some hypothesis that approximates the target slightly better than the trivial predictor. We then show, in different settings and for different learning algorithms, that there is a dependence between the success of the algorithm in learning a target class, and having weak approximation using a "simple" class. That is, we show that problems which are hard to weakly approximate using "simple" classes (e.g. shallow neural networks) are also hard to learn using the more "complex" class (e.g. deep neural network). The above is true, even for target classes that can be exactly represented by the "complex" hypothesis class.

Our main contributions are as follows:

1. **Gradient descent:** A target class of functions cannot be learned by gradient descent on deep neural networks, if 3-layer neural network cannot weakly approximate it.
2. **SQ algorithms:** A target class of functions can be learned by deep neural networks using any statistical query algorithm if and only if they can be weakly approximated by a kernel class of functions.
3. We consider two known examples of target classes which separate between hypothesis classes, and as a corollary of the previous results, we get that these target classes cannot be learned, even by hypothesis classes that can perfectly represent them. Namely, (1) Telgarski's triangle function [Telgarsky \(2016\)](#) cannot be learned using gradient descent; (2) Parity functions cannot be learned by an efficient SQ-algorithm.
4. We show a specific target class that separates between 2-layer and 3-layer neural networks, and prove that this target class cannot be efficiently learned by any SQ-algorithm.

These results show that the power of deep neural networks to approximate target functions is not enough. In order for deep networks to learn some target function, simpler models such as shallow

networks or kernel classes should be able to approximate the target function, at least in a weak sense.

In [Malach and Shalev-Shwartz \(2019a\)](#) it was conjectured that a target class which cannot be approximated by a shallow neural network cannot be learned using gradient-based methods, even when using deep neural networks. Here we give a partial positive answer to this conjecture, for the specific case of learning with gradient descent, and for the approximation capabilities of 3-layer neural networks.

1.1. Related Works

Hardness of learning with GD. Several recent works have shown specific scenarios where gradient-based methods fails. In [Shalev-Shwartz et al. \(2017\)](#); [Shamir \(2018\)](#) several of failures of gradient descent are shown, including distributions that are hard to learn, and certain hard target functions. In [Yehudai and Shamir \(2020\)](#) it was shown that even for the case of learning the simplest neural network, containing a single neuron, and in the realizable case, there are distributions and activation functions which are hard to learn with gradient methods. In [Malach and Shalev-Shwartz \(2019a\)](#) it was shown that for certain fractal distribution, learnability using gradient descent on deep neural networks depends on whether the target can be approximated using shallow networks. Our work can be seen as a generalization of this work to a much larger class of distributions, and to different learning setups (i.e. SQ-algorithms).

Separation Results. In [Telgarsky \(2016\)](#), a family of target functions was introduced that can be realized by depth- n neural networks with polynomial width, but cannot be realized by depth- $n^{1/3}$ neural networks unless the width is exponential. Several works [Yehudai and Shamir \(2019\)](#); [Allen-Zhu and Li \(2019\)](#); [Kamath et al. \(2020\)](#) have shown separation between kernel methods and shallow neural networks. In particular, in [Yehudai and Shamir \(2019\)](#) it is shown that kernel methods (including NTK) cannot efficiently approximate a single ReLU neuron, while this problem can be learned with gradient methods using 2-layer network (e.g. [Yehudai and Shamir \(2020\)](#); [Frei et al. \(2020\)](#)).

Several works have shown target functions that can be well approximated by 3-layer neural networks, while they cannot be approximated by 2-layer networks unless their width is exponential. Such works include [Eldan and Shamir \(2016\)](#); [Safran and Shamir \(2017\)](#); [Safran et al. \(2019\)](#) where isotropic functions are considered, [Daniely \(2017\)](#) where a composition of inner product with complex function which cannot be approximated by low degree polynomials, and [Malach and Shalev-Shwartz \(2019b\)](#) where boolean functions are considered. In [Vardi and Shamir \(2020\)](#) it is shown that there are natural proof barriers for proving such depth separation results for depth larger than 4.

SQ results. Since its introduction in a seminal work by Kearns [Kearns \(1998\)](#), the statistical-query (SQ) framework has been extensively studied in various works. Unlike standard PAC learning, where the learner has access to a set of sampled examples, in SQ learning the learner can use statistical properties of the data, but not individual examples. These statistical properties are provided via access to an oracle, which given some query on the distribution, returns an approximate evaluation of the query. In his original work, Kearns demonstrated that an SQ algorithm can be easily adapted to a noise-robust learning algorithm [Kearns \(1998\)](#). The work of [Blum et al. \(1994\)](#) introduced the SQ-dimension, a statistical measure of the target class, that can be used to characterize weak learning with statistical-queries. Such characterization has been extended to other variants of the SQ

framework, including strong learning and distribution-free learning [Feldman \(2012\)](#); [Simon \(2007\)](#); [Szörényi \(2009\)](#); [Feldman and Kanade \(2012\)](#). Importantly, the SQ framework has been used to derive lower bounds on complexity of learning various problems, for example learning parities [Kearns \(1998\)](#) and neural networks [Goel et al. \(2020\)](#).

2. Definitions and Notations

2.1. Preliminaries

We denote by \mathcal{X}_n the input space and by \mathcal{Y} the label space. We denote vectors in bold. We denote the sign function as $\text{sign}(x) = 1$ if $x \geq 0$ and -1 otherwise. We focus on binary classification tasks, so $\mathcal{Y} = \{\pm 1\}$. We consider two classes of functions:

- The **target class**, denoted by \mathcal{F}_n , which is a class of functions from \mathcal{X}_n to \mathcal{Y} that labels the underlying distribution, and which the learner needs to approximate.
- The **hypothesis class**, denoted by \mathcal{H}_n , which is a class of functions from \mathcal{X}_n to \mathbb{R} from which the learner can choose its hypothesis.

We denote by $\ell : \mathcal{Y} \times \mathbb{R} \rightarrow \mathbb{R}$ our loss function. Since we consider classification tasks, we assume throughout the paper that ℓ is the hinge-loss, namely $\ell(y, \hat{y}) = \max\{1 - y\hat{y}, 0\}$.

For some distribution \mathcal{D} over \mathcal{X}_n and some target function $f : \mathcal{X}_n \rightarrow \mathcal{Y}$ ($f \in \mathcal{F}_n$) we denote by $f(\mathcal{D})$ the distribution over $\mathcal{X}_n \times \mathcal{Y}$ where $\mathbf{x} \sim \mathcal{D}$ and $y = f(\mathbf{x})$. For some hypothesis $h : \mathcal{X}_n \rightarrow \mathbb{R}$ ($h \in \mathcal{H}_n$) we denote the loss of h on the distribution $f(\mathcal{D})$ by:

$$L_{f(\mathcal{D})}(h) = \mathbb{E}_{(\mathbf{x}, y) \sim f(\mathcal{D})} \ell(y, h(\mathbf{x})) = \mathbb{E}_{\mathbf{x} \sim \mathcal{D}} \ell(f(\mathbf{x}), h(\mathbf{x}))$$

We say that a function $r : \mathbb{N} \rightarrow \mathbb{N}$ is super-polynomial if for every polynomial p we have $\lim_{n \rightarrow \infty} p(n)/r(n) = 0$. For $\mathbf{x}_0 \in \mathbb{R}^d$ and $r > 0$ we define the ball of radius r around \mathbf{x}_0 as $B_r(\mathbf{x}_0) = \{\mathbf{x} \in \mathbb{R}^d : \|\mathbf{x} - \mathbf{x}_0\| \leq r\}$. We denote by $[n]$ the set $\{1, \dots, n\}$ for $n \in \mathbb{N}$.

We will use the following definition of neural networks with bounded width:

Definition 1 *Let $d \in \mathbb{N}$ be the data input dimension. We define a neural network of depth k and width at most p as $h : \mathbb{R}^d \rightarrow \mathbb{R}$ with $h(\mathbf{x}) = \mathbf{h}^{(k)} \circ \dots \circ \mathbf{h}^{(1)}(\mathbf{x})$ where:*

- $\mathbf{h}^{(1)} = \sigma(W^{(1)}\mathbf{x} + \mathbf{b}^{(1)})$ for $W^{(1)} \in \mathbb{R}^{p_1 \times d}$, $\mathbf{b}^{(1)} \in \mathbb{R}^{p_1}$
- $\mathbf{h}^{(i)} = \sigma(W^{(i)}\mathbf{h}^{(i-1)} + \mathbf{b}^{(i)})$ for $W^{(i)} \in \mathbb{R}^{p_i \times p_{i-1}}$, $\mathbf{b}^{(i)} \in \mathbb{R}^{p_i}$ for $i = 2, \dots, k-1$
- $\mathbf{h}^{(k)} = W^{(k)}\mathbf{h}^{(k-1)} + \mathbf{b}^{(k)}$ for $W^{(k)} \in \mathbb{R}^{1 \times p_{k-1}}$, $\mathbf{b}^{(k)} \in \mathbb{R}^1$

where $\sigma : \mathbb{R} \rightarrow \mathbb{R}$ is some non-linear function, and $p_i \leq p$ for all i .

2.2. Weak Approximation and Weak Dependence

Consider the problem of learning a distribution labeled by some target class \mathcal{F} , using a learning algorithm that can output a hypothesis in \mathcal{H} . The goal of the learning algorithm is to find a hypothesis $h \in \mathcal{H}$ that minimizes the loss $L_{f(\mathcal{D})}(h)$, when given access to examples from a distribution $f(\mathcal{D})$ for some $f \in \mathcal{F}$.

We will consider in the paper two problems which are connected. First, we want to understand whether the fact that the algorithm is forced to output a hypothesis in \mathcal{H} limits its ability to *approximate* \mathcal{F} . What we could hope for is that the hypothesis class can be expressive enough in order to approximate the target up to some small accuracy:

Definition 2 Let $\mathcal{H} = \{\mathcal{H}_n\}_{n \in \mathbb{N}}$ be a sequence of hypothesis classes, let $\mathcal{F} = \{\mathcal{F}_n\}_{n \in \mathbb{N}}$ be a sequence of target classes and let $\mathcal{D} = \{\mathcal{D}_n\}_{n \in \mathbb{N}}$ be a sequence of distributions over $\{\mathcal{X}_n\}_{n \in \mathbb{N}}$. We say that \mathcal{H} ϵ -**approximates** $(\mathcal{F}, \mathcal{D})$ if there exists $\epsilon \in [0, 1)$ such that for all $n \in \mathbb{N}$:

$$\sup_{f \in \mathcal{F}_n} \inf_{h \in \mathcal{H}_n} L_{f(\mathcal{D}_n)}(h) \leq \epsilon.$$

We say $(\mathcal{F}, \mathcal{D})$ is **realizable** by \mathcal{H} if \mathcal{H} can 0-approximate it.

Note that in the above definition we did not require ϵ to be small, just that it's a constant which does not depend on n . In some cases we cannot guarantee realizability, or even ϵ -approximation for a constant ϵ . The minimal requirement in this case is that the hypothesis class can approximate the target just a bit better than the trivial approximation:

Definition 3 Let $\mathcal{H} = \{\mathcal{H}_n\}_{n \in \mathbb{N}}$ be a sequence of hypothesis classes, let $\mathcal{F} = \{\mathcal{F}_n\}_{n \in \mathbb{N}}$ a sequence of target classes and let $\mathcal{D} = \{\mathcal{D}_n\}_{n \in \mathbb{N}}$ be a sequence of distributions over $\{\mathcal{X}_n\}_{n \in \mathbb{N}}$. We say that \mathcal{H} **weakly approximates** \mathcal{F} with respect to \mathcal{D} , if there exists some polynomial p such that for all $n \in \mathbb{N}$:

$$\sup_{f \in \mathcal{F}_n} \inf_{h \in \mathcal{H}_n} L_{f(\mathcal{D}_n)}(h) \leq 1 - 1/|p(n)|$$

Since we use the hinge loss, the loss on the zero hypothesis (i.e. output 0 for every input) is exactly 1. This means that the weak approximation requirement in this case is that the hypothesis class can approximate the target better than the trivial classifier, at least up to an inverse polynomial.

The main goal of this paper is to explore the relation between learnability and approximation. To do so, we define the notion of *weak dependence* between some algorithm \mathcal{A} and a hypothesis class \mathcal{H} :

Definition 4 Let \mathcal{A} be a learning algorithm and let $\mathcal{H} = \{\mathcal{H}_n\}_{n \in \mathbb{N}}$ be a sequence of hypothesis classes. We say that \mathcal{A} **weakly depends** on \mathcal{H} , if every class-distribution pair $(\mathcal{F}, \mathcal{D})$ that cannot be weakly approximated by \mathcal{H} , cannot be efficiently learned by \mathcal{A} .

We leave the exact definition of efficient learnability for the next sections, as it depends on the specific setting of learning that we consider. Note that the algorithm \mathcal{A} does not necessarily output a hypothesis from \mathcal{H} . In general, \mathcal{A} will output a hypothesis from a "more expressive" class than \mathcal{H} , e.g. where \mathcal{H} consists of 3-layer neural networks while \mathcal{A} outputs a deep neural network.

Clearly, if some learning algorithm \mathcal{A} outputs a hypothesis from some class \mathcal{H} , then in order for \mathcal{A} to succeed in learning, it must hold that \mathcal{H} can (weakly) approximate the target class \mathcal{F} . So, any algorithm \mathcal{A} which outputs a hypothesis from a class \mathcal{H} *weakly depends* on the hypothesis class of its output. However, it turns out that some algorithms weakly depend on hypothesis classes that are very different, and sometimes much "weaker", than the class that is being learned by the algorithm. We next show some notable examples of such dependencies.

We note that in all the above definitions we considered a sequence of target classes \mathcal{F} parameterized by some parameter n . In the literature, there are several kinds of depth separation results, where some parameter n tends to infinity, and this parameter corresponds to some property of the problem. For example, in [Eldan and Shamir \(2016\)](#); [Safran and Shamir \(2017\)](#) n is the input dimension, while in [Telgarsky \(2016\)](#) n is the depth of the network. A major benefit of our terminology is that we give a single definition for both kinds of depth separation.

3. Gradient Descent Weakly Depends on Shallow Neural-Networks

In this section we focus on the gradient descent algorithm. We show that under certain technical assumptions, the gradient descent algorithm applied to *deep* neural networks weakly depends on shallow 3-layer neural networks.

