

OPTIMAL APPROXIMATION BY RELU MLPS
OF MAXIMAL REGULARITY

BRIDGING THE GAP BETWEEN APPROXIMATION AND
LEARNING VIA OPTIMAL APPROXIMATION BY RELU MLPs
OF MAXIMAL REGULARITY

By RUIYANG HONG, BS

A Thesis Submitted to the School of Graduate Studies in Partial
Fulfillment of the Requirements for
the Degree Master of Science

McMaster University
MASTER OF SCIENCE (2024)
Hamilton, Ontario, Canada (Mathematics)

TITLE: Bridging the Gap Between Approximation and Learning
via Optimal Approximation by ReLU MLPs of Maximal
Regularity

AUTHOR: Ruiyang Hong
BS (Mathematics),
McMaster University, Hamilton, Canada

SUPERVISOR: Anastasis Kratsios

NUMBER OF PAGES: 138

Abstract

The foundations of deep learning are supported by the seemingly opposing perspectives of approximation or learning theory. The former advocates for large/expressive models that need not generalize, while the latter considers classes that generalize but may be too small/constrained to be universal approximators. Motivated by real-world deep learning implementations that are both expressive and statistically reliable, we ask: "Is there a class of neural networks that is both large enough to be universal but structured enough to generalize?"

This paper constructively provides a positive answer to this question by identifying a highly structured class of ReLU multilayer perceptions (MLPs), which are optimal function approximators and are statistically well-behaved. We show that any L -Lipschitz function from $[0, 1]^d$ to $[-n, n]$ can be approximated to a uniform $Ld/(2n)$ error on $[0, 1]^d$ with a sparsely connected L -Lipschitz ReLU MLP of width $\mathcal{O}(dn^d)$, depth $\mathcal{O}(\log(d))$, with $\mathcal{O}(dn^d)$ nonzero parameters, and whose weights and biases take values in $\{0, \pm 1/2\}$ except in the first and last layers which instead have magnitude at-most n . Unlike previously known "large" classes of universal ReLU MLPs, the empirical Rademacher complexity of our class remains bounded even when its depth

and width become arbitrarily large. Further, our class of MLPs achieves a near-optimal sample complexity of $\mathcal{O}(\log(N)/\sqrt{N})$ when given N i.i.d. normalized sub-Gaussian training samples.

We achieve this by avoiding the standard approach to constructing optimal ReLU approximators, which sacrifices regularity by relying on small spikes. Instead, we introduce a new construction that perfectly fits together linear pieces using Kuhn triangulations and avoids these small spikes.

Acknowledgements

My greatest thanks goes to my supervisor, Dr. Anastasis Kratsios, for his patient and valuable guidance throughout my graduate studies. I am grateful for the great opportunity that he provided to conduct this research, and it is the supportive and lively atmosphere he offered that made this research possible.

I would also like to thank the Department of Mathematics and Statistics at McMaster University for its emotional and financial support. My research benefited a lot from the wonderful work environment in the department that fosters ideas and inspiration.

Finally, I would like to thank all my family and friends for their continuous support.

Chapter 1

Introduction

The foundations of deep learning are typically investigated from two seemingly opposing perspectives: either analytically in terms of the approximation power of neural networks [83, 61, 16, 70, 71] or statistically by guaranteeing that these models generalize beyond the training data [6, 7]. On the one hand, the (universal) approximation literature studies MLPs that approximate continuous functions using a minimal number of neurons, and it typically overlooks the regularity needed for them to exhibit reliable statistical behaviour. In contrast, the learning theory literature studies restricted classes of MLPs which generalize beyond their training data by limiting their expressivity either through weights and biases restrictions [54, 49, 81] or Lipschitz constraints [25]. Practical deep-learning implementations, however, exhibit both of these characteristics: They are powerful approximators and reliably generalize. This means that the deep learning theory community still has not identified a class of MLPs that reflects the analytic and statistical properties of real-world deep learning implementations.

This paper addresses this gap between theory and practice by identifying a sub-class of deep ReLU MLPs that exhibits both *optimal approximation rates* when approximating continuous functions on $[0, 1]^d$ and which generalizes well, even when overparameterized. Informally, our main finding is that deep ReLU MLPs of maximal regularity are optimal (universal) approximators. Our sub-class of maximally regular ReLU MLPs exhibits several familiar properties of real-world deep learning implementations, such as sample interpolation [35, 72] and small derivatives [50, 30].

We find that it is enough to analyze the sub-class of real-valued ReLU MLPs $\mathcal{NN}_{\Delta, W}^{L, n}$ on \mathbb{R}^d whose *size* is constrained by a prescribed depth Δ and width W , and whose *regularity* is limited by restricting their maximal Lipschitz constant. Furthermore, the trainable parameters in $\mathcal{NN}_{\Delta, W}^{L, n}$ are highly structured in that, all the weights and biases in their first and last layer are at most n , while in all other intermediate layers, all biases are 0 and all weights are in $\{0, \pm 1/2\}$.

Our main result shows that structural constraints, imposed by L and n , does not hinder the optimal approximation power of $\mathcal{NN}_{\Delta, W}^{L, n}$.

Theorem 1.0.1 (Optimal Approximation by ReLU MLPs with Maximal Regularity).

For any L -Lipschitz function $f : [0, 1]^d \rightarrow [-n, n]$, there exists an ***L-Lipschitz*** ReLU MLP Φ in $\mathcal{NN}_{W, D}^{L, n}$ with width $W \leq 8d(\lceil \frac{Ln}{2} \rceil + 1)^d$, depth $D \leq \lceil \log_2 d \rceil + 4$, at most $16d(\lceil \frac{Ln}{2} \rceil + 1)^d$ nonzero parameters, which satisfies

$$\max_{\mathbf{x} \in [0, 1]^d} |f(\mathbf{x}) - \Phi(\mathbf{x})| \leq \frac{d}{n}.$$

If we relax the Lipschitz and weights constraints in Theorem 1.0.1, by sending both L and n to infinity, then we recover the optimal approximation guarantees for *unconstrained* ReLU MLPs, which are now well-known in the approximation theory

literature. However, Theorem 1.0.1 shows that optimal approximation rates can be achieved while simultaneously imposing regularity on Φ .

Our main result, namely Theorem 1.0.1, confirmed that the class of MLPs $\mathcal{NN}_{\Delta, W}^{L, n}$ is rich/large enough to be a universal approximator. The statistical viability of this class is guaranteed by our second result, which shows this class is tame/small enough to generalize well.

We consider training data for a classification task $\{(X_n, Y_n)\}_{n=1}^N$ in $\mathbb{R}^d \times [0, 1]$. Our second result quantifies the (uniform) *generalization gap* for the class $\mathcal{NN}_{\Delta, W}^{L, n}$, which is defined as the largest absolute difference between the true risk $\mathcal{R}(\Phi)$, computed over the true distribution $(X_1, Y_1) \sim \mathbb{P}$, and the empirical risk $\mathcal{R}_N(\Phi)$, computed on the training data, for any ReLU MLP Φ in $\mathcal{NN}_{\Delta, W}^{L, n}$, where

$$\mathcal{R}(\Phi) \stackrel{\text{def.}}{=} \mathbb{E}_{(X, Y) \sim \mathbb{P}}[\ell(\Phi(X), Y)], \quad \mathcal{R}_N(\Phi) \stackrel{\text{def.}}{=} \frac{1}{N} \sum_{n=1}^N \ell(\Phi(X_n), Y_n).$$


Theorem 1.0.2 (Nearly Optimal Sample Complexity Without Explosion for Deep and Wide MLPs). *Given the sample set $\{(X_n, Y_n)\}_{n=1}^N$ where $\{X_n\}_{n=1}^N$ are i.i.d. centered, sub-Gaussian random variables with normalized covariance $\mathbb{E}[X_1 X_1^\top] = \frac{1}{N} \mathbf{I}_d$, and \mathbf{I}_d is the $d \times d$ identity matrix then, for each $\delta \in (0, 1)$*

$$\sup_{\Phi \in \mathcal{NN}_{\Delta, W}^{L, 1}} |\mathcal{R}(\Phi) - \mathcal{R}_N(\Phi)| \in \tilde{\mathcal{O}} \left(\frac{\sqrt{\log(4/\delta)}}{\sqrt{N}} + \min \left\{ \frac{W^{3\Delta/2}}{2^\Delta \sqrt{N}}, \frac{L^{d/(d+3)}}{d^{+3} \sqrt{N}} \right\} \right)$$

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

Our generalization bound in Theorem 1.0.2 guarantees a nearly optimal sample complexity of $\mathcal{O}(\log(N)/\sqrt{N})$ observed in parametric generalization bounds of ReLU MLPs, e.g. [54, 6, 7]. However, unlike parametric generalization bounds, our bound

in Theorem 1.0.2 converges even in the overparameterized regime where Δ and W are allowed to be arbitrarily large (compared to the sample size N). Typically, only non-parametric generalization bounds, such as e.g. [33], do not explode when Δ and W are taken to be arbitrarily large since they can account for the enumerable parametric symmetries in large neural networks; see e.g. [17, 2] for a discussion on these parametric symmetries. However, unlike the non-parametric bounds, the sample complexity of generalization bound for $\mathcal{NN}_{\Delta, W}^{L, n}$ converges at a *dimension-free* rate.



(a) Our Strategy - “Maximally Regular” Networks: The networks constructed in Theorem 1.0.1 interpolate a given grid linearly with the optimal slope between any two grid points. These ReLU MLPs constructed from finite sample values of the target function can never have Lipschitz constant (or, more generally, modulus of regularity, as shown in Lemma 7.5.4) exceeding that of the target function.

(b) Standard Approach - “Spiky” Networks: Following methods such as [83, 70], ReLU MLPs are constructed by first computing piecewise constant approximators of the target function matching its values at grid points (**flat black lines**) outside of so-called *trifling regions* (shaded in red). The piecewise constant approximators are then “glued together” via linear interpolation (steep red lines).

Figure 1.1: How our method works: Our method (left) against the usual approach (right) of [83, 70] for ReLU MLPs achieving the optimal convergence rate when approximating a (blue dashed) target function. Both methods memorize the value of the target function (**black dots**) at a specific set of grid points.

What issue does our construction resolve? The trouble in standard construction (right) is that as approximation becomes more accurate, these small trifling regions become very small, which can result in extremely steep red interpolating segments. Using our Kuhn triangulation-based construction (for the multi-dimensional case), we can construct optimal ReLU approximators with *no trifling regions*, meaning that we do not need these steep red segments; therefore, our ReLU approximators are not irregular.

Our construction of universal MLPs, with one-dimensional cartoons illustrated by Figure 1.1, relies on a new proof technique based on the *Kuhn triangulation* of the hypercube $[0, 1]^d$, see [41], and not on the trifling regions perfected by [70]. This new geometric construction allows us to construct approximating ReLU MLPs of maximal regularity. Surprisingly, Theorem 6.1.1 shows that there is no other triangulation (up to reflections) which can be used to construct an approximating ReLU MLP with a minimal Lipschitz constant; thus, our new construction is essentially unique.

Remark 1.0.1 (Approximation of Smooth Functions). The approximate piecewise polynomial variant of this construction sketched in Figure 1.1b, which applies for smooth enough functions and which was developed by [84] and subsequently refined by [85, 61, 46], uses approximate implementations of piecewise polynomials instead of piecewise constant functions “outside trifling regions”. Nevertheless, these methods still “glue together” several of these (approximate) local polynomial approximators using steep interpolators on these little “trifling regions”. Thus, they, too, still can have little regions where the resulting ReLU MLPs have very steep linear segments.