First, let us define the algorithm that is being used. Let $g_\theta(\mathbf{x}) : \mathbb{R}^d \rightarrow \mathbb{R}$ be some function parameterized by a vector $\theta \in \mathbb{R}^r$, let $f : \mathbb{R}^d \rightarrow \mathbb{R}$ be a target function and \mathcal{D} a distribution over \mathbb{R}^d . Suppose we initialize the parameter vector at $\theta_0 \in \mathbb{R}^r$, for a learning rate $\eta > 0$ the **gradient descent** algorithm iteratively computes θ_t by the following rule:

$$\theta_t = \theta_{t-1} - \eta \nabla_{\theta_{t-1}} L_{f(\mathcal{D})}(g_{\theta_{t-1}}).$$

For example, we can think of $g_\theta(\mathbf{x})$ as a neural network with r parameters, taking as input d -dimensional data. We define when a class of functions is *not* weakly learnable by gradient descent in the following way:

Definition 5 Let $\mathcal{F} = \{f_n : \mathbb{R}^d \rightarrow \mathbb{R}\}_{n \in \mathbb{N}}$ a sequence of target classes, \mathcal{D} a distributions over $\mathcal{X} \subseteq \mathbb{R}^d$, $g = \{g_\theta^n(\mathbf{x}) : \mathcal{X} \rightarrow \mathbb{R}\}_{n \in \mathbb{N}}$ a sequence of functions parameterized by a vector $\theta \in \mathbb{R}^{q(n)}$ where $q(n)$ is some polynomial, and $\theta_0 = \{\theta_0^n \in \mathbb{R}^{q(n)}\}_{n \in \mathbb{N}}$ be a sequence of initialization points. Let $T(n), \eta(n) : \mathbb{N} \rightarrow \mathbb{R}$ be two polynomials. We say that \mathcal{F} is **not weakly learnable by gradient descent** with respect to \mathcal{D} using the functions g at initialization θ_0 , if there exists a super polynomial function $\alpha : \mathbb{N} \rightarrow \mathbb{N}$ such that for every $n \in \mathbb{N}$, running gradient descent on $g_{\theta_0^n}^n$ for $T(n)$ iterations, and any learning rate $\eta \leq \eta(n)$ we have that

$$L_{f_n(\mathcal{D})}(g_{\theta_0^n}^n) - L_{f_n(\mathcal{D})}(g_{\theta_T^n}^n) \leq \frac{1}{\alpha(n)}.$$

Here, for ease of notation we denote by $g_{\theta_0^n}^n$ the function g_θ^n initialized at θ_0^n , and by $g_{\theta_T^n}^n$ the function g_θ^n after T iterations of gradient descent, initialized at θ_0^n

The definition contains many parameters, but it is actually quite intuitive. In simple words, gradient descent is unable to learn a target class (and distribution) if after a polynomial number of iterations, the loss stays very close to the loss at the initialization. We give this definition in negation (“not weakly learnable”) because getting the loss away from its initialized value is a necessary condition to learn (assuming the loss is not so good at initialization), but it is not sufficient. Note that our only requirement of the learned function g is that the number of optimized parameters is polynomial, as optimizing a super-polynomial number of parameters is practically intractable. Although we give this definition with a constant learning rate, all the results can be readily extended to GD with variable learning rates, as long as they are all smaller than 1.

Next, we focus on the initialization scheme that is being used, we consider the following initialization of the parameter θ :

Definition 6 Let $g_\theta(\mathbf{x}) : \mathbb{R}^d \rightarrow \mathbb{R}$ be some function parameterized by $\theta \in \mathbb{R}^r$. We say that $\theta_0 \in \mathbb{R}^r$ is an *L-standard initialization* if there is a $\rho > 0$ such that for every $\theta \in B_\rho(\theta_0)$:

- $g_\theta(\mathbf{x})$ is an L -Lipschitz function of θ
- Each coordinate of $\nabla_\theta(g_\theta(\mathbf{x}))$ is an L -Lipschitz function of \mathbf{x} with $\sup_{\mathbf{x} \in [0,1]^d} \nabla_\theta(g_\theta^n)_i(\mathbf{x}) \leq L$ for every $i \in [r]$.

In Appendix A we show that Xavier initialization for depth- k neural networks is w.h.p an L -standard initialization, where $L = O(d)$ (d being the input dimension) and $\rho = \frac{1}{k}$.

Using the above definitions, we can show that gradient descent on a class of functions g weakly depends on the hypothesis class of 3-layer neural networks with bounded width.

Theorem 7 Let $\alpha : \mathbb{N} \rightarrow \mathbb{N}$ be a super-polynomial function, $d \in \mathbb{N}$ the input dimension for the data, and \mathcal{D} the uniform distribution on $[0, 1]^d$. Let $\mathcal{F} = \{f_n : [0, 1]^d \rightarrow \mathbb{R}\}_{n \in \mathbb{N}}$ be a sequence of functions, with $|f_n(\mathbf{x})| \leq C$ for all n for some constant $C > 0$. Assume that the sequence of hypothesis classes of 3-layer neural networks with ReLU activations and width at most $\alpha(n)^d \cdot 2d$ cannot weakly approximate \mathcal{F} in the sense that for any $n \in \mathbb{N}$, $\min_{h \in \mathcal{H}_n} L_{f_n}(h) \geq 1 - \alpha(n)^{-1}$. Let $g = \{g_\theta^n(\mathbf{x})\}_{n \in \mathbb{N}}$ be a function sequence parameterized by a vector $\theta \in \mathbb{R}^{p(n)}$ with a polynomial p , and assume we initialize at $\theta_0 \in \mathbb{R}^{p(n)}$ which is an L -standard initialization θ_0 with $\rho \geq \frac{1}{n}$. Then, running gradient descent for $T + 1$ iterations with learning rate η we have that:

$$L_{f_n}(g_{\theta_0}^n) - L_{f_n}(g_{\theta_{T+1}}^n) \leq \frac{21L^2C^2 \max\{1, \eta^2\} \sqrt{dp(n)}}{\alpha(n)} T^2,$$

The full proof can be found in Appendix B. The proof intuition is as follows: We show that the correlation between each coordinate of $\nabla_\theta g_\theta^n$ and f_n cannot be too large. We do that by using an approximation of Lipschitz functions with 3-layer neural networks. Using an argument regarding the optimization process of gradient descent and the Lipschitz initialization assumption on g_θ^n , we show that even after a polynomial number of iterations, the correlation between g_θ^n and f_n must remain small. This shows that the loss after a polynomial number of iterations cannot be too far away from the loss at the initialization.

Remark 8 Some remarks on applying a similar analysis for other variants of gradient-descent:

1. Our result holds for “exact” gradient-descent, where the update of the weights is done using the exact value of the population gradient. We note that, with some further assumptions on the initialization procedure, this result can be extended to “noisy” gradient-descent, where an i.i.d. noise is added to the weights at each iteration of gradient-descent. Indeed, we show that in the area of the initialization, the gradients are extremely small. Adding noise at each iteration is equivalent to re-initializing the model in a new (random) initialization, and then applying the extremely small gradients observed so far. So, if such random initialization is L -standard w.h.p., then a similar result can be derived for “noisy” gradient-descent.
2. We note that showing a similar result for SGD, where the gradients are calculated based on examples sampled from the distribution, is far more tricky. In fact, in [Abbe and Sandon \(2020\)](#)

it is shown that neural networks trained using SGD with a batch size of 1 can implement any poly-time PAC algorithm. Hence, showing hardness results in this setting is as complicated as showing hardness results on PAC learning, which involves relying on unproved computational hardness assumptions (e.g., cryptographic hardness).

We get the following as an immediate corollary from Thm. 7 by assuming that both η and T are at most polynomial in n :

Corollary 9 *Under the same assumptions as in Thm. 7, the function sequence \mathcal{F} is not weakly learnable by g with gradient descent that is initialized at any L -standard initialization with $\rho \geq \frac{1}{n}$, and run for a polynomial number of iterations, with a learning rate at most polynomial in n .*

Remark 10 *A couple of remarks about the assumptions of Thm. 7:*

1. *The notion of weak learnability in the theorem is the same as in Definition 3. Here we omitted the $\max_{f \in \mathcal{F}_n}$ since for any n , \mathcal{F}_n contains only a single function. Also, in the theorem we explicitly used the assumption that there is a super polynomial function which bounds the approximation.*
2. *Definition 6 can be satisfied by neural networks with differentiable activations, but not with the ReLU activation. To apply the theorem specifically to ReLU activation would require revising Lemma 24 to the derivative of a neural network with ReLU activation. We believe this can be done if we assume a uniform distribution on $[0, 1]^d$, and leave it for future work.*
3. *The assumption that $\rho \geq \frac{1}{n}$ is weaker than assuming that ρ is a constant, since here we allowed the radius for which the initialization is well behaved to get smaller with n .*

One caveat of the theorem is the exponential dependence of network width in the dimension. This dependency is due to the approximation of high dimensional Lipschitz functions using shallow networks. It can be seen that a target class of functions \mathcal{F} which satisfies the requirement of the theorem, has a Lipschitz constant which is exponential in d . The unlearnability itself comes from the fact that the target function has a large Lipschitz constant, although in these results we treat the input dimension d as a constant, while the depth n varies. In Vardi et al. (2021) it is shown that it may not be possible to find a function with a similar depth separation property that has a polynomial Lipschitz constant.

With that said, we immediately get from Thm. 7 and Thm. 22 the following:

Corollary 11 *Let \mathcal{A} be the following learning algorithm: For any n , initialize a neural network with depth n and width $p(n)$ (for $p(n)$ polynomial) using standard Xavier initialization. Then, train the neural network with gradient descent with step size $\eta \leq 1$. Then \mathcal{A} weakly depends on the hypothesis class of 3-layer neural networks with super-polynomial width.*

In Sec. 5 we will use this corollary to give an example of a function that, although it can be realized by depth n neural networks, it cannot be learned by them.

4. Statistical-Query Algorithms Weakly Depend on Kernel Classes

In this section we relate learnability using Statistical-Query (SQ) algorithms and approximation using a kernel class. Specifically, we show that efficient learning in the SQ model weakly depends on kernel classes. We start by defining weak learnability in the SQ model, following the definitions of Kearns (1998). First, for some function f distribution \mathcal{D} and tolerance parameter $\tau > 0$, we define the statistical-query oracle $\text{SQ}_\tau(f, \mathcal{D})$ to be an oracle which accepts queries of the form $q : \mathcal{X} \times \mathcal{Y} \rightarrow [-1, 1]$ and returns some value v such that $|v - \mathbb{E}_{\mathbf{x} \sim \mathcal{D}} q(\mathbf{x}, f(\mathbf{x}))| \leq \tau$.

Definition 12 Let $\mathcal{F} = \{\mathcal{F}_n\}_{n \in \mathbb{N}}$ be a sequence of target classes and let $\mathcal{D} = \{\mathcal{D}_n\}_{n \in \mathbb{N}}$ be a sequence of distributions over $\{\mathcal{X}_n\}_{n \in \mathbb{N}}$. We say that \mathcal{F} is **weakly learnable** with respect to \mathcal{D} , if there exists a sequence of algorithms $\mathcal{A} = \{\mathcal{A}_n\}_{n \in \mathbb{N}}$ and polynomials p, q, r such that for every $f \in \mathcal{F}_n$, the algorithm \mathcal{A}_n returns a hypothesis h such that:

$$L_{f(\mathcal{D}_n)}(h) \leq 1 - 1/|p(n)|$$

using at most $q(n)$ queries to $\text{SQ}_{1/r(n)}(f, \mathcal{D}_n)$.

Our main result in this section shows the weak dependence between SQ algorithms and the class of functions over a polynomial-size kernel space. We define a polynomial-size kernel class as follows:

Definition 13 A sequence of hypothesis classes $\mathcal{H} := \{\mathcal{H}_n\}_{n \in \mathbb{N}}$ is a **polynomial-size kernel class** if there exist polynomials p, q , and a sequence of mappings $\Psi_n : \mathcal{X}_n \rightarrow [-1, 1]^{p(n)}$ such that:

$$\mathcal{H}_n = \{\mathbf{x} \mapsto \langle \Psi_n(\mathbf{x}), \mathbf{w} \rangle : \|\mathbf{w}\|_2 \leq q(n)\}$$

Although this class is significantly less expressive than neural networks, here we show that it is possible to weakly learn a target class of functions with any SQ algorithm if and only if it can be weakly approximated by the kernel class:

Theorem 14 Let $\mathcal{F} = \{\mathcal{F}_n\}_{n \in \mathbb{N}}$ a sequence of target classes and let $\mathcal{D} = \{\mathcal{D}_n\}_{n \in \mathbb{N}}$ be a sequence of distributions over $\{\mathcal{X}_n\}_{n \in \mathbb{N}}$. Then, there exists an efficient statistical-query algorithm \mathcal{A} that weakly learns \mathcal{F} if and only if there exists a polynomial-size kernel class \mathcal{H} that weakly approximates \mathcal{F} with respect to \mathcal{D} .

The proof intuition is to show that a polynomial size kernel class can weakly approximate \mathcal{F} if and only if the SQ dimension of \mathcal{F} is polynomial. We also use a modified form of a known result which shows that the number of queries required for any SQ algorithm to learn a class of functions depend polynomially on the SQ dimension of this class of functions. The full proof can be found in Appendix C. The following corollary immediately follows from Thm. 14:

Corollary 15 Every SQ-algorithm \mathcal{A} **weakly depends** on some kernel class \mathcal{H} .

Note that SQ-algorithms are of course not limited to learning only kernel classes. In fact, almost any learning algorithm that has been studied in the machine learning literature can be implemented in the SQ framework. However, the above corollary states that although SQ-algorithms can potentially learn very complex function classes, they are limited to learning only classes that can be weakly approximated using a kernel class.

5. Strong Separation Between Hypothesis Classes

So far, we saw examples of *weak dependence* between algorithms and hypothesis classes. Our results suggest that in some cases, learning a complex hypothesis class (for example, deep neural networks) depends on having a weak approximation using a “simpler” class. So, from an optimization perspective, there is no gap between weak learnability of the “complex” class and weak approximation of the “simpler” class - if the “simple” class cannot (weakly) approximate, the “complex” class cannot be learned.

However, we can review the same question from an approximation perspective. Namely, we can consider distributions that can be expressed by the “complex” class and cannot be weakly approximated by the “simple” class. In this section we show that when we disregard optimization and consider only approximation capacity, we get such extreme gap between the “simple” and the “complex” class. In this case we say that there is a *strong separation* between the two classes. We define this formally as follows:

Definition 16 Let $\mathcal{H} = \{\mathcal{H}_n\}_{n \in \mathbb{N}}$, $\mathcal{H}' = \{\mathcal{H}'_n\}_{n \in \mathbb{N}}$ be two sequences of hypothesis classes, let $\mathcal{F} = \{\mathcal{F}_n\}_{n \in \mathbb{N}}$ a sequence of target classes and let $\mathcal{D} = \{\mathcal{D}_n\}_{n \in \mathbb{N}}$ be a sequence of distributions over $\{\mathcal{X}_n\}_{n \in \mathbb{N}}$. We say that $(\mathcal{F}, \mathcal{D})$ **strongly separates** \mathcal{H} from \mathcal{H}' , if \mathcal{H}' can ϵ -approximate it, but it cannot be weakly approximated by \mathcal{H} with respect to \mathcal{D} .

The construction of a function which strongly separates two hypothesis classes suggests that one class is significantly more expressive than the other. However, the fact that a hypothesis class can express a target function does not imply that a learning algorithm which outputs a hypothesis from this class can *learn* the target function.

Corollary 17 Let $\mathcal{H} = \{\mathcal{H}_n\}_{n \in \mathbb{N}}$, $\mathcal{H}' = \{\mathcal{H}'_n\}_{n \in \mathbb{N}}$ be two sequences of hypothesis classes, let $\mathcal{F} = \{\mathcal{F}_n\}_{n \in \mathbb{N}}$ a sequence of target classes and let $\mathcal{D} = \{\mathcal{D}_n\}_{n \in \mathbb{N}}$ be a sequence of distributions over $\{\mathcal{X}_n\}_{n \in \mathbb{N}}$. Assume that $(\mathcal{F}, \mathcal{D})$ strongly separates \mathcal{H} from \mathcal{H}' , then:

- If \mathcal{H} is the class of 3-layer neural networks with super-polynomial width, and $\mathcal{X}_n = [0, 1]^d$ for all n , then $(\mathcal{F}, \mathcal{D})$ is not learnable by gradient descent, even when it outputs a hypothesis from \mathcal{H}' .
- If \mathcal{H} is some polynomial size kernel class, and $\mathcal{X}_n = \{\pm 1\}^n$, then $(\mathcal{F}, \mathcal{D})$ is not learnable by any SQ algorithm.

The proof follows directly from the assumption that $(\mathcal{F}, \mathcal{D})$ cannot be weakly approximated by \mathcal{H} , and using either Thm. 14 or Thm. 7.

5.1. Examples of strong separation

In the following subsection we will consider known examples of functions which strongly separates hypothesis classes. As a result from the previous subsection, we will get that although the target functions can be realized by neural networks, they cannot be learned by standard learning algorithms.

TELGARSKI’S FUNCTION

Here we assume the input space is $\mathcal{X}_n = [0, 1]^d$ for all n . We define the following family of functions on d dimensional vectors:

$$f_n(\mathbf{x}) = \begin{cases} 1 & \exists t \in \mathbb{N}, x_1 \in [\frac{2t}{2^n}, \frac{2t+1}{2^n}] \\ -1 & \text{otherwise} \end{cases}$$

Let the input distribution \mathcal{D} to be the uniform distribution over \mathcal{X}_n , and let $\mathcal{F}_n := \{f_n\}$. This is the sign of Telgarsky’s triangle function in d dimensions (see [Telgarsky \(2016\)](#))¹.

Telgarsky shows that the function above cannot be approximated using a neural network with less than $n^{1/3}$ layers up to a constant accuracy. This result can be extended using the same methods to show that this function cannot be approximated using neural networks with less than $n^{1/2}$ layers, even up to *polynomially non-trivial* accuracy. This means that shallow neural network cannot approximate this function significantly better than the trivial predictor. Hence we have the following strong separation between deep and shallow neural networks with constant input dimension:

Theorem 18 *Let $k_1, k_2 : \mathbb{N} \rightarrow \mathbb{N}$ some functions such that $k_1(n) = n$ and $k_2(n) \leq \sqrt{n}$. Then $(\mathcal{F}, \mathcal{D})$ strongly separates the sign of polynomial-size depth- k_2 dimension- d networks from the sign of polynomial-size depth- k_1 networks. In particular, for any neural network N with depth $k_2(n)$ and width $p(n)$ we have that:*

$$L_{f_n(\mathcal{D}_n)}(N) \geq 1 - \frac{2^{\sqrt{n}}(2p(n))^{\sqrt{n}}}{2^n}.$$

This result is similar to the results from [Telgarsky \(2016\)](#). For completeness, we give here the full proof in our terminology of strong separation. The full proofs are in [Appendix D.1](#).

Note that in the above separation results the dimension is constant, and the separation is between varying depth neural networks. Also, we use the sign function of the network for ease of the proofs. It is possible to drop this assumption and show hardness of approximating f_n using the same proof techniques as in [Telgarsky \(2016\)](#).

Combining [Thm. 18](#) and [Corollary 17](#) we have shown that Telgarski’s function is not weakly learnable by gradient descent with standard initialization. This is true even if the algorithm is performed over deep neural networks which can represent Telgarski’s function.

We note that in our proofs we formally show that the gradient of the objective is exponentially small. One might consider normalized gradient descent, where the norm of the gradient at each iteration is normalized to some fixed value, hence avoiding the problem of small gradients. This solution is impractical, since for finite precision machines which are used in practice, the gradient is so small that it virtually equal to zero.

PARITY FUNCTIONS

Assume the input space is $\mathcal{X}_n = \{\pm 1\}^n$. The separation here will be with respect to the input dimension, as opposed to the previous example where the input dimension was fixed.

1. To be more precise, let $g(x)$ be Telgarski’s triangle function with 2^n jumps, defined using composition of ReLUs (Lemma 3.10(1) from [Telgarsky \(2016\)](#)). Then $f_n(x) = \text{sign}(g(x) - 0.5)$.

We will use **parity functions** over n -bits: For some subset $I \subseteq [n]$, denote by $f_I(\mathbf{x}) = \prod_{i \in I} x_i$, the parity over the bits of I , let $\mathcal{F}_n = \{f_I : I \subseteq [n]\}$ and \mathcal{D}_n the uniform distribution on \mathcal{X}_n . The following result was shown in several previous works (e.g. see [Daniely and Malach \(2020\)](#)):

Theorem 19 *Let \mathcal{H} be the class of depth-two networks with ReLU activation with polynomial width and weight magnitude, and let $\mathcal{F} = \{\mathcal{F}_n\}_{n \in \mathbb{N}}$ be the class of parity functions and $\{\mathcal{D}_n\}_{n \in \mathbb{N}}$ the uniform distribution of \mathcal{X}_n . Then, $(\mathcal{F}, \mathcal{D})$ strongly separates any polynomial-size kernel class from \mathcal{H} .*

For completeness, we give the full proof in Appendix D.2. Combining Thm. 19 and Corollary 17 shows that parity functions are not learnable by any efficient SQ algorithm.

6. Depth-2 and Depth-3 Neural Networks

Many previous works showed separation between depth-2 and depth-3 neural networks (e.g. [Eldan and Shamir \(2016\)](#); [Safran and Shamir \(2017\)](#); [Daniely \(2017\)](#); [Safran et al. \(2019\)](#)). These results usually construct a function that can be approximated up to very small accuracy with a 3-layer network and polynomial width, while it cannot be approximated by a 2-layer network, unless its width is exponential.

We believe that a necessary condition for a target class to be learned by an SQ-algorithm is to be weakly approximated by the class of 2-layer neural networks. That is, we conjecture that any rich enough² class of functions that cannot be approximated by shallow neural networks of polynomial width, cannot be learned by any SQ algorithm. If this conjecture is true, then any “rich” target class that strongly separates 2-layer from 3-layer neural networks, cannot be learned by an SQ-algorithm. This is true, even if the SQ-algorithm outputs a 3-layer (or deeper) neural network.

In what follows we will show a specific family of functions that strongly separates 2-layer from 3-layer neural networks, and is not weakly learnable by any SQ-algorithm.

6.1. An Example of Unlearnability and Separation For 2-layer and 3-layer Networks

Let $\mathcal{X}_n = \{\pm 1\}^n$, in this section the input space will be $\mathcal{X}_n \times \mathcal{X}_n$, and the output space is $\mathcal{Y} = \{\pm 1\}$. Fix some $\mathbf{z}' \in \mathcal{X}_n$, we define the following function³:

$$F_{\mathbf{z}'}(\mathbf{x}, \mathbf{z}) := \prod_{i=1}^n (x_i \vee z_i \vee z'_i) = \prod_{i \in I(\mathbf{z}')} (x_i \vee z_i),$$

where $I(\mathbf{z}') \subseteq [n]$ is the set of coordinates for which $z'_i = 1$. Define the following family of functions: $\mathcal{F} := \{F_{\mathbf{z}'} : \mathbf{z}' \in \mathcal{X}\} \subset \mathcal{Y}^{\mathcal{X} \times \mathcal{X}}$.

Theorem 20 *Let \mathcal{D}_n be a uniform distribution over $\mathcal{X}_n \times \mathcal{X}_n$. Then the family of function \mathcal{F} defined above is not weakly learnable by any SQ-algorithm.*

-
2. Note that in general, any single function f can be trivially “learned” by the SQ algorithm that always returns the function f . Therefore, for some class to be hard to learn using SQ algorithms, it needs to be “rich enough”, i.e. contain a large and diverse set of functions.
 3. If the input space is $\{0, 1\}^n$ this function would be equivalent to $\langle \mathbf{x}, \mathbf{z} \rangle \pmod 2$, where the inner products is taken over the coordinates for which $z'_i = 1$. We keep the input space as $\{\pm 1\}^n$ to be consistent with the previous sections.

The proof intuition is to show that class \mathcal{F} has large SQdim, and hence cannot be learned using SQ-algorithms, regardless of the chosen hypothesis class. We show this by defining the following inner-product: $\langle F_{\mathbf{z}'}, F_{\mathbf{z}''} \rangle = \mathbb{E}_{\mathbf{x}, \mathbf{z}} [F_{\mathbf{z}'}(\mathbf{x}, \mathbf{z}) F_{\mathbf{z}''}(\mathbf{x}, \mathbf{z})]$, and explicitly finding the vectors in \mathcal{X}_n that realize the SQ-dimension. The full proof can be found in Appendix E.1.

Suppose we fix some \mathbf{z}' such that its number of coordinates that equal to +1 is $\Omega(n)$. In this case, we can show that the function $F_{\mathbf{z}'}(\mathbf{x}, \mathbf{z})$ cannot be weakly approximated by 2-layer neural networks with polynomial width. On the other hand, for any \mathbf{z}' , the function $F_{\mathbf{z}'}(\mathbf{x}, \mathbf{z})$ can be realized by a 3-layer network with polynomial width and ReLU activation (see e.g. [Malach and Shalev-Shwartz \(2019b\)](#)). This is summed up in the following:

Theorem 21 *Let \mathcal{H} be a polynomial-size depth-three network class with ReLU activation. Let $F_n : \mathcal{X}_n \times \mathcal{X}_n$ with $F_n(\mathbf{x}, \mathbf{z}) = \prod_{i=1}^n (x_i \vee z_i)$ and let \mathcal{D}_n be the uniform distribution over $\mathcal{X}_n \times \mathcal{X}_n$. Then $(\{F_n\}, \{\mathcal{D}_n\})$ strongly separates any polynomial-size depth-two network class from \mathcal{H} .*

The full proof can be found in Appendix E.2. Although, for simplicity, we proved the theorem for \mathbf{z}' which equal to +1 in all the coordinates, the proof is the same for any \mathbf{z}' with $\Omega(n)$ coordinates which equal to +1.

To conclude, we have shown an example of a family of functions which strongly separates depth-2 from depth-3 neural networks, which cannot be learned by any SQ-algorithm, even if it outputs a 3-layer network which can represent the function.

7. Discussion and Future Work

In this work we have shown that there is an intricate connection between the approximation capabilities of "simple" hypotheses classes, and learning capabilities of certain algorithms, even when used on a "complex" hypothesis class that can perfectly represent the target. We have shown this connection appears in two learning setups: The first is that being able to successfully weakly approximate using a 3-layer neural is a necessary condition to successful learning with gradient descent. The second is that to be able to weakly learn a target class (even in a weak sense) by any SQ-algorithm is possible if and only if this target class can be weakly approximated by some kernel class.

We have also discussed two known examples of target classes which strongly separates hypothesis classes. As a consequence of the dependence between approximation and learnability, these target classes (namely, Telgarski's triangle function and parity functions) cannot be efficiently learned, even when a "complex" architecture which can represent the functions is used. Finally, we have shown a specific example for a target class which separates between 2-layer and 3-layer networks, that cannot be learned by any SQ-algorithm.

We note that our work does not cover the separation results from [Eldan and Shamir \(2016\)](#); [Safran and Shamir \(2017\)](#); [Daniely \(2017\)](#). This is because the target functions that are introduced seem to be weakly approximable by 2-layer networks (e.g. see Figure 1 in [Safran and Shamir \(2017\)](#)). It is an open question whether these function can be learned up to arbitrary accuracy using 3-layer (or deeper) neural networks.

Another interesting future direction is find more learning setups for which there is a dependence between approximation with a "simple" hypothesis class, and learning with a more "complex" hypothesis class. Such setups include other gradient methods, e.g. stochastic gradient descent, or momentum, and separation between shallow and deep neural networks where the input dimension is not constant.

ACKNOWLEDGMENTS

This research is supported by the European Research Council (TheoryDL project), and by European Research Council (ERC) grant 754705.

References

- Emmanuel Abbe and Colin Sandon. Poly-time universality and limitations of deep learning. *arXiv preprint arXiv:2001.02992*, 2020.
- Zeyuan Allen-Zhu and Yuanzhi Li. What can resnet learn efficiently, going beyond kernels? In *Advances in Neural Information Processing Systems*, pages 9017–9028, 2019.
- Avrim Blum, Merrick Furst, Jeffrey Jackson, Michael Kearns, Yishay Mansour, and Steven Rudich. Weakly learning dnf and characterizing statistical query learning using fourier analysis. In *Proceedings of the twenty-sixth annual ACM symposium on Theory of computing*, pages 253–262, 1994.
- George Cybenko. Approximation by superpositions of a sigmoidal function. *Mathematics of control, signals and systems*, 2(4):303–314, 1989.
- Amit Daniely. Depth separation for neural networks. In *Conference on Learning Theory*, pages 690–696, 2017.
- Amit Daniely and Eran Malach. Learning parities with neural networks. *arXiv preprint arXiv:2002.07400*, 2020.
- Simon Du and Wei Hu. Width provably matters in optimization for deep linear neural networks. In *International Conference on Machine Learning*, pages 1655–1664, 2019.
- Ronen Eldan and Ohad Shamir. The power of depth for feedforward neural networks. In *Conference on learning theory*, pages 907–940, 2016.
- Vitaly Feldman. A complete characterization of statistical query learning with applications to evolvability. *Journal of Computer and System Sciences*, 78(5):1444–1459, 2012.
- Vitaly Feldman and Varun Kanade. Computational bounds on statistical query learning. In *Conference on Learning Theory*, pages 16–1, 2012.
- Spencer Frei, Yuan Cao, and Quanquan Gu. Agnostic learning of a single neuron with gradient descent. *arXiv preprint arXiv:2005.14426*, 2020.
- Xavier Glorot and Yoshua Bengio. Understanding the difficulty of training deep feedforward neural networks. In *Proceedings of the thirteenth international conference on artificial intelligence and statistics*, pages 249–256, 2010.
- Surbhi Goel, Aravind Gollakota, Zhihan Jin, Sushrut Karmalkar, and Adam Klivans. Superpolynomial lower bounds for learning one-layer neural networks using gradient descent. In *International Conference on Machine Learning*, pages 3587–3596. PMLR, 2020.