Remark 1.0.2 (Approximation in Lipschitz Norm). Theorem 1.0.1 guarantees that any bounded L -Lipschitz function can be uniformly approximated by bounded ReLU MLPs whose Lipschitz constant *never exceeds* L . Thus, the approximation happens within the compact (thus separable), by Arzelà-Ascoli theorem, set of uniformly bounded L -Lipschitz functions in the Banach space of continuous functions on the d -dimensional cube $C([0, 1]^d)$. Our result *does not* claim that one can approximate any Lipschitz function $f : [0, 1]^d \rightarrow \mathbb{R}$ with ReLU MLPs with respect to the Lipschitz norm $\|f\|_{Lip} \stackrel{\text{def.}}{=} \max_{\mathbf{x} \in [0, 1]^d} |f(\mathbf{x})| + \text{Lip}(f)$; which is impossible as the space of the Banach space of Lipschitz functions on $[0, 1]^d$ with this norm is not separable whereas

the set of ReLU MLPs is.

1.1 Further Results

Our analysis yields additional results concerning the expressivity of regular ReLU MLPs.

Memorization/Interpolation with ReLU MLPs of Minimal Lipschitz Constant The memorization (i.e. interpolation) problem for the MLP model dates back, at least, to [65, Theorem 5.1] where the author showed, amongst other things, that real-valued shallow MLPs with non-polynomial activation function and width n can memorize/interpolate $n + 1$ points in a domain of \mathbb{R}^d . The following is a shallow version of our main *universal interpolation* result (Theorem 7.4.1), which guarantees interpolation of N data points with minimal parameter usage and minimal Lipschitz constant is possible with ReLU MLPs depending on $\mathcal{O}(N)$ trainable parameters, constant depth, and $\mathcal{O}(\sqrt{N})$ width.

Theorem 1.1.1 (Optimal Interpolation with ReLU MLPs of Maximal Regularity (Shallow Version)). *Let $((x_n, y_n))_{n=1}^N$ be distinct pairs of training datapoints in $\mathbb{R} \times \mathbb{R}$ (with $x_n < x_{n+1}$ for $n = 1, \dots, N - 1$). There exists a ReLU MLP $\Phi : \mathbb{R} \rightarrow \mathbb{R}$ of width at most $2\lceil\sqrt{N}\rceil$, depth 2, and with at most $2N + 8\lceil\sqrt{N}\rceil$ nonzero parameters such that*

$$\Phi(x_n) = y_n \quad \text{for } n = 1, \dots, N$$

Furthermore, Φ is linear on the intervals $[x_n, x_{n+1}]$ for $n = 1, \dots, N - 1$, and constant on $(-\infty, x_1]$ and $[x_N, \infty)$.

This (shallow) version of our main memorization/interpolation theorem matches the optimal (in the sense of VC-dimension) parameter usage of a ReLU interpolator, as in [78]. The construction of [78], however, can easily be seen to have a large Lipschitz constant, whereas the Lipschitz constant of our interpolating ReLU MLPs cannot be improved as they are exactly the best piecewise linear interpolator of the training data. It is worth noting that, in [62, Theorem 9.6], the authors recently constructed interpolating ReLU MLPs, which also obtain the optimal Lipschitz constant (with respect to the ℓ^1 norm) in the multi-dimensional case. Those networks, however, require depth $\mathcal{O}(\log(N))$ and width $\mathcal{O}(N)$, making them sub-optimal in terms of parameter usage. Instead, the memorizers of Theorem 1.1.1 are optimal for both literature streams as they both have minimal parameter usage and minimal Lipschitz constant.

Our main one-dimensional *universal interpolation theorem* (Theorem 7.4.1) is a deep version of Theorem 1.1.1, where the user can *exactly specify* the depth $L \geq 4$ as well as the layer widths of the interpolating ReLU MLP, subject to the restriction that no MLP layer (bottleneck) is less than 12 and the total layer widths are at-least $\mathcal{O}(\sqrt{N}/L)$. What is most interesting about that result is that, unlike the memorizers constructed in [80], which require very large information bottlenecks if the training data points are close together, both Theorem 1.1.1 and its deep generalization in Theorem 7.4.1, show that no such restriction is needed. Moreover, the result shows that for most depth and width specifications, one can interpolate the training data.

Optimal Global Lipschitz Constant The full version of Theorem 1.0.1, namely Theorem 4.1.1, allows for functions of arbitrary regularity. Furthermore, Corollary 5.1.1 guarantees that Φ can be chosen to be globally L -Lipschitz, not only

L -Lipschitz on $[0, 1]^d$, by increasing its depth by 1. A consequence of this is Corollary 5.1.2, which shows that for any L -Lipschitz function there is an at-most L -Lipschitz ReLU MLP of comparable depth and width to that of Theorem 1.0.1, such that

$$\sup_{\mathbf{x} \in \mathbb{R}^d} \underbrace{|f(\mathbf{x}) - \Phi(\mathbf{x})|}_{\text{Approximation Error in } [0,1]^d} - L \underbrace{\max_{i=1, \dots, n} [(x_i - 1)_+ + (-x_i)_+]}_{\text{Extrapolation Error beyond } [0,1]^d} \lesssim \frac{1}{n}$$

where for each $u \in \mathbb{R}$, $u_+ \stackrel{\text{def.}}{=} \max\{0, u\}$ and \lesssim hides a dimensional constant of the order of $\mathcal{O}(Ld)$; here $\max_{i=1, \dots, n} [(x_i - 1)_+ + (-x_i)_+]$ is simply the ℓ^∞ distance to from a point $\mathbf{x} \in \mathbb{R}^d$ to the hypercube $[0, 1]^d$. This is an additive formulation of the multiplicative global approximation result of [13].

Best Achievable Approximation of Discontinuous Target Function Our main version (Theorem 4.1.1) of Theorem 1.0.2 even applies to “regular” but discontinuous target functions. In these cases, the approximation error does not converge to 0 but, rather, to the minimal achievable approximation error (details in Section 3.2.2). This latter result is rather interesting since it allows us to quantify the best *uniform* approximation of discontinuous target functions instead of having to rely on an L^p , for $1 \leq p < \infty$, relaxation of the notion of approximability or via randomization as in [40]. However, as one would expect from the Uniform Limit Theorem; see e.g. [51, Theorem 21.6], this approximation error *does not converge to 0* if the target function is genuinely discontinuous; however, we can quantify how small it can be made.

Additional Structure of the weights and biases in the class $\mathcal{NN}_{\Delta, W}^{L, n}$ Proposition 4.1.1 provides details on the weights and biases, as well as the encoding scheme, used to construct the universal sub-class of $\mathcal{NN}_{\Delta, W}^{L, n}$ described in Theorem 1.0.1.

1.1.1 Organization of Paper

Our paper is organized as follows. Section 2 overviews related results in the approximation theory literature, focusing on our main result (Theorem 1.0.2). Section 3 contains the preliminary notation, terminology, and background required for the formulation of our main results and their proofs. Section 4 contains our main results, these are the full version of our “regular” approximation theorem (Theorem 1.0.1) as well as the version of our main generalization bound (Theorem 4.2.1) with explicit constants.

Section 5 discusses consequences of our main result, such as extrapolation rates (Corollary 5.1.2), global Lipschitz regularity (Corollary 5.1.1). A discussion on how our result fit into the modern theory landscape is given in Section 6; where we also discuss the uniqueness (up to symmetries) of the Kuhn triangulation we used to construct our optimal ReLU MLPs of maximal regularity. This later result shows that our new geometric argument is essentially unique and cannot be improved.

All proofs are relegated to Section 7.

Chapter 2

Related Literature

2.1 Lipschitzness in Neural Networks

Neural networks of a prescribed Lipschitz regularity are common in various areas of deep learning; with applications ranging from generative adversarial learning [4, 12, 39], conditional distribution estimation [10], to certifiable deep learning [18], amongst many applications. This has led to several optimization pipelines, e.g. [54, 76, 66, 25, 58], and architectural designs, e.g. [45, 44, 3], enforcing Lipschitzness of trained neural networks. Various computational tools, e.g. [36, 11, 82], have also been developed to efficiently estimate the Lipschitz constant of neural networks. Despite extensive work, there is no guarantee that there will be no loss of model expressivity when imposing Lipschitz constraints; our main result fills this gap, thus adding additional theoretical foundations to deep learning areas relying on Lipschitz neural networks. We mention that the Lipschitz constant of (untrained) randomly initialized neural networks has recently been studied in [20].

2.2 Approximators with Parameter Restrictions

To the best of the authors' knowledge, the tightest available approximation results in the literature, which provide weight and bias size limits, are given in [61, Theorem 3.1]. There, the authors consider an L^p -type approximation, for $0 < p < \infty$, of any α -Hölder function on $[0, 1]^d$, with α -Hölder coefficient at-most 1 is shown to be possible using a ReLU MLP depending on $\mathcal{O}(\varepsilon^{-\alpha/d})$ parameters, organized into at-most $(2 + \lceil \log_2(\alpha) \rceil)(11 + \alpha/d)$ layers, and whose weights all belong to $[-\varepsilon^{-s}, \varepsilon^{-s}]$; where $s \geq 1$ is an integer depending at-least on the Hölder exponent (α) of the target function and on the dimension. For simplicity, examining the proofs of [61, Lemmata A.3] one sees that $s \geq 7$. Now, [57, Theorem 1] guarantees that the width of any universal (in the L^p sense) class of ReLU MLPs must have a width at least $d + 1$. Therefore, these observations, together with the fact that the operator norm $\|\mathbf{A}\|_{op:2 \rightarrow 2}$ of any $d \times d$ matrix by its componentwise 2-norm bound (and the elementary bound on $\|\cdot\|_2 \leq \|\cdot\|_1$): $\|\mathbf{A}\|_{op:2 \rightarrow 2} \leq d \max_{i,j=1,\dots,d} |\mathbf{A}_{i,j}|$ implies that the upper-bound on the Lipschitz constant $\text{Lip}(\hat{f})$ of ReLU MLP L^p ($p < \infty$) approximator \hat{f} which can be deduced from [61, Theorem 3.1] must be at least

$$\frac{d^{(2+\lceil \log_2(\alpha) \rceil)(11+\alpha/d)}}{\varepsilon^{7(2+\lceil \log_2(\alpha) \rceil)(11+\alpha/d)}} \geq \frac{d^{33}}{\varepsilon^{231+21\alpha/d}}. \quad (2.2.1)$$

Comparing the Lipschitz constant of a neural network whose weights achieve the lower-bound in (2.2.1), to the Lipschitz constant of our main result (Theorem 1.0.1) shows that the latter significantly improve the guarantees in the literature; as (2.2.1) diverges as the approximation error tends to 0 while our Lipschitz constant remain at the optimum.

Approximation of Lipschitz functions with values in $[-1, 1]$ by MLPs with the ReQU activation function and weights in $[-1, 1]$ is possible, however, there are no such guarantees for MLPs with the most standard ReLU activation function. We mention the work of [9] which controls the statistical properties of ReLU MLPs with bounded weights and biases using their best approximation (assuming it is exogenously bounded) and a correction term.

2.3 Approximation of Functions and Their Derivatives

The approximation of a function and its derivatives has drawn significant attention in the deep learning for partial differential equations (PDEs) literature [56, 28, 15, 14, 29, 43, 29, 47, 23, 68]. Guarantees that networks can approximate a function while also approximating its (at least first) derivative date back to [32] with more recent quantitative guarantees being given by [27, 28, 9, 52] (not all of which use the ReLU activation function). Recall that, by the mean value theorem, every once continuously differentiable function is Lipschitz on $[0, 1]^d$ with its Lipschitz constant given by the maximum norm of its gradient thereon; and, as a partial converse, every Lipschitz function on \mathbb{R}^d is differentiable almost-everywhere on $[0, 1]^d$; see [19, Theorem 3.1.6]. Thus, for the subclass of Lipschitz functions which are once continuously differentiable, these results guarantee that MLPs can approximate these maps on $[0, 1]^p$ while also approximating their Lipschitz constant.

However, there is no guarantee that MLPs can approximate these functions while exactly implementing their Lipschitz constant even locally. More generally, there is

no result that MLPs can approximate functions of lower regularity (e.g. Hölder or sub-Hölder functions) while also exactly matching their Lipschitz constant *globally*.

2.4 Global Universal Approximation

Theorem 4.1.1 provides exactly this guarantee showing, in particular, that ReLU MLPs can *globally* implement concave moduli of continuity of any uniformly continuous functions (e.g. Hölder or Lipschitz functions) while locally approximating them in $[0, 1]^d$. A global guarantee of the extrapolation rate for an approximation on $[0, 1]^d$ will also be provided in Corollary 5.1.2, and this is possible due to our guarantee that we may exactly and globally match the modulus of continuity of the target function being approximated.

We note that there do exist *qualitative* global approximation theorems in the deep learning [13], reservoir computing [26], and Stone-Weierstrass type [21] approximation literature. However, although each of those results is qualitative, it is currently not known what the extrapolation rates are for deep learning models, which are only guaranteed to provide a uniform approximation on $[0, 1]^d$.

2.5 Non-Linear Widths

In constructive approximation, one typically quantifies the “hardness to approximate” a compact class K of functions in $C([0, 1]^d)$ by the size, so-called widths, of the set K . Broadly speaking, most widths can be divided into either one of two classes. One of them are *distance*-type widths, where one measures the size of K in relation to its distance from relevant low-dimensional objects, describing what can be implemented

by the relevant numerical scheme. The other are *encoding-decoding*-type widths, where one quantifies the size of K by the recover-error incurred by first encoding its elements into a low-dimensional parameter space and subsequently decoding them back into the function space $C([0, 1]^d)$; all the while trying to match them losslessly.

When the objective is to compress K by projecting it down onto the closest $m \in \mathbb{N}_+$ -dimensional linear subspace L_m , yielding lower bounds on linear approximation algorithms, one obtains various notions of Kolmogorov linear widths $d_m^{lin}(K) \stackrel{\text{def.}}{=} \inf_{L_m} \sup_{f \in K} \inf_{g \in L_m} \|f - g\|_{L^\infty([0,1]^d)}$ of [38]; see e.g. [64]. However, as most contemporary numerical methods, such as sparse wavelet compression in signal processing or deep neural networks in machine learning, are non-linear, then one often considers non-linear quantification of the size of such classes K .

The main result of this paper can be interpreted as a highly non-linear width of distance type. Here, we are interested in the *best approximation* of a compact class of functions $K \subset C([0, 1]^d)$ by the class $\mathcal{NN}_{\Delta, W}^{L, n}$. Thus, we are studying the width

$$d_{\Delta, W}^{L, n}(K) \stackrel{\text{def.}}{=} \sup_{f \in K} \inf_{\Phi \in \mathcal{NN}_{\Delta, W}^{L, n}} \|f - \Phi\|_{L^\infty([0,1]^d)}. \quad (2.5.1)$$

This is a highly non-linear width since $\mathcal{NN}_{\Delta, W}^{L, n}$ is contained in the intersection of the infinite-dimensional compact (by Arzela-Ascoli) convex set $\text{Lip}(\mathbb{R}^d, \mathbb{R}, L)$ of real-valued functions on \mathbb{R}^d intersected with the finite-dimensional highly non-convex (see [60, Theorem 2.1]) but closed (see [59, Theorem 3.8]) set of ReLU MLPs of a given depth and width. Thus, $d_{\Delta, W}^{L, n}(K)$ is a highly non-convex analogue of the Kolmogorov $m \in \mathcal{O}(\Delta W^2)$ -dimensional linear widths, where m is the maximal number of parameters defining such a neural network.

2.6 Deficits of Parametric Generalization Bounds for Overparameterized MLPs

Under the assumptions of Theorem 1.0.2, the Rademacher complexity bounds of [8, Theorems 8 and 12] and [6], together with some result on random matrices with independent rows in [79, Theorem 4.6.1] (see Appendix 7.8 for details) can only imply the following much weaker bound: for each $0 < \delta < 1$ the following holds

$$\sup_{\Phi \in \mathcal{NN}_{\Delta, W}^{L, n}} |\mathcal{R}(\Phi) - \mathcal{R}_N(\Phi)| \in \mathcal{O}\left(\frac{\sqrt{\log(4/\delta)}}{\sqrt{N}}\right) + \tilde{\mathcal{O}}\left(\frac{1}{N^{3/2}} + \frac{W^{3\Delta/2}\sqrt{\log(4/\delta)}}{\sqrt{N}}\right) \quad (2.6.1)$$

with probability at least $1 - \delta$. We also note that, the nearly optimal VC-bounds derived in [7] for ReLU MLPs also diverge when W and Δ tend to infinity.

Chapter 3

Preliminaries

This section contains the preliminaries needed to formulate our results. These include both the notation used in this manuscript, as well as the background terminology relating to multilayer perceptrons.

3.1 Notation

We use $\mathbb{R}, \mathbb{N}, \mathbb{N}_+, \mathbb{Z}$ to denote the set of real numbers, non-negative integers, positive integers, and integers, respectively.

For any $x \in \mathbb{R}$, let $\lfloor x \rfloor \stackrel{\text{def.}}{=} \max\{n \in \mathbb{Z} : n \leq x\}$ denote the floor of x , and $\lceil x \rceil \stackrel{\text{def.}}{=} \min\{n \in \mathbb{Z} : n \geq x\}$ denote the ceiling of x . For any $n \in \mathbb{N}_+$, let $[n]$ denote the set of integers $\{1, 2, \dots, n\}$.

We denote the rectified linear unit (ReLU) activation function by $\sigma : \mathbb{R} \ni t \mapsto \max\{0, t\} \in \mathbb{R}$. It will always be applied componentwise to any vector, by which we mean $\sigma(\mathbf{x}) \stackrel{\text{def.}}{=} (\sigma(x_i))_{i=1}^d$ for each $\mathbf{x} = (x_1, \dots, x_d) \in \mathbb{R}^d$ and every $d \in \mathbb{N}_+$.

Vectors in dimensions above 2 are denoted by bold lowercase letters. Matrices are

denoted by bold uppercase letters.

3.2 Uniform Regularity

We first overview the standard notions of uniform continuity before introducing their discontinuous generalizations for which our approximation results hold.

3.2.1 Uniform, Lipschitz, and Hölder Continuity

A map $\omega : [0, \infty) \rightarrow [0, \infty)$ is called a modulus of continuity if ω is monotonically increasing and continuous (from the right) at 0 with $\omega(0) = 0$. Let (X, ρ_X) and (Y, ρ_Y) be metric spaces and $f : X \rightarrow Y$. We say that ω is a modulus of continuity for f if: for each $x, \tilde{x} \in X$

$$\rho_Y(f(x), f(\tilde{x})) \leq \omega(\rho_X(x, \tilde{x})).$$

We often consider uniformly continuous functions which are either Lipschitz and Hölder continuous. If $L \geq 0$, $0 < \alpha \leq 1$, and $\omega(x) = Lx^\alpha$, then we will say that f is α -Hölder with constant L . The class of all such $f : X \supseteq \Omega \rightarrow \mathbb{R}$ is denoted by $\text{Hol}(L, \alpha, \Omega)$. If $\alpha = 1$, we say that f is L -Lipschitz. Given any real-valued function f on a subset X of \mathbb{R}^d , we use $\text{Lip}(f)$ to denote its optimal Lipschitz constant with respect to the restriction of the ℓ^1 -norm to X ; i.e.

$$\text{Lip}(f) \stackrel{\text{def.}}{=} \sup_{\substack{\mathbf{x}, \mathbf{y} \in X \\ \mathbf{x} \neq \mathbf{y}}} \frac{|f(\mathbf{x}) - f(\mathbf{y})|}{\|\mathbf{x} - \mathbf{y}\|_1}.$$

We are mostly interested in the case where X is a subset of the space (\mathbb{R}^d, ℓ_1) and

(Y, ρ_Y) is a Euclidean line; where, for $1 \leq p \leq \infty$, ℓ_p indicates that we are equipping \mathbb{R}^d with the metric induced by the ℓ^p norm $\|\mathbf{x}\|_p \stackrel{\text{def.}}{=} \sum_{i=1}^p |x_i|^p$ if p is finite and $\|\mathbf{x}\|_\infty \stackrel{\text{def.}}{=} \max_{i=1, \dots, d} |x_i|$ otherwise.

3.2.2 Uniform Regularity, Best Achievable Approximation Error, and Discontinuity

We consider the following broad class of (possibly discontinuous) target functions, including uniformly continuous functions. The reader which is only interested in approximation guarantees for continuous functions, *where an asymptotic error of 0 is achievable by ReLU MLPs*, is encouraged to skip this section and swap the term “modulus of regularity” for “concave modulus of continuity” in the remainder of the paper. We rely on the following weakened version of a concave modulus of continuity.

Definition 3.2.1 (Modulus of Regularity). Let $I = [0, \infty)$ or $I = [0, T]$ for some $T \geq 0$. A function $\omega : I \rightarrow [0, \infty)$ is called a modulus of regularity if:

1. $\omega(0) = 0$,
2. ω is monotone increasing,
3. ω is concave.

Moduli of regularity allow us to quantify the best achievable approximation error, when uniformly approximating discontinuous target functions. Before formalizing this, we recall that concave functions on bounded domains $\Omega \subset \mathbb{R}^d$ are continuous on their interior but need not be on the boundary $\partial\Omega$ of Ω . For example, the indicator function $I_{(0, \infty)}$ of $(0, \infty)$ is concave on $\Omega = [0, \infty)$, continuous on $(0, \infty)$, and has a

discontinuity on $\partial\Omega = \{0\}$. We will often consider the *best achievable approximation error* $\omega^* \in [0, \infty)$, when minimizing such a modulus, which is given by

$$\omega^* \stackrel{\text{def.}}{=} \lim_{t \downarrow 0} \omega(t).$$

Example 3.2.1 (Concave Moduli of Continuity are Moduli of Regularity). If $\omega : [0, \infty) \rightarrow [0, \infty)$ is a concave modulus of continuity, then ω is a modulus of regularity, with $\omega(0) = 0$, and ω is continuous at 0; i.e. $\lim_{t \downarrow 0} \omega(t) = 0$ and in particular $\omega^* = 0$.

The previous example illustrates the main difference between moduli of continuity and moduli of regularity; namely, the value of 0 need not be achievable by minimizing ω from the right. Indeed, if ω is a modulus of regularity, then it is non-negative and it fixes 0; thus, $\omega(t) \geq \omega(0) = 0$ for all positive values of t . Consequentially,

$$\omega(0) = 0 \leq \omega^* \tag{3.2.1}$$

for any modulus of regularity. This inequality can be strict for discontinuous functions.

Example 3.2.2 (Beyond Moduli of Continuity). Fix $M > 0$ and $L \geq 0$. The function $\omega : [0, \infty) \rightarrow [0, \infty)$ given for each $t \in [0, \infty)$ by

$$\omega(t) \stackrel{\text{def.}}{=} \begin{cases} M + Lt & \text{if } t > 0 \\ 0 & \text{if } t = 0 \end{cases}$$

is a modulus of regularity but not a modulus of continuity, and the inequality in (3.2.1) is strict.

The class of continuous functions between metric spaces are precisely those functions for which: for every compact subset of their domain, there is a modulus of continuity bounding the distance between the images of all pairs of points therein, see e.g. [51, Theorem 27.6], and such that *the images of arbitrarily close points are themselves arbitrarily close*.