- Pritish Kamath, Omar Montasser, and Nathan Srebro. Approximate is good enough: Probabilistic variants of dimensional and margin complexity. *arXiv preprint arXiv:2003.04180*, 2020.
- Michael Kearns. Efficient noise-tolerant learning from statistical queries. *Journal of the ACM (JACM)*, 45(6):983–1006, 1998.
- 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.
- Eran Malach and Shai Shalev-Shwartz. Is deeper better only when shallow is good? In *Advances in Neural Information Processing Systems*, pages 6429–6438, 2019a.
- Eran Malach and Shai Shalev-Shwartz. Learning boolean circuits with neural networks. *arXiv preprint arXiv:1910.11923*, 2019b.
- Ali Rahimi and Benjamin Recht. Uniform approximation of functions with random bases. In *2008 46th Annual Allerton Conference on Communication, Control, and Computing*, pages 555–561. IEEE, 2008.
- Itay Safran and Ohad Shamir. Depth-width tradeoffs in approximating natural functions with neural networks. In *International Conference on Machine Learning*, pages 2979–2987. PMLR, 2017.
- Itay Safran, Ronen Eldan, and Ohad Shamir. Depth separations in neural networks: What is actually being separated? *arXiv preprint arXiv:1904.06984*, 2019.
- Shai Shalev-Shwartz, Ohad Shamir, and Shaked Shammah. Failures of gradient-based deep learning. In *International Conference on Machine Learning*, pages 3067–3075, 2017.
- Ohad Shamir. Distribution-specific hardness of learning neural networks. *The Journal of Machine Learning Research*, 19(1):1135–1163, 2018.
- Hans Ulrich Simon. A characterization of strong learnability in the statistical query model. In *Annual Symposium on Theoretical Aspects of Computer Science*, pages 393–404. Springer, 2007.
- Yitong Sun, Anna Gilbert, and Ambuj Tewari. On the approximation properties of random relu features. *arXiv preprint arXiv:1810.04374*, 2018.
- Balázs Szörényi. Characterizing statistical query learning: simplified notions and proofs. In *International Conference on Algorithmic Learning Theory*, pages 186–200. Springer, 2009.
- Matus Telgarsky. Benefits of depth in neural networks. *arXiv preprint arXiv:1602.04485*, 2016.
- Gal Vardi and Ohad Shamir. Neural networks with small weights and depth-separation barriers. *arXiv preprint arXiv:2006.00625*, 2020.
- Gal Vardi, Daniel Reichman, Toniann Pitassi, and Ohad Shamir. Size and depth separation in approximating natural functions with neural networks. *Unpublished manuscript*, 2021.
- 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 6598–6608, 2019.

Appendix A. Discussion on L -standard initializations

The assumption of an L -standard initialization is required, although sometimes implicitly, in order to learn with gradient methods. The intuition for the first part of the definition is that moving the parameters of the functions by a small amount (e.g. doing a single gradient step) doesn't change the value of the function too drastically. Without this requirement, it will be hard to predict the outcome of even a single gradient step with a small step size. The second part of the definition is that the gradient is well behaved with respect to the data. This means that the gradient w.r.t two close data points is similar, without this requirement small perturbations of the data could also change the gradient direction too significantly. The boundness requirement is technical and can be obtained by bounding the data domain.

All of the requirements of L -standard initialization are satisfied by using standard Xavier initialization [Glorot and Bengio \(2010\)](#). We get a bound on the Lipschitz constant that depends on the norm of \mathbf{x} (which in our case can be bounded by d), and in a radius that depends on the depth of the network.

Theorem 22 *Let σ be a 1-Lipschitz activation function and let $d > 0$ be the data dimension. Suppose we initialize a neural network of depth k and width m using Xavier initialization, and assume that $m > k^2$. Then w.p $> 1 - e^{-\Omega(m/k)}$, this is an L -standard initialization with radius $\rho = \frac{1}{k}$*

In Xavier initialization, each weight entry in a weight matrix of width m is drawn from $\mathcal{N}(0, \frac{1}{m})$ (up to a constant factor). For a depth k width m network it will be easier to view this initialization as if each weight is drawn from a standard Gaussian distribution, and the network is multiplied by the normalization term $\frac{1}{m^{k/2}}$ (which is also done in [Du and Hu \(2019\)](#)).

For simplicity of the proof we assume the network doesn't have bias terms, and that the width of all the weight matrices is the same. The main technical part of the proof is using Lemma 6.1 from [Du and Hu \(2019\)](#) which uses tail bounds on certain random variables to bound the norm of the multiplication of random matrices where each entry is drawn from a standard Gaussian initialization.

Proof We can write the network as:

$$N(\mathbf{x}) = W_k \sigma(W_{k-1} \cdots \sigma(W_1 \mathbf{x}) \cdots),$$

where $\mathbf{x} \in \mathbb{R}^d$, $W_1 \in \mathbb{R}^{m \times d}$, $W_i \in \mathbb{R}^{m \times m}$ for $i = 2, \dots, k-1$ and $W_k \in \mathbb{R}^{m \times 1}$. For the first part of the definition, recall that the Lipschitz constant of a composition of functions can be bounded by the multiplication of the Lipschitz constants of the functions. Since σ is 1-Lipschitz, and a neural network is a composition of matrix multiplication with the activation, it is enough to show the Lipschitzness for the linear network function:

$$N'(\mathbf{x}) = W_k \cdots W_1 \mathbf{x}.$$

Let W_1, \dots, W_k be the weight matrices at initialization, and A_1, \dots, A_k be some perturbation matrices with $\max_i \|A_i\| \leq k$. Suppose for this part of the proof that $\|\mathbf{x}\| = 1$ (we will deal with \mathbf{x}

with other norms later). Denote by $W_{(-i_1, \dots, -i_\ell)}$ for $i_1, \dots, i_\ell \in \{1, \dots, k\}$ the multiplication of the matrices W_k, \dots, W_1 without the i_1, \dots, i_ℓ matrices. Then we have:

$$\begin{aligned} & \| (W_k + A_l) \cdots (W_1 + A_1) \mathbf{x} - W_k \cdots W_1 \mathbf{x} \| \\ & \leq \sum_{i=1}^k \| A_i W_{(-i)} \mathbf{x} \| + \sum_{i \neq j} \| A_i A_j W_{(-i, -j)} \mathbf{x} \| + \dots \end{aligned} \quad (1)$$

where at each sum there are ℓ indices for $\ell = 1, \dots, k$ for matrices which are left out of the multiplication. For a single sum we have:

$$\sum_{i_1 \neq \dots \neq i_\ell} \| A_{i_1} \cdots A_{i_\ell} \cdot W_{(-i_1, \dots, -i_\ell)} \mathbf{x} \| \leq \max_i \| A_i \|^\ell \cdot \sum_{i_1 \neq \dots \neq i_\ell} \| W_{(-i_1, \dots, -i_\ell)} \mathbf{x} \|^2$$

We use Lemma 6.1 from [Du and Hu \(2019\)](#) to get that w.p $> 1 - e^{-\Omega(m/k)}$ we have that $\| W_{(-i_1, \dots, -i_\ell)} \mathbf{x} \|^2 \leq 1.1 m^{\frac{k-\ell}{2}}$. Hence, we can bound the above term by:

$$\begin{aligned} 1.1 \cdot \max_i \| A_i \|^\ell \cdot m^{\frac{k-\ell}{2}} \binom{k}{\ell} & \leq 1.1 \cdot \frac{1}{k} \cdot m^{\frac{k-\ell}{2}} \cdot k^\ell \\ & \leq \frac{1.1}{k} m^{k/2} \left(\frac{k}{\sqrt{m}} \right)^\ell \leq \frac{1.1 m^{k/2}}{k} \end{aligned}$$

where we used the assumption that $m > k^2$. Since there are k such sums we have that the above can be bounded by $1.1 m^{k/2}$, dividing by the normalization term, we get that a Lipschitz constant of 1.1 with probability $> 1 - e^{-\Omega(m/k)}$. If $\| \mathbf{x} \| \neq 1$ then we can divide Eq. (1) by $\| \mathbf{x} \|^2$ to get a Lipschitz constant and follow the proof in the same manner to get a Lipschitz constant of $1.1 \| \mathbf{x} \|^2$. This finishes the first part of the definition.

For the second part, the Lipschitzness condition follows from the fact that

$$\frac{\partial N'}{\partial W_i} = (W_k \cdots W_{i+1})^\top (W_{i-1} \cdots W_1 \mathbf{x})^\top,$$

from Lemma 6.1 from [Du and Hu \(2019\)](#) and the same reasoning as in the proof of the first part. In this manner we get a Lipschitz constant of 1.1 w.h.p. For boundness condition, using the assumption that $\mathbf{x} \in [0, 1]^d$ we get that $\max_{\mathbf{x} \in [0, 1]^d} \| \mathbf{x} \|^2 = d$, hence we have that the supremum over \mathbf{x} of the gradient will be bounded by $1.1 d$ for every coordinate of the gradient. \blacksquare

Appendix B. Proof of Thm. 7

The following two lemmas will be necessary in order to approximate a Lipschitz function using a shallow neural network.

Lemma 23 *Let $\gamma > 0$ and $A := [a_1, b_1] \times \cdots \times [a_d, b_d] \subseteq \mathbb{R}^d$. Then there exists a 3-layer neural network $N(\mathbf{x})$ with depth 2 and width $2d$ such that $N(\mathbf{x}) = 1$ for $\mathbf{x} \in [a_1 + \gamma, b_1 - \gamma] \times \cdots \times [a_d + \gamma, b_d - \gamma]$, $N(\mathbf{x}) = 0$ for $\mathbf{x} \notin A$ and $|N(\mathbf{x})| \leq 1$ for all $\mathbf{x} \in \mathbb{R}^d$.*

4. Although the matrix W_k is of a different dimension, the proof can be easily extended to deal with this case.

Proof We define:

$$N(\mathbf{x}) = \sigma \left(1 - \frac{1}{\gamma} \sum_{i=1}^d \sigma(a_i + \gamma - x_i) - \frac{1}{\gamma} \sum_{i=1}^d \sigma(x_i - b_i + \gamma) \right),$$

it is a 3-layer neural network with width $2d$. Note that for every $i \in [d]$, if $x_i < a_i$ or $x_i > b_i$ then $N(\mathbf{x}) = 0$. Also, if $\mathbf{x} \in [a_1 + \gamma, b_1 - \gamma] \times \cdots \times [a_d + \gamma, b_d - \gamma]$ then $N(\mathbf{x}) = 1$, and finally for values of \mathbf{x} not specified above, $N(\mathbf{x})$ interpolates between 0 and 1, hence $|N(\mathbf{x})| \leq 1$ for these values. \blacksquare

Lemma 24 Let $h : [0, 1]^d \rightarrow \mathbb{R}$ an L -Lipschitz function with $\sup_{\mathbf{x} \in [0, 1]^d} |h(\mathbf{x})| \leq C$, and $n \in \mathbb{N}$. Then there is a 3-layer neural network $N(\mathbf{x})$ with width $n^d \cdot 2d$ such that $\int_{[0, 1]^d} |N(\mathbf{x}) - h(\mathbf{x})| d\mathbf{x} \leq \frac{2C + L\sqrt{d}}{n^d}$.

Proof First, we split the hypercube $[0, 1]^d$ into n^d smaller and equally sized hypercubes, in the following form: for every $i_1, \dots, i_d \in \{0, \dots, n-1\}$ we define the hypercube $\left[\frac{i_1}{n}, \frac{i_1+1}{n}\right] \times \cdots \times \left[\frac{i_d}{n}, \frac{i_d+1}{n}\right]$, there are n^d such hypercubes, each with volume n^{-d} . Denote these hypercubes as A_1, \dots, A_M for $M = n^d$. For each A_i , pick any $\mathbf{x}_i \in A_i$ and let $c_i = h(\mathbf{x}_i)$. For each $A_i = \left[\frac{i_1}{n}, \frac{i_1+1}{n}\right] \times \cdots \times \left[\frac{i_d}{n}, \frac{i_d+1}{n}\right]$ we use Lemma 23 with $\gamma = n^{-2d}$ to get a neural network $N_i(\mathbf{x})$ such that $N_i(\mathbf{x}) = 0$ for $\mathbf{x} \notin A_i$, $N_i(\mathbf{x})=1$ for $\mathbf{x} \in \left[\frac{i_1}{n} + n^{-2d}, \frac{i_1+1}{n} - n^{-2d}\right] \times \cdots \times \left[\frac{i_d}{n} + n^{-2d}, \frac{i_d+1}{n} - n^{-2d}\right]$, and $|N_i(\mathbf{x})| \leq 1$ for every $\mathbf{x} \in [0, 1]^d$. Then we have that:

$$\int_{A_i} |h(\mathbf{x}) - c_i \cdot N_i(\mathbf{x})| d\mathbf{x} \leq 2Cn^{-2d} + L\sqrt{d}n^{-2d},$$

where we used that inside A_i , if $N_i(\mathbf{x}) = c_i$ then $|h(\mathbf{x}) - c_i \cdot N_i(\mathbf{x})| \leq \frac{L\sqrt{d}}{n^d}$ since $h(\mathbf{x})$ is L -Lipschitz, otherwise $|h(\mathbf{x}) - c_i \cdot N_i(\mathbf{x})| \leq 2C$ but the area for which this happens is at most n^{-2d} . Define $N(\mathbf{x}) := \sum_{i=1}^M c_i N_i(\mathbf{x})$, then $N(\mathbf{x})$ is a 3-layer neural network with width $n^d \cdot 2d$, and we get that:

$$\int_{[0, 1]^d} |N(\mathbf{x}) - h(\mathbf{x})| d\mathbf{x} \leq \sum_{i=1}^M \int_{A_i} |N_i(\mathbf{x}) - h(\mathbf{x})| d\mathbf{x} \leq n^d \cdot (2Cn^{-2d} + L\sqrt{d}n^{-2d}) \leq \frac{2C + L\sqrt{d}}{n^d}. \quad \blacksquare$$

We are now ready to prove the main theorem:

Proof [Proof of Thm. 7] Fix $n \in \mathbb{N}$, in the proof we denote $\theta_0 := \theta_0^n$ and θ_t to be θ_0 after t iterations of gradient descent, and denote $\nabla_T = \nabla_{\theta_T} (L_{f_n}(g_{\theta_T}^n))$. We have that:

$$\begin{aligned} \|\theta_{T+1} - \theta_0\|^2 &= \|\theta_T - \eta \nabla_T - \theta_0\|^2 = \|\theta_T - \theta_0\|^2 + \eta^2 \|\nabla_T\|^2 - 2\eta \langle \nabla_T, \theta_T - \theta_0 \rangle \\ &\leq \|\theta_T - \theta_0\|^2 + \eta^2 \|\nabla_T\|^2 + 2\eta \|\nabla_T\| \cdot \|\theta_T - \theta_0\| \\ &\leq \cdots \leq \sum_{t=0}^T \eta^2 \|\nabla_t\|^2 + 2\eta \|\nabla_t\| \cdot \|\theta_t - \theta_0\| \end{aligned} \quad (2)$$

We will first bound the norm of the gradient at each iteration. Suppose that $\|\theta_t - \theta_0\| \leq \frac{1}{n}$, then, by the assumption on the initialization, each coordinate of $\nabla_{\theta_t} (g_{\theta_t}^n(\mathbf{x}))$ is an L -Lipschitz function from $[0, 1]^d$ to \mathbb{R} with $\sup_{\mathbf{x} \in [0, 1]^d} (\nabla_{\theta_t} (g_{\theta_t}^n(\mathbf{x})))_i \leq L$ for $i \in [p(n)]$. By Lemma 24 there is a 3-layer neural network $N_i : [0, 1]^d \rightarrow \mathbb{R}$ with width $\alpha(n)^d 2d$ such that

$$\int_{[0, 1]^d} |N_i(\mathbf{x}) - (\nabla_{\theta_t} g_{\theta_t}^n(\mathbf{x}))_i| d\mathbf{x} \leq \frac{4L + 2L\sqrt{d}}{\alpha(n)^d} \leq \frac{6L\sqrt{d}}{\alpha(n)^d}.$$

For every $i \in [p(n)]$ (each coordinate of $\nabla_{\theta_t} (g_{\theta_t}^n(\mathbf{x}))$) we have that:

$$\begin{aligned} & \left| \mathbb{E}_{\mathbf{x} \sim U([0, 1]^d)} [(\nabla_{\theta_t} g_{\theta_t}^n(\mathbf{x}))_i f_n(\mathbf{x})] \right| \\ & \leq \left| \mathbb{E}_{\mathbf{x} \sim U([0, 1]^d)} [N_i(\mathbf{x}) f_n(\mathbf{x})] \right| + \mathbb{E}_{\mathbf{x} \sim U([0, 1]^d)} [|(\nabla_{\theta_t} g_{\theta_t}^n(\mathbf{x}))_i - N_i(\mathbf{x})| \cdot |f_n(\mathbf{x})|] \\ & \leq \frac{1}{\alpha(n)} + \frac{6LC\sqrt{d}}{\alpha(n)^d} \leq \frac{7LC\sqrt{d}}{\alpha(n)}. \end{aligned} \quad (3)$$

Here we used the assumption that ReLU neural networks with at most $\alpha(n)^d 2d$ cannot weakly approximate f_n in the sense stated in the theorem, which for the case of hinge loss, means that the correlation between the two functions is bounded by $\alpha(n)^{-1}$.