If one relaxes the second condition, we are left with the following functions considered in our paper. Our interest in this class stems from (3.2.1), wherein our main result (Theorem 4.1.1) shows that for ω -regular function can be approximated by a ReLU MLP with at-most ω -regularity, up to the best achievable approximation error ω^* . In the very special case where the target function is additionally continuous, we can guarantee the usual type of conclusion: our ReLU MLP approximators can achieve *zero approximation error* asymptotically.



Figure 3.1: The heaviside function h is ω -regular with $\omega(t) = 1$ if $t > 0$ and 0 otherwise.

Definition 3.2.2 (ω -Regular Functions). Let ω be a modulus of regularity, (X, ρ) and (Y, ρ') be metric spaces. A map $f : X \rightarrow Y$ is said to have ω -regularity at a pair of points $x, x' \in X$ if

$$\rho'(f(x), f(x')) \leq \omega(\rho(x, x')).$$

If, moreover, f has ω -regularity at all pairs of points $x, x' \in X$ then, we say that f is

ω -regular (on X). The set of all ω -regular functions from (X, ρ) to $(\mathbb{R}, |\cdot|)$ is denoted by $\text{Lip}_\omega(X)$.

We will say that a function is *regular* if it is ω -regular for some modulus of regularity ω . Geometric examples of discontinuous regular functions include quasi-isometric [42] and certain coarse embeddings [55, 48]. Discontinuous regular maps can also be constricted as additive perturbations of Lipschitz functions by some “deterministic bounded noise”, as follows.

Example 3.2.3 (Additive Perturbations of Lipschitz Functions). Fix $L, M \geq 0$ and fix any function $\eta : \mathbb{R} \rightarrow [-M/2, M/2]$. For any L -Lipschitz function $f : \mathbb{R}^d \rightarrow \mathbb{R}$ define the perturbed function $f_\eta : \mathbb{R}^d \rightarrow \mathbb{R}$ as sending any $\mathbf{x} \in \mathbb{R}^d$ to

$$f_\eta(\mathbf{x}) \stackrel{\text{def.}}{=} f(\mathbf{x}) + \eta(\mathbf{x}).$$

Then, f_η is ω -regular with modulus of regularity as in Example 3.2.2.

Regular functions need not be well-behaved measure-theoretically. For instance, in the context of Example 3.2.3 if A is the Vitali set in \mathbb{R} and $\eta = \frac{M}{2} I_A$ then we see that f_η need not even be Lebesgue measurable even if f was Lipschitz.

Example 3.2.4 (Not all Functions are Regular). Not all functions are regular. The following map is not regular: $f : [0, 1] \rightarrow [0, \infty)$ given for each $x \in [0, 1]$ by

$$f(x) \stackrel{\text{def.}}{=} \begin{cases} \frac{1}{x} & \text{if } x > 0 \\ 0 & \text{if } x = 0. \end{cases}$$

However, any bounded function is regular. Nevertheless, their best achievable

approximation error can be large.

3.3 Multilayer Perceptrons (MLPs) with ReLU Activation Function

We now define multilayer perceptions. We then formulate notions of parameteric and functional regularity often encountered in the literature considered herein.

Definition 3.3.1 (Multilayer Perceptrons with ReLU Activation Function (ReLU MLPs)). Let $\Delta \in \mathbb{N}_+$ and consider a multi-index $\mathbf{d} \stackrel{\text{def.}}{=} [d_1, \dots, d_{\Delta+1}] \in \mathbb{N}_+^d$. The class $\mathcal{NN}(\mathbf{d})$ consists of all multilayer perceptrons with ReLU activation function (ReLU MLPs) $\Phi : \mathbb{R}^{d_1} \rightarrow \mathbb{R}^{d_{\Delta+1}}$ admitting the following iterative representation

$$\begin{aligned} \Phi(\mathbf{x}) &= \mathbf{W}^{(\Delta)} \mathbf{x}^{(\Delta)} + \mathbf{b}^{(\Delta)} \\ \mathbf{x}^{(l+1)} &\stackrel{\text{def.}}{=} \text{ReLU} \bullet (\mathbf{W}^{(l)} \mathbf{x}^{(l)} + \mathbf{b}^{(l)}) \quad \text{for } l = 1, \dots, \Delta - 1 \\ \mathbf{x}^{(1)} &\stackrel{\text{def.}}{=} \mathbf{x}. \end{aligned} \tag{3.3.1}$$

where for $l = 1, \dots, \Delta$, $\mathbf{W}^{(l)}$ is a $d_{l+1} \times d_l$ -matrix and $\mathbf{b}^{(l)} \in \mathbb{R}^{d_{l+1}}$, and $\text{ReLU} \bullet$ denotes componentwise application of the ReLU function. We denote $\text{widthvec}(f) \stackrel{\text{def.}}{=} [d_2, \dots, d_{\Delta}]$.

Given $L, W, d, D \in \mathbb{N}_+$, we use $\mathcal{NN}(d, D)$ to denote the class of maps $f : \mathbb{R}^d \rightarrow \mathbb{R}^D$, $\mathcal{NN}_{\Delta, W}(d, D)$ the subset of $\mathcal{NN}(d, D)$ of ReLU MLPs with depth at-most Δ and width at-most W , and $\mathcal{NN}_{\Delta, W}^L(d, D)$ the subset of $\mathcal{NN}_{\Delta, W}(d, D)$ of L -Lipschitz ReLU MLPs therein. When clear from the context, we suppress the notational dependence on d and D .

Given a ReLU MLP f with representation 3.3.1, we call the integers $\Delta - 1$ and $\max_{l=2,\dots,\Delta} d_l$ its depth and width, respectively. Let $P(\mathbf{d}) \stackrel{\text{def.}}{=} \sum_{l=1}^{\Delta} d_l(d_{l+1} + 1)$. When discussing the stability of the Φ on the *parameters* defining it via this representation, we will rely on the following vectorization of the parameters defining its weight matrices (the $\mathbf{W}^{(l)}$) and its biases (the $\mathbf{b}^{(l)}$)

$$\mathbb{R}^{P(\mathbf{d})} \ni \theta \iff (\mathbf{W}^{(l)}, \mathbf{b}^{(l)})_{l=1}^{\Delta} \in \prod_{l=1}^{\Delta} (\mathbb{R}^{d_{l+1} \times d_l} \times \mathbb{R}^{d_{l+1}}). \quad (3.3.2)$$

In particular, this vectorization of the parameters of Φ allows us to define the maximum norm of a set of weights and biases defining Φ , in representation (3.3.1), as the ℓ^∞ norm of its parameter vector θ in (3.3.2); via

$$|\Phi|_{\text{par}} \stackrel{\text{def.}}{=} \max_{i=1,\dots,P(\mathbf{d})} |\theta_i|, \quad (3.3.3)$$

where the operation $|\Phi|_{\text{par}}$ is defined *given* the representation of Φ in (3.3.1). That is, different *representations* of the same function Φ may have different maximum parameter sizes, meaning that, $|\Phi|_{\text{par}}$ is not a canonical intrinsic quantification of the regularity of Φ . However, this is *not* the case for the optimal Lipschitz constant of Φ which is independent of any parameterization thereof. Thus, the optimal Lipschitz constant of Φ is a more natural measure of regularity of Φ than norms on its parameter vectors; given a particular representation.

Chapter 4

Main Results

We now present our main approximation result (Theorem 1.0.1) and then the full-version of our statistical guarantee (Theorem 4.2.1).

4.1 Approximation Guarantees

The following is the full version of our optimal approximation theorem with maximal regularity.

Theorem 4.1.1 (Optimal Regular Approximation by Sample-Interpolating ReLU MLPs). *Let f be a function from $[0, 1]^d$ to \mathbb{R} , and let $\omega : [0, d] \rightarrow [0, \infty)$ be a modulus of regularity of f . Then, for any $n \in \mathbb{N}_+$, there exists a ReLU MLP Φ on $[0, 1]^d$ with width at most $8d(n + 1)^d$, depth at most $\lceil \log_2 d \rceil + 4$, at most $16d(n + 1)^d$ nonzero parameters, which satisfies the approximation guarantee*

$$\|f - \Phi\|_{L^\infty([0,1]^d)} \leq \omega\left(\frac{d}{2n}\right)$$

Furthermore, ω is a modulus of regularity of Φ on $[0, 1]^d$.

Remark 4.1.1. Theorem 4.1.1 provided guarantees on the modulus of regularity of Φ on $[0, 1]^d$; this guarantee can be extended to all of \mathbb{R}^d at a cost of only a factor of $\mathcal{O}(1)$ more non-zero parameters (Corollary 5.1.1 below). In this way, one may obtain a global extrapolation variant of this result (Corollary 5.1.2 below).

The ReLU MLPs constructed in our main approximation theorem are constructed from a finite number of samples from the target function at specific grid points. Given *degree of freedom* $n \in \mathbb{N}_+$, we draw

$$N(n) \stackrel{\text{def.}}{=} (1 + n)^d$$

samples from the target function the following grid points in $[0, 1]^d$

$$\mathbb{X}_n \stackrel{\text{def.}}{=} \{\mathbf{x} \in [0, 1]^d : x_1, \dots, x_d \in \{j/n\}_{j=0}^n\} \quad (4.1.1)$$

We define the sampling-based encoder \mathcal{E}_n by sending any function $f : \mathbb{R}^d \rightarrow \mathbb{R}$ to the following latent code in $\mathbb{R}^{N(n)}$

$$\mathcal{E}(f) \stackrel{\text{def.}}{=} (f(\mathbf{x}_i))_{i=1}^{N(n)}.$$

Our main universal approximation theorem guarantee, generalizing Theorem 1.0.2, shows that our regular ReLU approximator of f can be constructed from the $N(n)$ sample values $\mathcal{E}(f)$, and its construction only relies on the information of f contained in these samples.

We rephrase the above main theorem in the context of encoder-decoder as follows.

Theorem 4.1.2 (Theorem 4.1.1, Encoder-Decoder Formulation). *Let ω be a modulus of regularity. For every $n \in \mathbb{N}_+$, there exists a decoder $\mathcal{D}_n : \mathbb{R}^{N(n)} \rightarrow \mathcal{NN}(d, D)$ such that: for each $f \in \text{Lip}_\omega([0, 1]^d)$, ω is a modulus of regularity of the ReLU MLP $\Phi \stackrel{\text{def.}}{=} \mathcal{D}_n \circ \mathcal{E}_n(f)$ on $[0, 1]^d$ which has width at most $8d(n+1)^d$, depth at most $\lceil \log_2 d \rceil + 4$, at most $16d(n+1)^d$ nonzero parameters, and it satisfies the approximation guarantee*

$$\max_{\mathbf{x} \in [0, 1]^d} |f(\mathbf{x}) - \Phi(\mathbf{x})| \leq \omega\left(\frac{d}{2n}\right)$$

as well as the sample-interpolation guarantee

$$f(\mathbf{x}_i) = \Phi(\mathbf{x}_i) \quad \text{for each } i = 1, \dots, N(n).$$

Moreover, the parameter norm, path-norm, are all recorded in Table 4.1.

Description of MLP	Estimate
Parameteric Complexity	
Depth	$\lceil \log_2(d) \rceil + 4$
Width	$8d(n + 1)^d$
Nonzero parameters	$16d(n + 1)^d$
Regularity	
Modulus of Regularity	$\Phi \in \text{Lip}_\omega(\{\mathbf{x}_i\}_{i=1}^{N(n)})$
Parameteric Regularity $ \Phi _{\text{par}}$	$\max\{n, \max_{i \in [N(n)]} f(\mathbf{x}_i) \}$
Misc. Structure	
Encoding Dimension $\dim(\text{dom}(\mathcal{D}))$	$(n + 1)^d$
Support	$\text{supp}(\Phi) \subseteq [-1/n, 1 + 1/n]^d$

Table 4.1: ReLU MLP (Φ) constructed in Theorem 4.1.2 approximating $f \in \text{Lip}_\omega(\{\mathbf{x}_i\}_{i=1}^{N(n)})$ on $[0, 1]^d$: All maxima are indexed over i in $\{1, \dots, N(n)\}$. The depth, width, and maximum parameter size $|\Phi|_{\text{par}}$ are defined in (3.3.3). The regularity of the encoder is with respect to the ℓ^∞ norm on $\mathbb{R}^{N(n)}$. By the support $\text{supp}(\Phi)$ of Φ , we mean the set points in \mathbb{R}^d where it is non-zero.