Using Eq. (3) we can now bound the norm of the gradient at each iteration, under the assumption that $\|\theta_t - \theta_0\| \leq \frac{1}{n}$:

$$\begin{aligned} \|\nabla_t\|^2 &= \left\| \frac{\partial}{\partial \theta} \mathbb{E}_{\mathbf{x} \sim U([0, 1]^d)} [\max\{0, 1 - g_{\theta_t}^n(\mathbf{x}) f_n(\mathbf{x})\}] \right\|^2 \\ &\leq \sum_{i=1}^{p(n)} \left| \mathbb{E}_{\mathbf{x} \sim U([0, 1]^d)} (\nabla_{\theta_t} g_{\theta_t}^n(\mathbf{x}))_i f_n(\mathbf{x}) \right|^2 \leq \frac{7LC\sqrt{d} \cdot p(n)}{\alpha(n)}. \end{aligned} \quad (4)$$

Denote $a := \frac{7LC\sqrt{d} \cdot p(n)}{\alpha(n)}$ and assume w.l.o.g that $a < 1$ (otherwise, take a larger n), combining Eq. (2) and Eq. (4) we get:

$$\|\theta_{T+1} - \theta_0\|^2 \leq T\eta^2 a + 2\eta\sqrt{a} \sum_{t=0}^T \|\theta_t - \theta_0\| \leq \max\{1, \eta^2\} \left(Ta + 2\sqrt{a} \sum_{t=0}^T \|\theta_t - \theta_0\| \right).$$

We will show using induction that $\|\theta_{T+1} - \theta_0\|^2 \leq \max\{1, \eta^2\} 3aT^2$. For $T = 0$ it is clear. Assume for all $t \leq T$, then we have:

$$\begin{aligned} \frac{1}{\max\{1, \eta^2\}} \cdot \|\theta_{T+1} - \theta_0\|^2 &\leq Ta + 2\sqrt{a} \sum_{t=0}^T \sqrt{3at^2} \leq Ta + 2\sqrt{3a} \sum_{t=0}^T t \\ &\leq Ta + \frac{2\sqrt{3a}T^2}{2} \leq T^2 a + \sqrt{3a}T^2 \leq 3aT^2. \end{aligned} \quad (5)$$

Hence, if $\max\{1, \eta^2\}3aT^2 \leq \frac{1}{n}$ then the above applies since the function $g_{\theta_T}^n(x)$ is Lipschitz and bounded, this applies for all $T \leq \sqrt{\frac{\alpha(n)}{21LC \max\{1, \eta^2\} \sqrt{dp(n)}}}$. Let T be bounded as above, then:

$$\begin{aligned}
L_{f_n}(g_{\theta_0}^n) &= \mathbb{E}_{\mathbf{x} \sim U([0,1]^d)} [\max\{0, 1 - g_{\theta_0}^n(\mathbf{x})f_n(\mathbf{x})\}] \\
&= \mathbb{E}_{\mathbf{x} \sim U([0,1]^d)} [\max\{0, 1 - g_{\theta_{T+1}}^n(\mathbf{x})f_n(\mathbf{x}) + (g_{\theta_{T+1}}^n(\mathbf{x}) - g_{\theta_0}^n(\mathbf{x}))f_n(\mathbf{x})\}] \\
&\leq \mathbb{E}_{\mathbf{x} \sim U([0,1]^d)} [\max\{0, 1 - g_{\theta_{T+1}}^n(\mathbf{x})f_n(\mathbf{x})\}] + \mathbb{E}_{\mathbf{x} \sim U([0,1]^d)} \left[|(g_{\theta_0}^n(\mathbf{x}) - g_{\theta_{T+1}}^n(\mathbf{x}))f_n(\mathbf{x})| \right] \\
&\leq L_{f_n}(g_{\theta_{T+1}}^n) + L\|\theta_{T+1} - \theta_0\| \cdot \mathbb{E}_{\mathbf{x} \sim U([0,1]^d)} [f_n(\mathbf{x})] \\
&\leq L_{f_n}(g_{\theta_{T+1}}^n) + \frac{21L^2C^2 \max\{1, \eta^2\} \sqrt{dp(n)}}{\alpha(n)} T^2, \tag{6}
\end{aligned}$$

where we used Eq. (5) and that $|f_n(\mathbf{x})| \leq C$ for all $\mathbf{x} \in [0, 1]^d$. This proves that:

$$L_{f_n}(g_{\theta_0}^n) - L_{f_n}(g_{\theta_T}^n) \leq \frac{21L^2C^2 \max\{1, \eta^2\} \sqrt{dp(n)}}{\alpha(n)} T^2,$$

and in particular, if T is polynomial in n , then \mathcal{F} is not weakly learnable with gradient descent. ■

Appendix C. Proof of Thm. 14

To prove the theorem, we first need the following technical lemma (due to Szörényi (2009)):

Lemma 25 *Let f_1, \dots, f_d such that $|\langle f_i, f_j \rangle| < \frac{1}{d}$ for every $i \neq j$. Fix some $h : \mathcal{X} \rightarrow [-1, 1]$, and some $\tau > 0$. Then, the number of functions from f_1, \dots, f_d for which $|\langle f_j, h \rangle| \geq \tau$ is at most $2(\tau^2 - \frac{1}{d})^{-1}$, i.e.:*

$$|\{j \in [d] : |\langle f_j, h \rangle| \geq \tau\}| \leq 2(\tau^2 - 1/d)^{-1}$$

Proof We define $A = \{j \in [d] : \langle f_j, h \rangle \geq \tau\}$, and note that:

$$\tau |A| \leq \sum_{j \in A} \langle h, f_j \rangle = \left\langle h, \sum_{j \in A} f_j \right\rangle$$

We also have:

$$\begin{aligned}
\left\langle h, \sum_{j \in A} f_j \right\rangle_{\mathcal{D}}^2 &\leq \|h\|_{\mathcal{D}}^2 \left\| \sum_{j \in A} f_j \right\|_{\mathcal{D}}^2 \leq \left\langle \sum_{j \in A} f_j, \sum_{j \in A} f_j \right\rangle = \sum_{j \in A} \left(1 + \sum_{j' \in A, j' \neq j} \langle f_j, f_{j'} \rangle \right) \\
&\leq |A| \left(1 + \frac{|A|}{d} \right) = |A| + \frac{|A|^2}{d}
\end{aligned}$$

So, we have $|A| \leq (\tau^2 - 1/d)^{-1}$. Similarly, for $A' = \{j \in [d] : \langle f_j, \Psi_i \rangle \leq -\tau\}$ we get $|A'| \leq (\tau^2 - 1/d)^{-1}$. Therefore, the required follows. ■

We use the following Theorem, which is an extension of the result in Blum et al. (1994):

Theorem 26 Let \mathcal{F} be a class of functions over \mathcal{X} and let \mathcal{D} be a distribution such that $SQdim(\mathcal{F}, \mathcal{D}) \geq d \geq 16$. Let ℓ be the hinge-loss. Then any statistical-query algorithm with tolerance at least $1/d^{1/3}$, needs at least $\frac{1}{8}d^{1/3}$ queries to learn \mathcal{F} with loss less than $1 - \frac{2}{\sqrt{d}}$.

Proof Let f_1, \dots, f_d be the maximal set of functions with $|\langle f_i, f_j \rangle_{\mathcal{D}}| \leq \frac{1}{d^3}$. Now, let $g : \{\pm 1\}^n \times \{\pm 1\} \rightarrow [-1, 1]$ be some statistical-query, and denote $C_g = \mathbb{E}_{\mathbf{x} \sim \mathcal{D}, y \sim \{\pm 1\}} [g(\mathbf{x}, y)]$. We say that some function f_i is consistent with C_g if:

$$\left| \mathbb{E}_{\mathbf{x} \sim \mathcal{D}} [g(\mathbf{x}, f_i(\mathbf{x}))] - C_g \right| \leq \frac{1}{d^{1/3}}$$

Denote g_+, g_- such that $g_+(\mathbf{x}) = g(\mathbf{x}, 1)$ and $g_-(\mathbf{x}) = g(\mathbf{x}, -1)$. Now, observe that:

$$\begin{aligned} \mathbb{E}_{\mathbf{x} \sim \mathcal{D}} [g(\mathbf{x}, f_i(\mathbf{x}))] - C_g &= \mathbb{E}_{\mathbf{x} \sim \mathcal{D}} \left[\mathbf{1}_{f_i(\mathbf{x})=1} g_+(\mathbf{x}) + \mathbf{1}_{f_i(\mathbf{x})=-1} g_-(\mathbf{x}) - \frac{1}{2} g_+(\mathbf{x}) - \frac{1}{2} g_-(\mathbf{x}) \right] \\ &= \mathbb{E}_{\mathbf{x} \sim \mathcal{D}} \left[\frac{1}{2} (1 + f_i(\mathbf{x})) g_+(\mathbf{x}) + \frac{1}{2} (1 - f_i(\mathbf{x})) g_-(\mathbf{x}) - \frac{1}{2} g_+(\mathbf{x}) - \frac{1}{2} g_-(\mathbf{x}) \right] \\ &= \frac{1}{2} \mathbb{E}_{\mathbf{x} \sim \mathcal{D}} [f_i(\mathbf{x}) g_+(\mathbf{x}) - f_i(\mathbf{x}) g_-(\mathbf{x})] \\ &= \frac{1}{2} \langle f_i, g_+ \rangle - \frac{1}{2} \langle f_i, g_- \rangle = \left\langle f_i, \frac{1}{2} (g_+ - g_-) \right\rangle \end{aligned}$$

Denote $\bar{g} = \frac{1}{2}(g_+ - g_-)$ and since $\bar{g}(\mathbf{x}) \in [-1, 1]$ we get from 25 that the number of functions from f_1, \dots, f_d that are not consistent with C_g is at most $2(1/d^{2/3} - 1/d)^{-1} \leq 4d^{2/3}$. Now, let \mathcal{A} be some statistical-query algorithm, and let g_1, \dots, g_k be k queries made by \mathcal{A} , upon receiving responses $C_{g_1}, \dots, C_{g_{k-1}}$, and let h be the hypothesis returned by \mathcal{A} after these queries. Now, the number of functions f_i that are consistent with all the responses $C_{g_1}, \dots, C_{g_{k-1}}$ is at most $4kd^{2/3}$. So, if \mathcal{A} makes at most $\frac{1}{8}d^{1/3}$ queries, then at least $1/2$ of the functions f_1, \dots, f_d are consistent with the responses, and assume w.l.o.g. that $f_1, \dots, f_{d/2}$ are consistent with the responses. Now, let \tilde{h} be the clipping of h to $[-1, 1]$, so:

$$\tilde{h}(\mathbf{x}) = \begin{cases} -1 & h(\mathbf{x}) < -1 \\ h(\mathbf{x}) & h(\mathbf{x}) \in [-1, 1] \\ 1 & h(\mathbf{x}) > 1 \end{cases}$$

Notice that for the hinge-loss ℓ we have for every \mathbf{x} and $y \in \{\pm 1\}$ that $\ell(h(\mathbf{x}), y) = \max\{1 - h(\mathbf{x})y, 0\} \geq 1 - y\tilde{h}(\mathbf{x})$. Indeed, we have the following cases:

- If $h(\mathbf{x}) \in [-1, 1]$ then $\max\{1 - h(\mathbf{x})y, 0\} = 1 - h(\mathbf{x})y = 1 - \tilde{h}(\mathbf{x})y$.
- If $h(\mathbf{x})y > 1$ then $\max\{1 - h(\mathbf{x})y, 0\} = 0 = 1 - y\tilde{h}(\mathbf{x})$.
- If $h(\mathbf{x})y < -1$ then $\max\{1 - h(\mathbf{x})y, 0\} = 1 - h(\mathbf{x})y \geq 1 - \tilde{h}(\mathbf{x})y$.

Therefore, the hinge-loss of h with respect to some function f_i is:

$$\begin{aligned} L_{f_i(\mathcal{D})}(h) &= \mathbb{E}_{\mathbf{x} \sim \mathcal{D}} [\ell(h(\mathbf{x}), f_i(\mathbf{x}))] \\ &\geq \mathbb{E}_{\mathbf{x} \sim \mathcal{D}} [1 - f_i(\mathbf{x})\tilde{h}(\mathbf{x})] \\ &= 1 - \langle f_i, \tilde{h} \rangle \end{aligned}$$

From Lemma 25, there are at most $(4/d - 1/d)^{-1} = d/3 < d/2$ functions in f_1, \dots, f_d with $\langle f_i, \tilde{h} \rangle \geq \frac{2}{\sqrt{d}}$. So, there exists a function f_i that is consistent with all the responses of the oracle and has $\langle f_i, \tilde{h} \rangle < \frac{2}{\sqrt{d}}$. For this function we get:

$$L_{f_i(\mathcal{D})}(h) \geq 1 - \langle f_i, \tilde{h} \rangle \geq 1 - \frac{2}{\sqrt{d}}$$

■

Proof of Theorem 14. Assume \mathcal{F} can be weakly approximated with respect to \mathcal{D} by a polynomial-size kernel class \mathcal{H} . Then, from Theorem 31, there exists a polynomial $p(n)$ such that $\text{SQdim}(\mathcal{F}_n, \mathcal{D}_n) \leq p(n)$. Then, the algorithm that returns f_i with $i = \arg \max_{j \in [d]} |\langle f_j, f^* \rangle|$ is an efficient statistical-query algorithm that weakly learns \mathcal{F} with respect to \mathcal{D} and the hinge-loss (since the returned function is a binary function, the hinge-loss and the zero-one loss are equivalent).

In the other direction, assume that \mathcal{F} is efficiently weakly learnable from statistical-queries. So, there exists some polynomial p and a statistical-query algorithm that returns a loss with error $\leq 1 - \frac{1}{p(n)}$ for every $f \in \mathcal{F}_n$, using a polynomial number of queries and polynomial tolerance. Therefore, from Theorem 26, there exists a (positive) polynomial q such that $\text{SQdim}(F_n, \mathcal{D}_n) \leq q(n)$. So, let $f_1, \dots, f_{q(n)}$ be the maximal set such that $|\langle f_i, f_j \rangle| \leq \frac{1}{q(n)^3}$ for every $i \neq j$. Denote $\Psi_n(\mathbf{x}) = (f_1(\mathbf{x}), \dots, f_{q(n)}(\mathbf{x}))$, and so we have $f_i \in \mathcal{H}_{\Psi_n}^1$ for every i . So, for every $f \in F$ we have:

$$\begin{aligned} \min_{h \in \mathcal{H}_{\Psi_n}^1} L_{f(\mathcal{D})}(h) &\leq \min_{i \in [q(n)]} L_{f(\mathcal{D})}(f_i) = \min_{i \in [q(n)]} \mathbb{E}_{\mathcal{D}} \ell(f_i(\mathbf{x}), f(\mathbf{x})) \\ &= \min_{i \in [q(n)]} \mathbb{E}_{\mathcal{D}} [1 - f_i(\mathbf{x})f(\mathbf{x})] = L_{\mathcal{D}}(\mathbf{0}) - \max_{i \in [q(n)]} |\langle f_i, f \rangle| \leq L_{\mathcal{D}}(\mathbf{0}) - \frac{1}{q(n)^3} \end{aligned}$$

■

Appendix D. Proofs from Sec. 5

D.1. Telgarski's Function

Recall we consider approximating the following function for $n \in \mathbb{N}$:

$$f_n(\mathbf{x}) = \begin{cases} 1 & \exists t \in \mathbb{N}, x_1 \in [\frac{2t}{2^n}, \frac{2t+1}{2^n}] \\ -1 & \text{otherwise} \end{cases}$$

In words, we split the interval $[0, 1]$ to 2^n intervals each of length $\frac{1}{2^n}$, on even intervals $f_n(x) = 1$ and on odd intervals $f_n(x) = -1$. For the rest of the proof we consider the distribution family $\mathcal{F} = \{\mathcal{F}_n\}_{n \in \mathbb{N}}$ defined by the functions f_n . We have that:

Lemma 27 *Let $\mathcal{H} = \{\mathcal{H}_n\}$ be the hypothesis classes defined by the sign function of ReLU neural networks with depth $2n$ and width 2, and \mathcal{F} as above. Then \mathcal{F} is realizable by \mathcal{H} .*

Proof Following Lemma 3.10 from [Telgarsky \(2016\)](#), we define $h' : \mathbb{R} \rightarrow \mathbb{R}$ by $h'(x) = \sigma(2\sigma(x) - 4\sigma(x - 1/2) - 1/2)$ where σ is the ReLU function. Note that h' interpolates between $(0, -1/2)$, $(1/2, 1/2)$ and $(1, -1/2)$. We define $h : \mathbb{R}^d \rightarrow \mathbb{R}$ by $h(\mathbf{x}) = \text{sign}(h' \circ \dots \circ h'(x_1))$ where the composition is done n times. Note that $h(\mathbf{x})$ is a $2n$ -layer neural network and that $h(\mathbf{x}) = f_n(\mathbf{x})$ (the first and last interval might be cut in half, in this case we can shift $x_1 \mapsto x_1 - 1/2^{n+1}$). \blacksquare

In order to show that the functions above cannot be weakly approximated with a shallow neural network, we show that a 1-dimensional ReLU network is a piecewise linear function, and bound the number of pieces.

Lemma 28 *Let $N(x) : [0, 1] \rightarrow \mathbb{R}$ be a ReLU network with depth L and width k . Then $N(x)$ is a piecewise linear function with at most $2^{L-1}k^L$ pieces.*

Proof First, let $g_1, g_2 : [0, 1] \rightarrow \mathbb{R}$ be two piecewise linear functions with corresponding k_1, k_2 pieces. Denote $h(x) = g_1(x) + g_2(x)$, we show that $h(x)$ is also a piecewise linear function with at most $k_1 + k_2$ pieces. Let A_1 (resp. A_2) be a partition of $[0, 1]$ such that on each interval in A_1 (resp. A_2) the function g_1 (resp. g_2) is linear. Take B to be a partition which is a refinement of A_1 and A_2 in the following way: Denote $A_1 = \{[a_0, a_1], \dots, [a_{k-1}, a_k]\}$ and $A_2 = \{[a'_0, a'_1], \dots, [a'_{k-1}, a'_k]\}$. We construct $B = \{[b_0, b_1], \dots, [b_{m-1}, b_m]\}$ such that $b_0 = a_0 = 0$ (note that also $a'_0 = 0$), $b_1 = \min\{a_1, a'_1\}$ and inductively for every $i > 1$ take $b_i = \min\{a_j, a'_{j'}\}$ where j and j' are the smallest index for which there is no $l, l' < i$ with $b_l = a_j$ and $b_{l'} = a'_{j'}$. Note that B has at most $k_1 + k_2$ intervals, since each boundary point of an interval in B contains exactly one unique boundary point from an interval in either A_1 or A_2 . Also, in each interval of B the function $h(x)$ is linear, because it is linear in both g_1 and g_2 . Hence $h(x)$ is piecewise linear with at most $k_1 + k_2$ pieces.

Next, we can write $N(x)$ in the following way: Let $a_0 = x$, $a_1 = \sigma(U_1 a_0 + b_1)$ where $U_1 \in \mathbb{R}^{k \times 1}$, $b_1 \in \mathbb{R}^k$, and for $1 < i < L$ $a_i = \sigma(U_i a_{i-1} + b_i)$ where $U_i \in \mathbb{R}^{k \times k}$, $b_i \in \mathbb{R}^k$, finally $a_L = U_L a_{L-1} + b_L$ for $U_L \in \mathbb{R}^{1 \times L}$, $b_L \in \mathbb{R}$. We show that each coordinate of a_i for $1 \leq i \leq L$ is a piecewise linear function of x using induction, and bound the amount of pieces by k^L . For $i = 0$ it is clear, assume it is true for $i - 1$. Because $a_i = \sigma(U_i a_{i-1} + b_i)$, the j -th coordinate of a_i is equal to $\sigma\left(\sum_{l=1}^k (U_i)_{jl}(a_{i-1})_l + (b_i)_l\right)$, by the induction hypothesis for every l , $(a_{i-1})_l$ is a piecewise linear function of x with at most k^{i-1} pieces. By the previous claim, $\sum_{l=1}^k (U_i)_{jl}(a_{i-1})_l + (b_i)_l$ is a sum of k piecewise linear functions, each with $(2k)^{i-1}$ pieces hence it has at most $2^{i-1}k^i$ pieces.

Finally note that since the ReLU function is a piecewise linear function with two pieces, composing it with a linear function is a piecewise linear function with at most two pieces. Hence, composing ReLU with a piecewise linear function with k pieces can turn every linear segment into at most two linear segments, meaning that the composition is a piecewise linear function with at most $2k$ pieces. This means that $\sigma\left(\sum_{l=1}^k (U_i)_{jl}(a_{i-1})_l + (b_i)_l\right)$ is a piecewise linear function with $(2k)^i$ pieces.

Using the above for $i = L$ we have that $N(x)$ is a piecewise linear function with at most $2^{L-1}k^L$ pieces (note that the last layer does not contain ReLU activation). \blacksquare

We can now prove that the functions above cannot be weakly approximated by neural networks with depth less than \sqrt{n} :

Theorem 29 Let $k : \mathbb{N} \rightarrow \mathbb{N}$ some function such that $k(n) \leq \sqrt{n}$. Then the sign function of every polynomial-size depth- $k(n)$ dimension- d network **cannot** weakly approximate $\mathcal{F} := \{\mathcal{F}_n\}_{n \in \mathbb{N}}$.

Proof Let $g : [0, 1] \rightarrow \{-1, 1\}$ be a function with at most k jumps (i.e. change of output from -1 to 1 or from 1 to -1), and let $n \in \mathbb{N}$ with $2^{n-1} > k$. We split the interval $[0, 1]$ into intervals $\left[\frac{2i}{2^n}, \frac{2(i+1)}{2^n}\right]$ for $i = 0, \dots, 2^{n-1} - 1$, there are 2^{n-1} such intervals. Since $g(x)$ changes sign at most k times there are at least $2^{n-1} - k$ intervals on which $g(x)$ is constant. Hence, we have that:

$$L_{\mathcal{F}_n}(g) = \mathbb{E}_{x \sim U([0,1])} [\max\{0, 1 - f_n(x) \cdot g(x)\}] \geq 1 \cdot \frac{(2^{n-1} - k)}{2^{n-1}} + 0 \cdot \frac{k}{2^{n-1}} = \frac{(2^{n-1} - k)}{2^{n-1}} \quad (7)$$

For any $\mathbf{y} \in \mathbb{R}^{d-1}$ we define $p_{\mathbf{y}} : \mathbb{R} \rightarrow \mathbb{R}^d$ by $p_{\mathbf{y}}(x) = (\mathbf{y}, x)$. Let $N(\mathbf{x})$ be a depth- L dimension- d network with width $p(n)$ for some polynomial p , and note that if we fix $\mathbf{y} \in \mathbb{R}^{d-1}$ then $N \circ p_{\mathbf{y}} : \mathbb{R} \rightarrow \mathbb{R}$ is a depth- L dimension-1 network with width at most $2p(n)$. Hence, by Lemma 28 $N \circ p_{\mathbf{y}}$ is a piecewise linear function with at most $2^{L-1}(2k)^L$ pieces, which shows that the number of jumps of its sign has the same bound. Since we assume the dimension is constant and the width is polynomial in n , we can assume that for large enough n , $p(n) > d$. Now, using Eq. (7) we get:

$$\begin{aligned} L_{f_n(\mathcal{D}_n)}(N) &= \mathbb{E}_{\mathbf{x} \sim U([0,1]^d)} [\max\{0, 1 - f_n(\mathbf{x}) \cdot N(\mathbf{x})\}] \\ &= \int_{\mathbf{y} \sim U([0,1]^{d-1})} \left(\int_{x \sim U([0,1])} \max\{0, 1 - f_n \circ p_{\mathbf{y}}(x) \cdot N \circ p_{\mathbf{y}}(x)\} dx \right) d\mathbf{y} \\ &\geq \frac{(2^{n-1} - 2^{\sqrt{n}-1}(2p(n))^{\sqrt{n}})}{2^{n-1}} = 1 - \frac{2^{\sqrt{n}}(2p(n))^{\sqrt{n}}}{2^n}. \end{aligned}$$

In particular, for any polynomial $q(n)$ we have that $L_{f_n(\mathcal{D}_n)}(N) \geq 1 - 1/q(n)$ ■

The proof of Thm. 18 now directly follows from Lemma 27 and Thm. 29.

D.2. Parity Functions

For the following result we use the notion of the statistical-query (SQ) dimension (first introduced in Blum et al. (1994)). This notion is a measure of complexity of a target class with respect to a given distribution, which counts the number of "almost-orthogonal" function in the target class with respect to the inner product induced by the distribution:

Definition 30 Let \mathcal{F}_n be some target class and let \mathcal{D}_n be some distribution over \mathcal{X}_n . We define the statistical-query dimension of $(\mathcal{F}_n, \mathcal{D}_n)$, denoted $SQdim(\mathcal{F}_n, \mathcal{D}_n)$, to be the largest number $d \in \mathbb{N}$ for which there exist functions $f_1, \dots, f_d \in \mathcal{F}_n$ such that for every $i \neq j$ we have:

$$\left| \langle f_i, f_j \rangle_{\mathcal{D}_n} \right| := \left| \mathbb{E}_{\mathbf{x} \sim \mathcal{D}_n} [f_i(\mathbf{x})f_j(\mathbf{x})] \right| < \frac{1}{d}$$

While the SQ-dimension was introduced as a measure of bounding the complexity of statistical-query algorithms, we show that this measure can also directly bound the approximation power of kernel classes. In particular, we show that if the SQ-dimension of some target class and distribution is super-polynomial, then they cannot be weakly approximated by a polynomial-size kernel class. In the following results we assume that $\mathcal{X}_n = \{\pm 1\}^n$.

Theorem 31 Let \mathcal{H} be some polynomial-size kernel class, let $\mathcal{F} = \{\mathcal{F}_n\}_{n \in \mathbb{N}}$ a sequence of target classes and let $\mathcal{D} = \{\mathcal{D}_n\}_{n \in \mathbb{N}}$ be a sequence of distributions over $\{\mathcal{X}_n\}_{n \in \mathbb{N}}$. Denote $d(n) = \text{SQdim}(\mathcal{F}_n, \mathcal{D}_n)$. Then, if $d(n)$ is super-polynomial, \mathcal{H} cannot weakly approximate \mathcal{F} with respect to \mathcal{D} .

The theorem directly connects between the ability of kernel classes to approximate a target class, and its SQ-dimension. The proof relies on the fact that for a fixed mapping $\Psi_n : \mathcal{X}_n \rightarrow \mathcal{Y}_n$ and a set of "almost-orthogonal" functions $f_1, \dots, f_{d(n)}$, it cannot happen that the correlation between Ψ_n and all of the f_i 's is large. Hence, given many such f_i 's, it is possible to find one which Ψ_n cannot approximate well, in other words, the more f_i 's there are, it is harder for Ψ_n to approximate them all at once.