Note that, the path-norm bound in Table 4.1 implies that most weights defining the ReLU MLP must be small. Indeed, a detailed inspection of the construction of Φ , see Remark 7.6.1, shows that most weights and biases defining Φ are in $\{0, \pm 1/2\}$.

A closer look at the proof of Theorem 4.1.1 reveals much more information about the structure of the weights and biases of the approximating ReLU MLP constructed therein, as well as their dependence on the target function. We summarize these additional facts in the following Proposition; which is proven during the course of our main result. The following result thus provides in-depth details of the structure of

the decoder mapping the latent code $\mathcal{E}(f)$ to our constructed ReLU MLP.

Proposition 4.1.1 (Estimates for the Weights and Biases in the Construction of Φ in Theorem 4.1.1). *Consider the setting of Theorem 4.1.1, represent Φ as in (3.3.1); note that $1 < \Delta \leq \lceil \log_2 d \rceil + 4$.*

- (i) **Target Dependant Layers:** *The weights $\mathbf{W}^{(\Delta)}$ depend on f ,*
- (ii) **Target Independents Layers:** *The weights $\{\mathbf{W}^{(l)}\}_{l=1}^{\Delta-1}$ do not depend on f but only depend on d , and the biases $\{\mathbf{b}^{(l)}\}_{l=2}^{\Delta}$ are all 0,*
- (iii) **Typical Weights:** *For each $l = 2, \dots, \Delta - 1$, $\mathbf{W}_{i,j}^{(l)} \in \{0, \pm 1/2\}$ for each i, j ,*
- (iv) **Initial and Terminal Weights:** $\|\mathbf{W}^{(1)}\|_{\ell^\infty} \leq 1, \|\mathbf{W}^{(\Delta)}\|_{\ell^\infty} \leq \max_{i \in [N(n)]} |f(\mathbf{x}_i)|, \|\mathbf{b}^{(l)}\|_{\ell^\infty} \leq n$.

We conclude this section by noting that the shape of the optimal ReLU MLPs in Theorem 4.1.1 can be very flexible. As the following variant of that result shows, we may specify ReLU MLPs of virtually any width or depth to obtain our optimal approximation guarantee with maximal regularity.

Proposition 4.1.2 (Theorem 4.1.1 with Variable Width and Depth). *Let f be a function from $[0, 1]^d$ to \mathbb{R} , and let $\omega : [0, d] \rightarrow [0, \infty)$ be a modulus of regularity of f . Then, for any master parameter $n \in \mathbb{N}_+$, depth parameter $L \in \mathbb{N}_+$, and width parameters $m_1, \dots, m_L \in \mathbb{N}_+$ satisfying*

$$m_1 + \dots + m_L = (n + 1)^d$$

there exists a ReLU MLP Φ on $[0, 1]^d$ with width at most $8d \max\{m_1, \dots, m_L\} + d + 2$, depth at most $L(\lceil \log_2 d \rceil + 4)$, at most $16d(n + 1)^d + L(d + 2)$ nonzero parameters,

which satisfies the approximation guarantee

$$\|f - \Phi\|_{L^\infty([0,1]^d)} \leq \omega\left(\frac{d}{2n}\right)$$

Furthermore, ω is a modulus of regularity of Φ on $[0, 1]^d$.

In Corollary 4.1.2, if we take $L = d$ and $m_1 = m_2 = \dots = m_d = \lceil (n+1)^d/d \rceil$, then we can conclude that the same approximator Φ in Theorem 4.1.1 can be implemented by another ReLU MLP with width $\mathcal{O}(n^d)$ and depth $\mathcal{O}(d \log d)$.

4.2 Statistical Guarantees

Next, we investigate the learning theoretic implications of our main result. We consider i.i.d. training data for a classification problem $(X_1, Y_1), \dots, (X_N, Y_N) \sim \mathbb{P}$ in $\mathbb{R}^d \times [0, 1]$ and a loss function $\ell : \mathbb{R} \times \mathbb{R} \rightarrow [0, 1]$. Our objective is to describe the worst-case *generalization gap* for any ReLU MLP Φ in the class studied in Theorem 4.1.1. The generalization gap for any such MLP is defined as the absolute difference between the true risk $\mathcal{R}(\Phi)$ over the (unseen) test set and the empirical risk $\mathcal{R}_N(\Phi)$ computed on the training data; where

$$\mathcal{R}(\Phi) \stackrel{\text{def.}}{=} \mathbb{E}_{(X,Y) \sim \mathbb{P}}[\ell(\Phi(X), Y)], \quad \mathcal{R}_N(\Phi) \stackrel{\text{def.}}{=} \frac{1}{N} \sum_{n=1}^N \ell(\Phi(X_n), Y_n).$$

Theorem 4.1.1 and Proposition 4.1.1 identifies a highly structured subclass of the set of neural networks with depth Δ and width W , for $\Delta, W \in \mathbb{N}_+$, which are *expressive enough* to approximate any L -Lipschitz functions with range in $[0, 1]$ to a uniform precision of $\omega(d/2n)$ for any prescribed $n \in \mathbb{N}_+$. Specifically, the class $\mathcal{NN}_{\Delta, W}^{L, n}$ consists

of all ReLU MLPs $\Phi : \mathbb{R}^d \rightarrow \mathbb{R}$ in $\mathcal{NN}_{\Delta, W}^L(d, 1)$, see Definition 3.3.1, for which the weights $\mathbf{W}^{(1)}, \dots, \mathbf{W}^{(\Delta)}$ in (3.3.1) satisfy

$$\|\mathbf{W}^{(1)}\|_\infty \leq 1, \max_{l=2, \dots, L-1} \|\mathbf{W}^{(l)}\|_\infty \leq \frac{1}{2}, \text{ and } \|\mathbf{W}^{(L)}\|_\infty \leq n. \quad (4.2.1)$$

Assumption 4.2.1 (Normalized Sub-Gaussian Training Data). The $\mathbb{R}^d \times [0, 1]$ -valued random variables $(X_1, Y_1), \dots, (X_N, Y_N)$ are i.i.d and X_1 is centered, sub-Gaussian, with normalized covariance

$$\mathbb{E}[X_n X_n^\top] = \frac{1}{N} \mathbf{I}_d$$

and are defined on a common probability space $(\Omega, \mathcal{A}, \mathbb{P})$; where \mathbf{I}_d is the $d \times d$ identity matrix.

Unlike the generalization bound in (2.6.1), deep and wide neural networks generalize well in practice. On the contrary, several recent results studying gradient dynamics of “infinitely wide” neural networks suggest that highly overparameterized neural networks (which may be extremely deep and wide) tend to generalize well and tend to converge to highly regular networks after being optimized by gradient descent on the training data. This is not captured by the generalization bound in (2.6.1) since, when either Δ or W become large, then the generalization bound in (2.6.1) diverges, even if L is held constant.

The divergence of (2.6.1) is counter-intuitive since $\mathcal{NN}_{\Delta, W}^{L, n}$ is contained in the class of L -Lipschitz functions on \mathbb{R}^d with image in $[0, 1]$; which is totally bounded by Arzelà-Ascoli theorem, and thus its Rademacher complexity should be controllable by Dudley’s entropy integral estimate (see e.g. [77, Corollary 2.2.8]). The reason is that the bound in (2.6.1) is inherited from bounding the Rademacher complexity of the

larger class $\mathcal{NN}_{\Delta,W}$ which becomes unbounded as the depth and width parameters grow. Upon noting that the Rademacher complexity of the class $\mathcal{NN}_{\Delta,W}^{L,n}$ must be no larger than that of the class of L -Lipschitz functions with values in $[0, 1]$ and no larger than that of the class $\mathcal{NN}_{\Delta,W}$, then using the bound on the former computed for instance in [33, Lemma 25], we may improve (2.6.1) so that it remains bounded as Δ and W tend to infinity.

Theorem 4.2.1 (Non-Exploding Generalization Bounds for Arbitrarily Deep and Wide ReLU MLPs). *Let $\Delta, W, n, d \in \mathbb{N}_+$, $L, L_\ell > 0$, and consider a L_ℓ -Lipschitz loss function $\ell : \mathbb{R}^2 \rightarrow [0, 1]$. Under Assumption 4.2.1, for each $\delta \in (0, 1)$ the generalization gap $\sup_{\Phi \in \mathcal{NN}_{\Delta,W}^{L,n}} |\mathcal{R}(\Phi) - \mathcal{R}_N(\Phi)|$ is bounded above by*

$$\underbrace{\frac{\sqrt{8 \log(4/\delta)}}{\sqrt{N}}}_{\text{Prob. Satisfaction}} + 2L_\ell \min \left\{ \underbrace{\frac{4}{N^{3/2}} + \frac{26 \log(N) \log(2W) W^{3\Delta/2}}{2^{\max\{0, \Delta-2\}} N}}_{\text{Parameter Space}} (\sqrt{d} + C_{d, X_1} \sqrt{N} + C_{d, X_1} \sqrt{\ln(4/\delta)}), \underbrace{C_d \frac{L^{d/3}}{N^{1/3}}}_{\text{Function Space}} \right\}$$

with probability at least $1 - \delta$; where $C_d \stackrel{\text{def.}}{=} (8(d+1)^2(16)^d)^{1/(d+3)} + \frac{2^{5/2} 16^{d/(d+3)}}{(18(d+1))^{(d+1)/(d+3)}}$ and $C_{d, X_1} > 0$ is a constant depending only on X_1 and d .

Chapter 5

Implications of Theorem 4.1.1

We now showcase the breadth of Theorem 4.1.1 by examining some of its implications from approximation theory to learning theory.

5.1 Approximation Theory: Extrapolation Rates

Our main result guarantees that the approximating ReLU MLP Φ has a (modulus of) regularity which is no worse than that of the target function. This allows us to deduce a sharp extrapolation rate for our approximator outside of $[0, 1]^d$.

We can thus directly obtain a version of a *global universal approximation theorem*. However, our result differs from that of [13], which leverage topologies stronger than the uniform convergence on compacts topology on $C(\mathbb{R}^d)$ but weaker than the uniform topology thereon on suitable subspaces; e.g. the strict topology on continuous bounded functions studied by [75, 21, 74]. Instead, we do not discount/compress our error using weight functions outside the unit cube but rather we quantify how fast it grows, using the moduli of *both* the target function and its neural network

approximator.