The proof of Theorem 31 is largely based on the following key lemma:

Lemma 32 Fix some $\Psi : \mathcal{X} \rightarrow [-1, 1]^N$, and define:

$$\mathcal{H}_\Psi^B = \{\mathbf{x} \rightarrow \langle \Psi(\mathbf{x}), \mathbf{w} \rangle : \|\mathbf{w}\|_2 \leq B\}$$

Then, if $d(n) > 3$, there exist $f_1, \dots, f_{d(n)} \in \mathcal{F}_n$ such that:

$$\mathbb{E}_{j \sim [d(n)]} \left[\min_{h \in \mathcal{H}_\Psi^B} L_{f_j(\mathcal{D}_n)}(h) \right] \geq 1 - \frac{\sqrt{5NB}}{d(n)^{1/12}}$$

Proof Let $f_1, \dots, f_{d(n)} \in \mathcal{F}_n$ be the set of functions realizing $\text{SQdim}(\mathcal{F}_n, \mathcal{D}_n)$, and we denote $d := d(n)$. For some $j \in [d]$, let $\mathcal{L}_j(\mathbf{w}) := L_{f_j(\mathcal{D}_n)}(\langle \Psi(\mathbf{x}), \mathbf{w} \rangle)$ and define the objective $G_j(\mathbf{w}) := \mathcal{L}_j(\mathbf{w}) + \frac{\lambda}{2} \|\mathbf{w}\|^2$. Observe that for every $i \in [N]$ we have:

$$\frac{\partial}{\partial w_i} G_j(0) = \mathbb{E}_{\mathbf{x} \sim \mathcal{D}_n} [f_j(\mathbf{x}) \Psi_i(\mathbf{x})]$$

And so:

$$\begin{aligned} \mathbb{E}_{j \sim [d]} [\|\nabla G_j(0)\|^2] &= \mathbb{E}_{j \sim [d]} \left[\sum_{i \in [N]} \left(\frac{\partial}{\partial w_i} G_j(0) \right)^2 \right] \\ &= \mathbb{E}_{j \sim [d]} \left[\sum_{i \in [N]} \mathbb{E}_{\mathbf{x} \sim \mathcal{D}} [f_j(\mathbf{x}) \Psi_i(\mathbf{x})]^2 \right] \\ &= \sum_{i \in [N]} \frac{1}{d} \sum_{j=1}^d \mathbb{E}_{\mathbf{x} \sim \mathcal{D}} [f_j(\mathbf{x}) \Psi_i(\mathbf{x})]^2 \end{aligned} \tag{8}$$

We define $A = \{j \in [d] : |\langle f_j, \Psi_i \rangle_{\mathcal{D}}| \geq \tau\}$, and from Lemma 25 we have $|A| \leq 2(\tau^2 - \frac{1}{d})^{-1}$. Therefore, for $\tau = d^{-1/3}$ we get:

$$\begin{aligned} \sum_{j=1}^d \mathbb{E}_{\mathbf{x} \sim \mathcal{D}} [f_j(\mathbf{x}) \Psi_i(\mathbf{x})]^2 &= \sum_{j \in A} \langle f_j, \Psi_i \rangle_{\mathcal{D}}^2 + \sum_{j \notin A} \langle f_j, \Psi_i \rangle_{\mathcal{D}}^2 \\ &\leq 2(\tau^2 - 1/d)^{-1} + d\tau \leq 2(d^{-2/3} - 1/d)^{-1} + d^{2/3} \leq 5d^{2/3} \end{aligned}$$

And plugging into (8) we get:

$$\mathbb{E}_{j \sim [d]} [\|\nabla G_j(0)\|^2] = \sum_{i \in [N]} \frac{1}{d} \sum_{j=1}^d \mathbb{E}_{\mathbf{x} \sim \mathcal{D}} [f_j(\mathbf{x}) \Psi_i(\mathbf{x})]^2 \leq \frac{5N}{d^{1/3}}$$

Using Jensen inequality we get:

$$\mathbb{E}_{j \sim [d]} [\|\nabla G_j(0)\|]^2 \leq \mathbb{E}_{j \sim [d]} [\|\nabla G_j(0)\|^2] \leq \frac{5N}{d^{1/3}} \quad (9)$$

Note that G_j is λ -strongly convex, and therefore, for every \mathbf{w}, \mathbf{u} we have:

$$\langle \nabla G_j(\mathbf{w}) - \nabla G_j(\mathbf{u}), \mathbf{w} - \mathbf{u} \rangle \geq \lambda \|\mathbf{w} - \mathbf{u}\|^2$$

Let $\mathbf{w}_j^* := \arg \min_{\mathbf{w}} G_j(\mathbf{w})$, and so $\nabla G_j(\mathbf{w}_j^*) = 0$. Using the above we get:

$$\lambda \|\mathbf{w}_j^*\|^2 \leq \langle \nabla G_j(\mathbf{w}_j^*) - \nabla G_j(0), \mathbf{w}_j^* \rangle \leq \|\nabla G_j(0)\| \|\mathbf{w}_j^*\| \Rightarrow \|\mathbf{w}_j^*\| \leq \frac{1}{\lambda} \|\nabla G_j(0)\|$$

Now, notice that \mathcal{L}_j is \sqrt{N} -Lipschitz, since:

$$\|\nabla \mathcal{L}_j(\mathbf{w})\| = \|\nabla \mathbb{E} [\ell(y, \langle \Psi(\mathbf{x}), \mathbf{w} \rangle)]\| \leq \mathbb{E} [|\ell'| \|\Psi(\mathbf{x})\|] \leq \sqrt{N}$$

Therefore, we get that:

$$1 - \mathcal{L}_j(\mathbf{w}_j^*) = \mathcal{L}_j(0) - \mathcal{L}_j(\mathbf{w}_j^*) \leq \sqrt{N} \|\mathbf{w}_j^*\| \leq \frac{\sqrt{N}}{\lambda} \|\nabla G_j(0)\| \quad (10)$$

Denote $\hat{\mathbf{w}}_j = \arg \min_{\|\mathbf{w}\| \leq B} \mathcal{L}_j(\mathbf{w})$, and by optimality of \mathbf{w}_j^* we have:

$$\mathcal{L}_I(\mathbf{w}_j^*) \leq \mathcal{L}_j(\mathbf{w}_j^*) + \frac{\lambda}{2} \|\mathbf{w}_j^*\|^2 \leq \mathcal{L}_j(\hat{\mathbf{w}}_j) + \frac{\lambda}{2} \|\hat{\mathbf{w}}_j\|^2 \leq \mathcal{L}_j(\hat{\mathbf{w}}_j) + \frac{\lambda B^2}{2} \quad (11)$$

From 10 and 11 we get:

$$1 - \frac{\sqrt{N}}{\lambda} \|\nabla G_j(0)\| \leq \mathcal{L}_j(\mathbf{w}_j^*) \leq \mathcal{L}_j(\hat{\mathbf{w}}_j) + \frac{\lambda B^2}{2} \quad (12)$$

Taking an expectation and plugging in 9 we get:

$$\mathbb{E}_{j \sim [d]} \left[\min_{h \in \mathcal{H}_{\Psi}^B} L_{f_j(\mathcal{D})}(h) \right] = \mathbb{E}_{j \sim [d]} [\mathcal{L}_j(\hat{\mathbf{w}}_j)] \geq 1 - \frac{\sqrt{N}}{\lambda} \mathbb{E}_{j \sim [d]} [\|\nabla G_j(0)\|] - \frac{\lambda B^2}{2} \geq 1 - \frac{\sqrt{5N}}{\lambda d^{1/6}} - \frac{\lambda B^2}{2}$$

Since this is true for all λ , taking $\lambda = \frac{\sqrt{2\sqrt{5N}}}{d^{1/12} B}$ we get:

$$\mathbb{E}_{j \sim [d]} \left[\min_{h \in \mathcal{H}_{\Psi}^B} L_{f_j(\mathcal{D})}(h) \right] \geq 1 - \frac{\sqrt{2\sqrt{5N}} B}{d^{1/12}}$$

■

Given the above Lemma, the proof of the Theorem 31 is immediate:

Proof of Thm. 31.

Fix some polynomial-size kernel class $\mathcal{H} := \{\mathcal{H}_n\}_{n \in \mathbb{N}}$, with mappings $\Psi_n : \mathcal{X}_n \rightarrow [-1, 1]^{p(n)}$ and $\mathcal{H}_n = \mathcal{H}_{\Psi_n}^{q(n)}$ for some polynomials p, q . Then, from Lemma 32, for every n we have:

$$\max_{f \in \mathcal{F}_n} \min_{h \in \mathcal{H}_n} L_{f(\mathcal{D}_n)}(h) \geq \mathbb{E}_{j \sim [d]} \min_{h \in \mathcal{H}_n} L_{f_j(\mathcal{D}_n)}(h) = \mathbb{E}_{j \sim [d]} \left[\min_{h \in \mathcal{H}_{\Psi}^{q(n)}} L_{f_j(\mathcal{D})}(h) \right] \geq 1 - \frac{\sqrt{5p(n)q(n)}}{d(n)^{1/12}}$$

Since for every \mathcal{D} we have $L_{\mathcal{D}}(\mathbf{0}) = 1$ (we use the hinge-loss), we get that:

$$\max_{f \in \mathcal{F}_n} \min_{h \in \mathcal{H}_n} L_{f(\mathcal{D}_n)}(\mathbf{0}) - L_{f(\mathcal{D}_n)}(h) \leq \frac{\sqrt{5p(n)q(n)}}{d(n)^{1/5}}$$

And so, for every polynomial r we have:

$$\sup_{n \in \mathbb{N}} \max_{f \in \mathcal{F}_n} \min_{h \in \mathcal{H}_n} |r(n)| (L_{f(\mathcal{D})}(\mathbf{0}) - L_{f(\mathcal{D})}(h)) \leq \sup_{n \in \mathbb{N}} \frac{\sqrt{5p(n)q(n)} |r(n)|}{d(n)^{1/12}} = 0$$

Which proves the theorem. ■

Now, we can use Theorem 31 to show a strong separation between **any** polynomial-size kernel class and the class of polynomial-size shallow neural-network. All we need is to find a target class of functions which have super polynomial SQ-dimension, but can be realized by 2-layer neural networks.

We will use **parity functions** over n -bits: For some subset $I \subseteq [n]$, denote by $f_I(\mathbf{x}) = \prod_{i \in I} x_i$, the parity over the bits of I , let $\mathcal{F}_n = \{f_I : I \subseteq [n]\}$ and \mathcal{D}_n the uniform distribution on \mathcal{X}_n .

Proof [Proof of Thm. 19] By Lemma 5 from [Shalev-Shwartz et al. \(2017\)](#) parity functions are realizable by 2-layer neural networks with width and weight magnitude linear in the input dimension. On the other hand, it is easy to see that every two parity functions on a different subset on \mathcal{X}_n are orthogonal, and there are 2^n such functions. Hence, by Thm. 31 parity functions are not weakly approximated by a polynomial-size kernel class. ■

Appendix E. Proofs from Sec. 6

E.1. Proof of Thm. 20

We first show use a concentration inequality to find a large set of vectors with pairwise small inner product.

Lemma 33 *There exists a set of vectors $\mathbf{z}^{(1)}, \dots, \mathbf{z}^{(d)} \in \mathcal{X}_n$ of size $d = 2^{n/12}$ such that for every $i \neq j$ we have $\sum_{t=1}^n \mathbf{1}\{z_t^{(i)} \neq z_t^{(j)}\} \geq \frac{n}{4}$.*

Proof Fix some d , draw d vectors uniformly from \mathcal{X} , and denote them $Z := (\mathbf{z}^{(1)}, \dots, \mathbf{z}^{(d)})$. Fix some $i \neq j$, and denote $S_{i,j} = \sum_{t=1}^n \mathbf{1}\{z_t^{(i)} \neq z_t^{(j)}\}$. Notice that: $\mathbb{E}_Z [S_{i,j}] = \sum_t \mathbb{P} [z_t^{(i)} \neq z_t^{(j)}] = \frac{n}{2}$, and by Hoeffding's inequality:

$$\mathbb{P} \left[S_{i,j} \leq \frac{n}{4} \right] \leq \exp \left(-\frac{n}{8} \right)$$

There are $\frac{d^2}{2}$ choices for $i \neq j$, and so using the union bound we get:

$$\mathbb{P} \left[\forall i \neq j : S_{i,j} > \frac{n}{4} \right] \geq 1 - \frac{d^2}{2} e^{-n/8}$$

Since $2^{4/3} \leq e$, choosing $d = 2^{n/12} \leq e^{n/16}$ we get that with probability at least $\frac{1}{2}$ over the choice of Z , for every $i \neq j$ we have $S_{i,j} \geq \frac{n}{4}$. Therefore, the required follows. \blacksquare

Proof of Theorem 20. Fix $\mathbf{z}^{(1)}, \dots, \mathbf{z}^{(d)} \in \mathcal{X}_n$ with $d = 2^{n/12}$ such that $\sum_{t=1}^n \mathbf{1}\{z_t^{(i)} \neq z_t^{(j)}\} \geq \frac{n}{4}$ for all $i \neq j$. Fix some $i \neq j$ and observe that:

$$\begin{aligned} |\langle F_{\mathbf{z}^{(i)}}, F_{\mathbf{z}^{(j)}} \rangle| &= \left| \mathbb{E}_{\mathbf{x}, \mathbf{z}} \left[\prod_{t \in I(\mathbf{z}^{(i)})} (x_t \vee z_t) \prod_{t \in I(\mathbf{z}^{(j)})} (x_t \vee z_t) \right] \right| \\ &= \left| \prod_{t \in I(\mathbf{z}^{(i)}) \Delta I(\mathbf{z}^{(j)})} \mathbb{E}(x_t \vee z_t) \right| = \prod_{t \in I(\mathbf{z}^{(i)}) \Delta I(\mathbf{z}^{(j)})} \frac{1}{2} \leq 2^{-n/4} = \frac{1}{d^3} \end{aligned}$$

Where $I \Delta J = I \setminus J \cup J \setminus I$, and using the fact that $|I(\mathbf{z}^{(i)}) \Delta I(\mathbf{z}^{(j)})| = |\{t \in [n] : z_t^{(i)} \neq z_t^{(j)}\}| \geq \frac{n}{4}$. This shows that $\text{SQdim}(\mathcal{F}, \mathcal{D}) \geq 2^{n/12}$, combining with Thm. 26 finishes the proof. \blacksquare

E.2. Proof of Thm. 21

We will prove the theorem in a more generalized setting. We start by showing that any target class with large SQ-dimension can be used to construct a target distribution that is hard to approximate using a 2-layer network. This construction is very natural: given a class of functions \mathcal{F}_n from \mathcal{X}_n to $\{\pm 1\}$, we consider a function from $\mathcal{X}_n \times \mathcal{Z}_n$ to $\{\pm 1\}$, for some space \mathcal{Z}_n , where we identify every function in $f' \in \mathcal{F}_n$ with an element $\varphi_n(f') \in \mathcal{Z}_n$. Then, the new function $f(\mathbf{x}, \mathbf{z})$ is just applying $\varphi_n^{-1}(\mathbf{z})$ on the input \mathbf{x} . We call this the induced function:

Definition 34 Let $\mathcal{X}_n, \mathcal{Z}_n$ be two input spaces. Let \mathcal{F}_n be some target class of functions from \mathcal{X}_n to \mathcal{Y} , with $|\mathcal{F}_n| = |\mathcal{Z}_n|$ and let $\varphi_n : \mathcal{Z}_n \rightarrow \mathcal{F}_n$ be some bijection. The tuple $(\mathcal{F}_n, \varphi_n)$ naturally induces a function $f : \mathcal{X}_n \times \mathcal{Z}_n \rightarrow \mathcal{Y}$, where $f(\mathbf{x}, \mathbf{z}) = \varphi_n(\mathbf{z})(\mathbf{x})$. We denote by $F(\mathcal{F}_n, \varphi_n) := f$ the induced function.