The first step towards this result is given by our first consequence of Theorem 4.1.1. Namely, this first corollary shows that we can easily extend the domain of the interpolator Φ in Theorem 4.1.1 to all of \mathbb{R}^d while maintaining its regularity. To formulate this consequence, we consider the distance $\text{dist}_{[0,1]^d}$ from any point to the d -dimensional unit cube, defined for each $\mathbf{x} \in \mathbb{R}^d$ by

$$\text{dist}_{[0,1]^d}(\mathbf{x}) \stackrel{\text{def.}}{=} \min_{\mathbf{z} \in [0,1]^d} \|\mathbf{z} - \mathbf{x}\|_1.$$

Corollary 5.1.1 (Theorem 4.1.1 with Optimal Global Regularity). *For any $n \in \mathbb{N}_+$ and $f : [0, 1]^d \rightarrow \mathbb{R}$ with modulus of regularity $\omega : [0, d] \rightarrow \mathbb{R}$, there exists a ReLU MLP $\Phi : \mathbb{R}^d \rightarrow \mathbb{R}$ with width at most $8d(n + 1)^d$, depth at most $\lceil \log_2 d \rceil + 5$, and no more than $18d(n + 1)^d$ nonzero parameters such that*

$$\|f - \Phi\|_{L^\infty([0,1]^d)} \leq \omega\left(\frac{d}{2n}\right)$$

Moreover, the following extension of ω is a modulus of regularity of Φ on \mathbb{R}^d :

$$\bar{\omega}(x) \stackrel{\text{def.}}{=} \begin{cases} \omega(x), & \text{if } x \in [0, d] \\ \omega(d), & \text{if } x > d \end{cases}, \quad \forall x \in [0, \infty) \quad (5.1.1)$$

Using Corollary 5.1.2, we are able to deduce the following “global universal approximation theorem” with *additive correction term*, instead of the usual *multiplicative corrective weight* used in the global approximation literature [13, Theorem 4.13], in the reservoir computing literature, e.g. [26, Corollary 9], or in the deep learning for dynamical systems literature; e.g. [1, Theorem 4.11].

Corollary 5.1.2 (Extrapolation Bounds for ReLU MLP Approximators). *For any $n \in \mathbb{N}_+$ and $f : \mathbb{R}^d \rightarrow \mathbb{R}$ with modulus of regularity $\omega : [0, d] \rightarrow \mathbb{R}$ on $[0, 1]^d$, there is a ReLU MLP $\Phi : \mathbb{R}^d \rightarrow \mathbb{R}$ with width at most $8d(n+1)^d$, depth at most $\lceil \log_2 d \rceil + 5$, and no more than $18d(n+1)^d$ nonzero parameters satisfying*

$$\sup_{\mathbf{x} \in \mathbb{R}^d} \underbrace{|f(\mathbf{x}) - \Phi(\mathbf{x})|}_{\text{Approximation in } [0,1]^d} = \underbrace{\bar{\omega}(\text{dist}_{[0,1]^d}(\mathbf{x}))}_{\text{Extrapolation: beyond } [0,1]^d} \leq \bar{\omega}\left(\frac{d}{2n}\right).$$

where $\bar{\omega}$ is defined in (5.1.1). Furthermore, $\bar{\omega}$ is a modulus of regularity of Φ on \mathbb{R}^d .

Chapter 6

Discussion

We now discuss some technical points of our main result, explaining how certain steps in our construction cannot be improved on. The proofs of each of these results are relegated to Section 7.9 below.

6.1 The Kuhn Triangulation is Unique for Regular Approximation

The multi-dimensional analogue of the cartoon of our one-dimensional regular ReLU approximator illustrated in Figure 1.1a, relies on a specific triangulation of the d -dimensional cube $[0, 1]^d$. This triangulation called the *Kuhn triangulation*, defined shortly afterwards, is the most “regular” triangulation in the sense that it is the only triangulation whose continuous piecewise linear approximators always preserve the regularity of the function that it approximates under some mild conditions; see [37, Lemma 1]. This triangulation is used in one of the key technical steps in proving

Theorem 4.1.1; namely, Lemma 7.5.4.




Figure 6.1: Kuhn triangulations of a 2D (left) and a 3D (right) unit cube.

Definition 6.1.1 (Kuhn triangulation). For any $d \in \mathbb{N}_+$, the $d!$ different simplices

$$\text{conv} \left(\left\{ \sum_{i=1}^k \mathbf{e}_{\tau(i)} : 0 \leq k \leq d \right\} \right), \quad \tau \in S_d \quad (6.1.1)$$

form a triangulation of $[0, 1]^d$, where S_d is the symmetric group on d letters, and $\mathbf{e}_1, \dots, \mathbf{e}_d$ are the standard basis vectors of \mathbb{R}^d . This triangulation is called the Kuhn triangulation.

It is natural to ask if another triangulation could have been used during the proof of our main result. Surprisingly, this is not the case as any other triangulation, other than reflections of the Kuhn triangulation, yield approximators with larger Lipschitz constants. This is the content of the following result which serves as a type of converse to Lemma 7.5.4.

Theorem 6.1.1 (Uniqueness of the Kuhn Triangulation - up to Reflections). *The Kuhn triangulation is the only triangulation (up to reflections) that makes Lemma 7.5.4 true for the case $n = 1$.*

6.2 Parameter Usage Compared to State-of-the-Art Approximation

This section contains a statement and derivation of the main result of the state-of-the-art (SOTA) approximation theorem for ReLU MLPs [70] which *includes estimates on the Lipschitz constant of the ReLU MLP constructed therein*. We emphasize that the optimality criterion in the following approximation theorem and its predecessor was parameter usage, not regularity.

Theorem 6.2.1 (SOTA Universal Approximation Theorem of [70] with Regularity Quantification). *For any $2 \leq n \in \mathbb{N}_+$ and $f \in \text{Hol}(\nu, \alpha, [0, 1]^d)$ with $\alpha \in (0, 1]$, we have:*

1. *If $d = 1$, then $\exists \Phi \in \mathcal{NN}(\#input = 1; \text{widthvec} = [2n + 1, 2n])$ with $\mathcal{O}(n^2)$ nonzero parameters such that*

$$\|f - \Phi\|_{L^\infty([0,1])} \leq \frac{\nu}{(2n^2)^\alpha}$$

$$\Phi \in \text{Hol}(\nu, \alpha, [0, 1])$$

2. *If $d > 1$, then $\exists \Phi \in \mathcal{NN}(\#input = d; \text{width} \leq \max\{(2n + 3)d, 6d + 3, 2[(n + 1)^{d/2}] + 2d\}; \text{depth} \leq 23d + 9)$ with $\mathcal{O}(dn^d)$ nonzero parameters such that*

$$\|f - \Phi\|_{L^\infty([0,1]^d)} \leq \nu \left(\frac{d^2}{n(2d + 1)} \right)^\alpha$$

$$\text{Lip}(\Phi) \leq \nu(n + 1)^d(2d + 1)d^\alpha$$

Observe that the minor improvement of the main result of [70], given in Theorem 6.2.1, yields significantly more irregular MLPs than our main theorem due to sharp descents between trifling regions (illustrated by the red bands in Figure 1.1b). In contrast, our Kuhn triangulation-based construction resolves this issue optimally since the ReLU MLPs constructed using it do not require any trifling region.

Remark 6.2.1 (Improvement of Constant for Variant of [70, Theorem 1.1]). In the multi-dimensional case, our version of the main result of [70] achieves the optimal approximation rate with a more efficient parameter usage than the original formulation in [70]. In our version of their result, we achieve the same approximation rate while dropping the dependence of the constant on d from exponential to linear.

6.3 Parameter Quantization

Part of the deep learning literature [78, 34] investigates quantized neural networks; meaning, that their parameters belong to the grid $\{a/2^j : j \in \{0, \dots, b\}, a \in \mathbb{Z}\}$. The reader will notice that most weights and biases in the ReLU MLP Φ constructed in Theorem 4.1.1 belong to the set $\{-n, -n + 1, \dots, n - 1, n\} \cup \{-1/2, 1/2\}$. Since Proposition 4.1.1 guarantees that the parameters of our MLP are remain bounded, then [63, Theorem 3] could be used to fully quantize the neural network with only a minor impact on its expressivity (approximation error). This is because that result provides tight estimates on the effect of perturbing the weights and biases of an MLP with fixed depth and width and bounded parameters. However, such perturbations need not preserve its Lipschitz regularity. We thus mention this as an interesting direction for future research.

6.4 Path Norms and Lipschitz Constants

We mention that $|\Phi|_{\text{par}}$, defined in (3.3.3), is closely related to the so-called “path-norm” seen in the literature, defined by

$$|\Phi|_{\text{path}} \stackrel{\text{def.}}{=} \prod_{l=1}^{\Delta} \|\mathbf{W}^{(l)}\|_{\text{op}}$$

where $\|\cdot\|_{\text{op}}$ denotes the spectral norm of a matrix. The path norm is typically used to quantify the regularity of a parametric representation of a neural network [53], and which is often used as a regularizer during training [54]. The path-norm provides readily computable upper bound on the Lipschitz constant of Φ , which can be easily computed from the parameter space of Φ . This is in stark contrast to the exact Lipschitz constant of ReLU MLPs, which can be difficult to exactly in practice [81]. Here, we mention that $|\Phi|_{\text{par}}$ can easily be estimated from the values in Table 4.1; however, one readily sees that the path norm bounded in this way can severely overestimate the Lipschitz constant of the function *implemented* by the ReLU MLP in Theorem 1.0.1.

Chapter 7

Proofs

7.1 Additional Notation During Proofs

During our proofs, we will also adopt the following notational conventions, in addition to those described in Section 3.1. A single layer of an MLP is considered as a column vector. For each $m \in \mathbb{N}_+$, the median of the real numbers x_1, x_2, \dots, x_m , denoted by $\text{median}(x_1, x_2, \dots, x_m)$, is

$$\text{median}(x_1, x_2, \dots, x_m) \stackrel{\text{def.}}{=} \begin{cases} x_{(m+1)/2} & \text{if } m \text{ is odd} \\ (x_{(m/2)} + x_{(m/2+1)})/2 & \text{if } m \text{ is even} \end{cases} \quad (7.1.1)$$

where $\{x_{(i)}\}_{i=1}^m = \{x_i\}_{i=1}^m$ and $x_{(1)} \leq \dots \leq x_{(m)}$. With a slight abuse of notation, we sometimes abbreviate a column vector in the following way: for example, if $\mathbf{x} = \begin{bmatrix} x_1 & x_2 \end{bmatrix}^\top$ and $\mathbf{y} = \begin{bmatrix} y_1 & y_2 & y_3 \end{bmatrix}^\top$, then

$$\begin{bmatrix} \mathbf{x} & \sigma(\pm a \pm b) & \mathbf{y} \end{bmatrix}^\top \stackrel{\text{def.}}{=} \begin{bmatrix} x_1 & x_2 & \sigma(a+b) & \sigma(a-b) & \sigma(-a+b) & \sigma(-a-b) & y_1 & y_2 & y_3 \end{bmatrix}^\top.$$

For any $p \in [1, \infty)$, the ℓ^p norm of a vector $\mathbf{x} = (x_1, x_2, \dots, x_d) \in \mathbb{R}^d$ is defined as

$$\|\mathbf{x}\|_p \stackrel{\text{def.}}{=} (|x_1|^p + |x_2|^p + \dots + |x_d|^p)^{1/p}$$

and its ℓ^∞ norm is defined as

$$\|\mathbf{x}\|_\infty \stackrel{\text{def.}}{=} \max\{|x_1|, |x_2|, \dots, |x_d|\}$$

The L^p norm of a function $f : \Omega \subseteq \mathbb{R}^d \rightarrow \mathbb{R}$ is defined as

$$\|f\|_{L^p(\Omega)} \stackrel{\text{def.}}{=} \left(\int_{\Omega} |f|^p d\mu \right)^{1/p}$$

where μ is the Lebesgue measure, and its L^∞ norm is defined as

$$\|f\|_{L^\infty(\Omega)} \stackrel{\text{def.}}{=} \inf\{C \geq 0 : |f| \leq C \text{ almost everywhere on } \Omega\}.$$

This paper uses the ℓ^1 norm on the domain; for instance, when defining Lipschitz constants and regularity. We use the (uniform) L^∞ norm when quantifying approximation estimates.

We use \mathcal{O} for the big O notation, which is written as $\tilde{\mathcal{O}}$ when logarithmic factors are ignored. We use Ω for the big Ω notation when it is clear from the context.

For a set S , let $\text{card}(S)$ denote its cardinality. For two sets S_1 and S_2 , let $S_1 \times S_2 \stackrel{\text{def.}}{=} \{(s_1, s_2) : s_1 \in S_1, s_2 \in S_2\}$ be the Cartesian product of S_1 and S_2 , and denote the Cartesian product of S_1 with itself d times for $d \in \mathbb{N}_+$ as S_1^d , i.e. $S_1^d \stackrel{\text{def.}}{=} \prod_{i=1}^d S_1 = \{(s_1, s_2, \dots, s_d) : s_1, s_2, \dots, s_d \in S_1\}$.

For a finite subset $S = \{\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_n\}$ of \mathbb{R}^d , let $\text{conv}(S)$ denote the convex hull

of S , i.e.

$$\text{conv}(S) \stackrel{\text{def.}}{=} \{\lambda_1 \mathbf{v}_1 + \lambda_2 \mathbf{v}_2 + \cdots + \lambda_n \mathbf{v}_n : \lambda_1, \dots, \lambda_n \geq 0, \lambda_1 + \lambda_2 + \cdots + \lambda_n = 1\}.$$

For a set of sets S , let $\cup S \stackrel{\text{def.}}{=} \bigcup_{S' \in S} S'$ denote the union of all the sets in S .

We say that a 1-dimensional function $f : \Omega \subseteq \mathbb{R} \rightarrow \mathbb{R}$ is continuous piecewise linear if there exists a finite collection of intervals I_1, \dots, I_m such that $\Omega \subseteq I_1 \cup \cdots \cup I_m$, and f is linear on I_1, \dots, I_m . We say that $x \in \Omega$ is a *break point* of f if the left and right-hand derivatives of f at x are different. For example, σ is a continuous piecewise linear function on \mathbb{R} , and $x = 0$ is the only break point of σ .

In what follows, given a ReLU MLP Φ with representation (3.3.1); when convenient, we will use $\text{depth}(\Phi)$ to denote its *depth* of a network Φ and $\text{width}(\Phi)$ to denote its width. Emulating [70], we use $\mathcal{NN}(c_1; c_2; \cdots; c_m)$ to denote the class of functions implemented by ReLU MLPs which satisfy conditions c_1, c_2, \dots, c_m . For example, $\mathcal{NN}(\#\text{input} = 2; \text{depth} \leq 3; \text{width} = 4; \#\text{output} = 5)$ denote the class of functions from \mathbb{R}^2 to \mathbb{R}^5 implemented by ReLU MLPs which have depth at most 3 and width equal to 4.

One example of this is the notation $\mathcal{NN}(\mathbf{d})$ introduced in Definition 3.3.1. In this case, the condition is $\text{widthvec}(\Phi) = \mathbf{d}$.

Compositional Notation During the course of our analysis, it will be convenient to describe ReLU MLPs via the role of each of their (sets of) layers. Specifically, the structure of a ReLU MLP Φ is represented in the following way: suppose $\Phi = \mathcal{L}_m \circ (\sigma \circ \mathcal{L}_{m-1}) \circ \cdots \circ (\sigma \circ \mathcal{L}_2) \circ (\sigma \circ \mathcal{L}_1)$ where the \mathcal{L}_i 's are affine transformations,

then we express it as

$$\begin{aligned}
\mathbf{x} &\implies (\sigma \circ \mathcal{L}_1)(\mathbf{x}) \implies (\sigma \circ \mathcal{L}_2) \circ (\sigma \circ \mathcal{L}_1)(\mathbf{x}) \\
&\implies (\sigma \circ \mathcal{L}_{m-1}) \circ \cdots \circ (\sigma \circ \mathcal{L}_2) \circ (\sigma \circ \mathcal{L}_1)(\mathbf{x}) \\
&\implies \mathcal{L}_m \circ (\sigma \circ \mathcal{L}_{m-1}) \circ \cdots \circ (\sigma \circ \mathcal{L}_2) \circ (\sigma \circ \mathcal{L}_1)(\mathbf{x}) \\
&= \Phi(\mathbf{x})
\end{aligned}$$

In other words, if $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_{m-1}$ are the 1, 2, \dots , $(m-1)$ -th hidden layer of Φ and \mathbf{x}_m is the output layer, then the structure of Φ is expressed as

$$\mathbf{x} \implies \mathbf{x}_1 \implies \mathbf{x}_2 \implies \cdots \implies \mathbf{x}_{m-1} \implies \mathbf{x}_m = \Phi(\mathbf{x})$$

If the structure of the network Φ with input \mathbf{x} and output $\Phi(\mathbf{x})$ is clear from the context, then in the following expression, we mean that \mathbf{x} and $\Phi(\mathbf{x})$ are connected via the network Φ :

$$\cdots \implies \begin{bmatrix} \vdots \\ \mathbf{x} \\ \vdots \end{bmatrix} \implies \begin{bmatrix} \vdots \\ \Phi(\mathbf{x}) \\ \vdots \end{bmatrix} \implies \cdots$$

In this case, there may be additional hidden layers between the two layers shown above. Finally, if one of the layers we constructed in fact did not use any activation functions, then it can be integrated with the layer after it, and does not require an extra layer to process. We color them in green to indicate that we can ignore these layers when estimating the size of the network. For example, if $\Phi \stackrel{\text{def.}}{=} \mathcal{L}_4 \circ (\sigma \circ \mathcal{L}_3) \circ (\mathcal{L}_2) \circ (\sigma \circ \mathcal{L}_1)$ is constructed as a network with three hidden layers $\mathbf{x}_1 \stackrel{\text{def.}}{=} \sigma \circ \mathcal{L}_1(\mathbf{x})$, $\mathbf{x}_2 \stackrel{\text{def.}}{=} \mathcal{L}_2(\mathbf{x}_1)$

and $\mathbf{x}_3 \stackrel{\text{def.}}{=} \sigma \circ \mathcal{L}_3(\mathbf{x}_2)$, and \mathbf{x}_2 does not use any activation functions, then Φ can be implemented by a network with only two hidden layers: $\Phi = \mathcal{L}_4 \circ (\sigma \circ (\mathcal{L}_3 \circ \mathcal{L}_2)) \circ (\sigma \circ \mathcal{L}_1)$, whose structure is expressed as

$$\mathbf{x} \implies \mathbf{x}_1 \implies \mathbf{x}_2 \implies \mathbf{x}_3 \implies \Phi(\mathbf{x})$$

7.2 Computational Lemmata

We now compile a sequence of lemmata showing that various key functions can be exactly implemented by “small” MLPs; such as the median function and piecewise linear memorizes. Moreover, when relevant, we show that the MLPs implementing these functions do not alter the regularity of their inputs (explained rigorously below).

7.3 The Median Function

This section focuses on the median function; defined in the following lemma. The implementation of the median function will allow us to construct MLPs with lighted absolute parameter usage (including constants) than in the available literature while matching the available optimal approximation rates.

Lemma 7.3.1. *Let $d \in \mathbb{N}_+$, and $x_1, x_2, \dots, x_{2d+1} \in \mathbb{R}$. Assume $x_1, \dots, x_{d+1} \in [y - \varepsilon, y + \varepsilon]$ for some $y \in \mathbb{R}$ and $\varepsilon > 0$. Then, $\text{median}(x_1, x_2, \dots, x_{2d+1}) \in [y - \varepsilon, y + \varepsilon]$.*

Proof. Let’s relabel $x_1, x_2, \dots, x_{2d+1}$ as $y_1 \leq y_2 \leq \dots \leq y_{2d+1}$. Then, at least one of x_1, x_2, \dots, x_{d+1} , denoted x_n , is in the set $\{y_1, y_2, \dots, y_{d+1}\}$, since the complement of this set only has d elements. Similarly, at least one of x_1, x_2, \dots, x_{d+1} , denoted x_m ,

is in the set $\{y_{d+1}, y_{d+2}, \dots, y_{2d+1}\}$. Therefore,

$$\text{median}(x_1, x_2, \dots, x_{2d+1}) = y_{d+1} \in [x_n, x_m] \subseteq [y - \varepsilon, y + \varepsilon].$$

□

Lemma 7.3.2. *For any $d \in \mathbb{N}_+$, the median function on $2d + 1$ non-negative inputs can be implemented by a ReLU MLP with width $6d + 3$ and depth $11d + 3$.*

Proof. We label the nodes as $n_1, n_2, \dots, n_{2d+1}$, and let

$$D(x) \stackrel{\text{def.}}{=} \sum_{i=1}^{2d+1} |x - n_i|$$

be the total distance from x to all the nodes. Then, D attains its minimum exactly at the median since there are odd number of nodes. We will use the following algorithm to compute the median:

1. Compute $x_0 \stackrel{\text{def.}}{=} \min\{n_1, n_2, \dots, n_{2d+1}\}$
2. Compute $s_0 \stackrel{\text{def.}}{=} \min\{D(n_1), D(n_2), \dots, D(n_{2d+1})\}$
3. For $i = 1, 2, \dots, d$, let $x_i = x_{i-1} + \frac{D(x_{i-1}) - s_0}{2d + 1 - 2i}$
4. Output x_d

We first show the correctness of this algorithm. We relabel the nodes as $m_1 \leq m_2 \leq \dots \leq m_{2d+1}$ and proceed by induction on i . For $i = 0$, $x_0 = m_1 \in [m_1, m_{d+1}]$. For $i = k$, suppose $x_k \in [m_{k+1}, m_{d+1}]$ for some $k \leq d - 1$, we will show that $x_{k+1} \in [m_{k+2}, m_{d+1}]$.

If $x_k \in [m_{k+1}, m_{k+2}]$, then noting that for $r \leq d$,

$$D(m_r) - D(m_{r+1}) = (2(d-r) + 1)(m_{r+1} - m_r)$$

and the same holds when m_{r+1} changes to any $x \in [m_r, m_{r+1}]$, in particular:

$$D(m_{k+1}) - D(x_k) = (2(d-r-1) + 1)(x_k - m_{k+1})$$

so we have

$$\begin{aligned} x_{k+1} &= x_k + \frac{D(x_k) - s_0}{2d-1-2k} = x_k + \frac{D(x_k) - D(m_{d+1})}{2(d-k)-1} \\ &= x_k + \frac{D(m_{k+1}) - (2(d-k-1) + 1)(x_k - m_{k+1}) - D(m_{d+1})}{2(d-k)-1} \\ &= x_k - (x_k - m_{k+1}) + \frac{D(m_{k+1}) - D(m_{d+1})}{2(d-k)-1} \\ &= m_{k+1} + \frac{1}{2(d-k)-1} \sum_{j=1}^{d-k} (2j-1)(m_{d+2-j} - m_{d+1-j}) \end{aligned}$$

then

$$\begin{aligned} x_{k+1} &\leq m_{k+1} + \sum_{j=1}^{d-k} (m_{d+2-j} - m_{d+1-j}) = m_{d+1} \\ x_{k+1} &\geq m_{k+1} + \frac{1}{2(d-k)-1} \sum_{j=d-k}^{d-k} (2j-1)(m_{d+2-j} - m_{d+1-j}) = m_{k+2} \end{aligned}$$

hence $x_{k+1} \in [m_{k+2}, m_{d+1}]$.

Otherwise, $x_k \in [m_{k+2}, m_{d+1}]$, then suppose $x_k \in [m_{k'+1}, m_{k'+2}]$ for some $k < k' \leq d-1$, so we know from above that

$$x_k + \frac{D(x_k) - D(m_{d+1})}{2d-1-2k'} \leq m_{d+1}$$

thus

$$m_{k+2} \leq x_k \leq x_{k+1} = x_k + \frac{D(x_k) - D(m_{d+1})}{2d - 1 - 2k} \leq x_k + \frac{D(x_k) - D(m_{d+1})}{2d - 1 - 2k'} \leq m_{d+1}$$

hence $x_{k+1} \in [m_{k+2}, m_{d+1}]$, so we completed the induction step.