A simple example is where $\mathcal{X}_n = \mathcal{Z}_n = \{\pm 1\}^n$, we can think of all the functions of the form $f_{\mathbf{z}} : \mathcal{X}_n \rightarrow \{\pm 1\}$ defined by $f_{\mathbf{z}}(\mathbf{x}) = \prod_{i=1}^n (x_i \vee z_i)$. It is natural to identify each function $f_{\mathbf{z}}$ with a vector $\mathbf{z} \in \mathcal{Z}_n$, hence the induced function $f : \mathcal{X}_n \times \mathcal{Z}_n \rightarrow \{\pm 1\}$ is defined similarly by $f(\mathbf{x}, \mathbf{z}) = \prod_{i=1}^n (x_i \vee z_i)$, where \mathbf{z} is an input vector instead of a constant.

We define a depth-two neural network $g : \mathcal{X} \times \mathcal{Z} \rightarrow \mathbb{R}$ to be any function of the form:

$$g(\mathbf{x}, \mathbf{z}) = \sum_{i=1}^k u_i \sigma \left(\langle \mathbf{w}^{(i)}, \mathbf{x} \rangle + \langle \mathbf{v}^{(i)}, \mathbf{z} \rangle + b_i \right)$$

where σ is some 1-Lipschitz activation. Note that this is the exact same definition of depth-2 neural network we had before, only here we explicitly split the weights to accommodate the two input vectors \mathbf{x} and \mathbf{z} .

We show that the induced function can be used to generate a distribution that is hard to weakly approximate using any depth-two network g :

Theorem 35 *Let $\mathcal{F} = \{\mathcal{F}_n\}_{n \in \mathbb{N}}$ be a sequence of target classes and let $\mathcal{D} = \{\mathcal{D}_n\}_{n \in \mathbb{N}}$ be a sequence of distributions over $\{\mathcal{X}_n\}_{n \in \mathbb{N}}$, and denote $d(n) = \text{SQdim}(\mathcal{F}_n, \mathcal{D}_n)$. For every n , let $\mathcal{Z}_n \subseteq \{\pm 1\}^n$ be some input space with $|\mathcal{Z}_n| = |\mathcal{F}_n|$, and fix a bijection $\varphi_n : \mathcal{Z}_n \rightarrow \mathcal{F}_n$. Then, if $d(n)$ is super-polynomial, there exists a sequence of distributions \mathcal{D}' over $\mathcal{X}_n \times \mathcal{Z}_n$ such that the sequence of induced target functions $\{F(\mathcal{F}_n, \varphi_n)\}_{n \in \mathbb{N}}$ cannot be weakly approximated by any polynomial-size depth-two networks class with respect to \mathcal{D}' .*

Proof [Proof of Thm. 21] The proof that $\{\mathcal{F}_n\}$ cannot be weakly approximated by 2-layer neural networks follows immediately from Thm. 35. We just need to take $\mathcal{Z}_n = \mathcal{X}_n = \{\pm 1\}^n$, and φ_n to be the natural bijection $\mathbf{z} \mapsto F_n(\mathbf{z}, \mathbf{x})$, and using the fact that parity over n -bits w.r.t the uniform distribution has an exponential SQ-dimension.

The proof that $\{\mathcal{F}_n\}$ can be realized by depth-3 neural networks is known, e.g. see [Malach and Shalev-Shwartz \(2019b\)](#). ■

We move on to proving Thm. 35. Fix some target class \mathcal{F}_n and some distribution \mathcal{D}_n such that $d(n) = \text{SQdim}(\mathcal{F}_n, \mathcal{D}_n)$. Fix some $\varphi_n : \mathcal{Z}_n \rightarrow \mathcal{F}_n$ and let $f := F(\mathcal{F}_n, \varphi_n)$. Let $f_1, \dots, f_{d(n)} \in \mathcal{F}_n$ be functions that realize the SQdim. We define a distribution \mathcal{D}'_n over $\mathcal{X}_n \times \mathcal{Z}_n$ such that

$$\mathcal{D}'_n(\mathbf{x}, \mathbf{z}) = \begin{cases} \frac{1}{d(n)} \mathcal{D}_n(\mathbf{x}) & \varphi_n(\mathbf{z}) \in \{f_1, \dots, f_{d(n)}\} \\ 0 & \text{otherwise} \end{cases}.$$

That is, we sample $\mathbf{x} \sim \mathcal{D}_n$ and sample uniformly $\mathbf{z} \sim \{\varphi_n^{-1}(f_1), \dots, \varphi_n^{-1}(f_{d(n)})\}$.

As a first step for proving the Theorem, we show that such distribution cannot be approximated by a polynomial-size depth-two neural network, where the weights \mathbf{v} take integer values:

Lemma 36 *Assume that there exists $\Delta > 0$ such that $v_j^{(i)} \in \Delta\mathbb{Z} := \{\Delta \cdot z : z \in \mathbb{Z}\}$ for every i, j and $\|\mathbf{u}^{(i)}\|, \|\mathbf{w}^{(i)}\|, \|\mathbf{v}^{(i)}\|, \|\mathbf{b}\| < R$.*

$$L_{f(\mathcal{D}'_n)}(g) = \mathbb{E}[\ell(g(\mathbf{x}, \mathbf{z}), f(\mathbf{x}, \mathbf{z}))] \geq 1 - \frac{3\sqrt{10k}R^{5/2}n^{3/4}}{\sqrt{\Delta}d(n)^{1/12}}$$

Proof For every $\mathbf{z} \in \mathcal{X}$ denote $j(\mathbf{z}) = \langle \mathbf{v}^{(i)}, \mathbf{z} \rangle$, and since $v_j^{(i)} \in [-R, R] \cap \Delta\mathbb{Z}$, we get that $j(\mathbf{z}) \in [-R\sqrt{n}, R\sqrt{n}] \cap \Delta\mathbb{Z}$. Indeed, fix some i and we have $\frac{1}{\Delta}\mathbf{v}^{(i)} \in \mathbb{Z}^n$, and since $\mathbf{z} \in \mathbb{Z}^n$ we have $\frac{1}{\Delta}j(\mathbf{z}) = \langle \frac{1}{\Delta}\mathbf{v}^{(i)}, \mathbf{z} \rangle \in \mathbb{Z}$. Define $\Psi_{i,j}(\mathbf{x}) = \frac{1}{3R\sqrt{n}}\sigma(\langle \mathbf{w}^{(i)}, \mathbf{x} \rangle + j + b_i)$ for every $i \in [k]$ and $j \in [-R\sqrt{n}, R\sqrt{n}] \cap \Delta\mathbb{Z}$, and note that:

$$|\Psi_{i,j}(\mathbf{x})| \leq \frac{1}{3R\sqrt{n}} \left| \langle \mathbf{w}^{(i)}, \mathbf{x} \rangle + j + b_i \right| \leq \frac{1}{3R\sqrt{n}} \left(\|\mathbf{w}^{(i)}\| \|\mathbf{x}\| + |j| + |b_i| \right) \leq 1$$

Notice that $[-R\sqrt{n}, R\sqrt{n}] \cap \Delta\mathbb{Z} \leq 2 \left\lfloor \frac{R\sqrt{n}}{\Delta} \right\rfloor$, and so there are at most $2 \left\lfloor \frac{R\sqrt{n}}{\Delta} \right\rfloor$ choices for j . Denote $N := 2k \left\lfloor \frac{R\sqrt{n}}{\Delta} \right\rfloor$ and let $\Psi : \mathcal{X} \rightarrow [-1, 1]^N$ defined as $\Psi(\mathbf{x}) = [\Psi_{i,j}(\mathbf{x})]_{i,j}$ (in vector form).

Denote $B = 3R^2\sqrt{n}$, and from Lemma 32 we have:

$$\begin{aligned}\mathbb{E}_{\mathcal{D}'_n} \left[\min_{\|\hat{\mathbf{u}}\| \leq B} \ell(\langle \hat{\mathbf{u}}, \Psi(\mathbf{x}) \rangle, f(\mathbf{x}, \mathbf{z})) \right] &= \mathbb{E}_{j \sim [d(n)]} \left[\min_{\|\hat{\mathbf{u}}\| \leq B} \ell(\langle \hat{\mathbf{u}}, \Psi(\mathbf{x}) \rangle, f_j(\mathbf{x})) \right] \\ &= \mathbb{E}_{j \sim [d(n)]} \left[\min_{h \in \mathcal{H}_\Psi^B} L_{\mathcal{D}, f_j}(h) \right] \geq 1 - \frac{\sqrt{5NB}}{d(n)^{1/12}}\end{aligned}$$

Notice that $g(\mathbf{x}, \mathbf{z}) = \sum_{i=1}^k 3R\sqrt{n}u_i\Psi_{i,j(\mathbf{z})}(\mathbf{x}) = \langle \mathbf{u}(\mathbf{z}), \Psi(\mathbf{x}) \rangle$ where $\mathbf{u}(\mathbf{z})_{i,j} = \begin{cases} 3R\sqrt{n}u_i & j = j(\mathbf{z}) \\ 0 & j \neq j(\mathbf{z}) \end{cases}$.

Since $\|\mathbf{u}(\mathbf{z})\| \leq 3R\sqrt{n}\|\mathbf{u}\| \leq B$ we get that:

$$\mathbb{E}[\ell(g(\mathbf{x}, \mathbf{z}), f(\mathbf{x}, \mathbf{z}))] = \mathbb{E}[\ell(\langle \mathbf{u}(\mathbf{z}), \Psi(\mathbf{x}) \rangle, f(\mathbf{x}, \mathbf{z}))] \geq \mathbb{E} \left[\min_{\|\hat{\mathbf{u}}\| \leq B} \ell(\langle \hat{\mathbf{u}}, \Psi(\mathbf{x}) \rangle, f(\mathbf{x}, \mathbf{z})) \right] \geq 1 - \frac{\sqrt{5NB}}{d(n)^{1/12}}$$

■

Now, we can extend this result to general polynomial-size depth-two networks. This is done by correctly rounding the weights to get integer values, and use the previous lemma.

Lemma 37 Assume $\|\mathbf{u}^{(i)}\|, \|\mathbf{w}^{(i)}\|, \|\mathbf{v}^{(i)}\|, \|\mathbf{b}\| \leq R$. Then:

$$L_{f(\mathcal{D})}(g) \geq 1 - \frac{6\sqrt{k}R^2n^{5/6}}{d(n)^{1/18}}$$

Proof Fix some $\Delta \in (0, 1)$, and let $\hat{\mathbf{v}}^{(i)} = \Delta \lfloor \frac{1}{\Delta} \mathbf{v}^{(i)} \rfloor \in \Delta\mathbb{Z}^n$, where $\lfloor \cdot \rfloor$ is taken element-wise. Notice that for every j we have:

$$\left| v_j^{(i)} - \hat{v}_j^{(i)} \right| = \left| v_j^{(i)} - \Delta \left\lfloor \frac{1}{\Delta} v_j^{(i)} \right\rfloor \right| = \Delta \left| \frac{1}{\Delta} v_j^{(i)} - \left\lfloor \frac{1}{\Delta} v_j^{(i)} \right\rfloor \right| \leq \Delta$$

Observe the following neural network:

$$\hat{g}(\mathbf{x}, \mathbf{z}) = \sum_{i=1}^k u_i \sigma \left(\langle \mathbf{w}^{(i)}, \mathbf{x} \rangle + \langle \hat{\mathbf{v}}^{(i)}, \mathbf{z} \rangle + b_i \right)$$

For every $\mathbf{x}, \mathbf{z} \in \mathcal{X}$, using Cauchy-Schwartz inequality, and the fact that σ is 1-Lipchitz:

$$\begin{aligned}|g(\mathbf{x}, \mathbf{z}) - \hat{g}(\mathbf{x}, \mathbf{z})| &\leq \|\mathbf{u}\| \sqrt{\sum_{i=1}^k \left| \sigma(\langle \mathbf{w}^{(i)}, \mathbf{x} \rangle + \langle \mathbf{v}^{(i)}, \mathbf{z} \rangle + b_i) - \sigma(\langle \mathbf{w}^{(i)}, \mathbf{x} \rangle + \langle \hat{\mathbf{v}}^{(i)}, \mathbf{z} \rangle + b_i) \right|^2} \\ &\leq \|\mathbf{u}\| \sqrt{\sum_{i=1}^k |\langle \mathbf{v}^{(i)}, \mathbf{z} \rangle - \langle \hat{\mathbf{v}}^{(i)}, \mathbf{z} \rangle|^2} \\ &\leq \|\mathbf{u}\| \sqrt{\sum_{i=1}^k \|\mathbf{v}^{(i)} - \hat{\mathbf{v}}^{(i)}\|^2 \|\mathbf{z}\|^2} \leq R\sqrt{k}\Delta n\end{aligned}$$

Now, by Lemma 36 we have:

$$L_{f(\mathcal{D}'_n)}(\hat{g}) \geq 1 - \frac{3\sqrt{10k}R^{5/2}n^{3/4}}{\sqrt{\Delta}d(n)^{1/12}}$$

And using the fact that ℓ is 1-Lipschitz we get:

$$\begin{aligned} L_{f(\mathcal{D}'_n)}(g) &= \mathbb{E} [\ell(g(\mathbf{x}, \mathbf{z}), f(\mathbf{x}, \mathbf{z}))] \\ &\geq \mathbb{E} [\ell(\hat{g}(\mathbf{x}, \mathbf{z}), f(\mathbf{x}, \mathbf{z}))] - \mathbb{E} [|\ell(g(\mathbf{x}, \mathbf{z}), f(\mathbf{x}, \mathbf{z})) - \ell(\hat{g}(\mathbf{x}, \mathbf{z}), f(\mathbf{x}, \mathbf{z}))|] \\ &\geq L_{f(\mathcal{D}'_n)}(\hat{g}) - \mathbb{E} [|g(\mathbf{x}, \mathbf{z}) - \hat{g}(\mathbf{x}, \mathbf{z})|] \\ &\geq 1 - \frac{3\sqrt{10k}R^{5/2}n^{3/4}}{\sqrt{\Delta}d(n)^{1/12}} - R\sqrt{k}\Delta n \end{aligned}$$

This is true for any $\Delta > 0$, so we choose $\Delta = \frac{\sqrt[3]{45/2}R}{d(n)^{1/18}n^{1/6}}$ and we get:

$$L_{f(\mathcal{D})}(g) \geq 1 - \frac{\sqrt[3]{180}\sqrt{k}R^2n^{5/6}}{d(n)^{1/18}}$$

■

Proof of Theorem 35. Immediate from Lemma 37.

■