Therefore, $x_i \in [m_{i+1}, m_{d+1}]$ for $i = 1, 2, \dots, d$, so $x_d \in [m_{d+1}, m_{d+1}]$, $x_d = m_{d+1}$, which is the median.

In the following, we will use this algorithm to construct our ReLU MLP. In the rest of this proof, column vectors represent layers of neurons, and functions apply to them entry-wise, whose outputs are still column vectors.

Step 1: Compute x_0 and s_0

It is easy to see that

$$x = \sigma(x) - \sigma(-x), \quad |x| = \sigma(x) + \sigma(-x), \quad \forall x \in \mathbb{R}$$

then

$$\min\{x, y\} = \frac{x + y}{2} - \frac{|x - y|}{2} = \frac{\sigma(x + y)}{2} - \frac{\sigma(-x - y)}{2} - \frac{\sigma(x - y)}{2} - \frac{\sigma(-x + y)}{2}, \quad \forall x, y \in \mathbb{R}$$

so we can use this formula to construct the following building block for computing s_0 : denote the column vector $\begin{bmatrix} n_1 & n_2 & \dots & n_{2d+1} \end{bmatrix}^\top$ as \mathbf{n} , and let \mathbf{n}_k be the column vector obtained from \mathbf{n} by deleting its k th row, then

$$D(n_k) = \begin{bmatrix} 1 & 1 & \dots & 1 \end{bmatrix} \cdot [n_k - \mathbf{n}_k] = \begin{bmatrix} 1 & 1 & \dots & 1 \end{bmatrix} \cdot \begin{bmatrix} \sigma(n_k - \mathbf{n}_k) \\ \sigma(\mathbf{n}_k - n_k) \end{bmatrix}$$

Using these results, we can compute x_0 and d_0 via the following network:

$$\begin{aligned}
\begin{bmatrix} \mathbf{n} \\ \min_{1 \leq i \leq k-1} n_i \\ \min_{1 \leq i \leq k-1} D(n_i) \end{bmatrix} &\implies \begin{bmatrix} \mathbf{n} \\ \min_{1 \leq i \leq k-1} n_i \\ \sigma(n_k - \mathbf{n}_k) \\ \sigma(\mathbf{n}_k - n_k) \\ \min_{1 \leq i \leq k-1} D(n_i) \end{bmatrix} \implies \begin{bmatrix} \mathbf{n} \\ \sigma\left(\pm \min_{1 \leq i \leq k-1} n_i \pm n_k\right) \\ D(n_k) \\ \min_{1 \leq i \leq k-1} D(n_i) \end{bmatrix} \\
&\implies \begin{bmatrix} \mathbf{n} \\ \min_{1 \leq i \leq k} n_i \\ \sigma\left(\pm \min_{1 \leq i \leq k-1} D(n_i) \pm D(n_k)\right) \end{bmatrix} \implies \begin{bmatrix} \mathbf{n} \\ \min_{1 \leq i \leq k} n_i \\ \min_{1 \leq i \leq k} D(n_i) \end{bmatrix}
\end{aligned}$$

where $\sigma(\pm a \pm b)$ abbreviates the four neurons $\sigma(a+b)$, $\sigma(a-b)$, $\sigma(-a+b)$ and $\sigma(-a-b)$.

Note that the terms without the σ are non-negative and thus are unaffected by the activation function σ . Therefore, by connecting $2d$ such networks we can output $x_0 = \min_{1 \leq i \leq k-1} n_i$ and $s_0 = \min_{1 \leq i \leq k-1} D(n_i)$. Its width is dominated by the second layer in the above, which is $(2d+1)+1+2d+2d+1 = 6d+3$, and its depth is $4(2d+1) = 8d+4$ (not counting the first layer, which is the input layer).

Step 2: Compute x_i for $i = 1, 2, \dots, d$

We use the following network:

$$\begin{bmatrix} \mathbf{n} \\ x_{k-1} \\ s_0 \end{bmatrix} \implies \begin{bmatrix} \mathbf{n} \\ \sigma(x_{k-1} - \mathbf{n}) \\ \sigma(\mathbf{n} - x_{k-1}) \\ x_{k-1} \\ s_0 \end{bmatrix} \implies \begin{bmatrix} \mathbf{n} \\ D(x_{k-1}) \\ x_{k-1} \\ s_0 \end{bmatrix} \implies \begin{bmatrix} \mathbf{n} \\ x_{k-1} + \frac{D(x_{k-1}) - s_0}{2d+1-2k} = x_k \\ s_0 \end{bmatrix}$$

Therefore, by connecting d such networks we can output x_d , which is the median as proven above. Its width is dominated by the second layer in the above, which is $(2d + 1) + 2d + 2d + 1 + 1 = 6d + 3$, and its depth is $3d - 1$ (not counting the first and the last layer, since the first layer is shared with the network in the previous step, and the last layer is the output).

Connecting the two network in steps 1 and 2 gives the desired ReLU MLP for computing the median of non-negative inputs. In total, this network has width $6d + 3$ and depth $(8d + 4) + (3d - 1) = 11d + 3$.

□

Lemma 7.3.3 (Median function preserves regularity). *If $f_1, f_2, \dots, f_n \in \text{Hol}(\nu, \alpha, X)$ are Lipschitz functions from a metric space X to \mathbb{R} , then $f \stackrel{\text{def.}}{=} \text{median}(f_1, f_2, \dots, f_n) \in \text{Hol}(\nu, \alpha, X)$.*

Proof. Take any $\mathbf{x}, \mathbf{y} \in X$, then $|f_i(\mathbf{x}) - f_i(\mathbf{y})| \leq \nu \|\mathbf{x} - \mathbf{y}\|_X^\alpha =: C$ for all i . Let $y_i : [0, 1] \rightarrow \mathbb{R}$ be a linear function such that $y_i(0) = f_i(\mathbf{x}), y_i(1) = f_i(\mathbf{y})$ for all i , then $y_i' \leq C$ for all i . For any $t \in [0, 1]$, let $p(t) = (s_1, s_2, \dots, s_n)$, where s_1, s_2, \dots, s_n is a permutation of $1, 2, \dots, n$ such that $f_{s_1}(t) \leq f_{s_2}(t) \leq \dots \leq f_{s_n}(t)$, and $s_i \leq s_j$ whenever $f_{s_i}(t) = f_{s_j}(t)$. Consider any $1 \leq i < j \leq n$, there are two cases:

1. If $f_i(\mathbf{x}) - f_i(\mathbf{y}) = f_j(\mathbf{x}) - f_j(\mathbf{y})$, then $y_i' \equiv y_j'$, $y_i(t) - y_j(t)$ is a constant, so the relative position of i and j in $p(t)$ does not change as t goes from 0 to 1;
2. If $f_i(\mathbf{x}) - f_i(\mathbf{y}) \neq f_j(\mathbf{x}) - f_j(\mathbf{y})$, then as t goes from 0 to 1, $y_i(t)$ and $y_j(t)$ coincide at most once, so the relative position of i and j in $p(t)$ changes at most once.

Therefore, $p(t)$ changes at most $\frac{n(n-1)}{2}$ times as t goes from 0 to 1, and the points where

$p(t)$ changes split $[0, 1]$ into finite intervals. Let $y(t)$ be the median of $y_1(t), \dots, y_n(t)$, then in each of these intervals:

1. If n is odd, then $y(t) = y_i(t)$ for some fixed i , so $|y'(t)| = |y'_i(t)| \leq C$;
2. If n is even, then $y(t) = \frac{y_i(t)}{2} + \frac{y_j(t)}{2}$ for some fixed i and j , so $|y'(t)| \leq \left| \left(\frac{y_i(t)}{2} \right)' \right| + \left| \left(\frac{y_j(t)}{2} \right)' \right| \leq \frac{C}{2} + \frac{C}{2} = C$.

Therefore, $|y'(t)| \leq C$ in each of these intervals, then $|f(\mathbf{x}) - f(\mathbf{y})| = |y(1) - y(0)| \leq C = \nu \|\mathbf{x} - \mathbf{y}\|_X^\alpha$. □

The next lemma shows that one dimensional continuous piecewise linear approximators have the same regularity as the function they approximate. See Lemma 7.5.4 for a generalization to higher dimensions.

Lemma 7.3.4 (Continuous piecewise linear approximators preserve regularity, and L^∞ error estimate). *Let f be a function from $[a, b]$ to \mathbb{R} with modulus of regularity ω , and let Φ be a continuous piecewise linear function on $[a, b]$ that passes through the points $\{(x_i, f(x_i))\}_{i=0}^n$ for $a = x_0 < x_1 < x_2 < \dots < x_n = b$, and is linear on $[x_0, x_1], [x_1, x_2], \dots, [x_{n-1}, x_n]$. Then, we have the error estimate*

$$\|f - \Phi\|_{L^\infty([a,b])} \leq \omega \left(\max_{i=0,1,\dots,n-1} \frac{|x_{i+1} - x_i|}{2} \right)$$

Moreover, ω is a modulus of regularity of Φ .

Proof. We first show that Φ has at least the same regularity as f . Take any $x, y \in [a, b]$, and assume without loss of generality that $x \leq y$. If x, y lie in the same interval in $[x_0, x_1], [x_1, x_2], \dots, [x_{n-1}, x_n]$, then $x, y \in [x_m, x_{m+1}]$ for some index m ,

by concavity of ω ,

$$\begin{aligned}
\omega(y-x) &= \omega\left(\left(1 - \frac{y-x}{x_{m+1}-x_m}\right) \cdot 0 + \frac{y-x}{x_{m+1}-x_m} \cdot (x_{m+1}-x_m)\right) \\
&\geq \left(1 - \frac{y-x}{x_{m+1}-x_m}\right) \omega(0) + \frac{y-x}{x_{m+1}-x_m} \omega(x_{m+1}-x_m) \\
&\geq \frac{y-x}{x_{m+1}-x_m} \omega(x_{m+1}-x_m) \\
&\geq \frac{y-x}{x_{m+1}-x_m} |f(x_{m+1}) - f(x_m)| \\
&= \frac{y-x}{x_{m+1}-x_m} |\Phi(x_{m+1}) - \Phi(x_m)| \\
&= |\Phi(y) - \Phi(x)|
\end{aligned}$$

where the last equality follows from the fact that Φ is linear on $[x_m, x_{m+1}]$. Otherwise, x, y lie in different intervals in $[x_0, x_1], [x_1, x_2], \dots, [x_{n-1}, x_n]$, then $x_p \leq x \leq x_{p+1} \leq x_q \leq y \leq x_{q+1}$ for some indices p, q , and let

$$A \stackrel{\text{def.}}{=} \frac{x-x_p}{x_{p+1}-x_p}, \quad B \stackrel{\text{def.}}{=} \frac{x_{q+1}-y}{x_{q+1}-x_q}$$

so $0 \leq A, B \leq 1$. Without loss of generality, assume that $A \leq B$. Then,

$$\Phi(x) = (1-A)\Phi(x_p) + A\Phi(x_{p+1}), \quad \Phi(y) = B\Phi(x_q) + (1-B)\Phi(x_{q+1})$$

Since

$$\begin{aligned}
y - x &= \left(\frac{x_{q+1} - y}{x_{q+1} - x_q} - \frac{x - x_p}{x_{p+1} - x_p} \right) (x_q - x_p) + \left(1 - \frac{x_{q+1} - y}{x_{q+1} - x_q} \right) (x_{q+1} - x_p) \\
&\quad + \frac{x - x_p}{x_{p+1} - x_p} (x_q - x_{p+1}) \\
&= (B - A)(x_q - x_p) + (1 - B)(x_{q+1} - x_p) + A(x_q - x_{p+1})
\end{aligned}$$

with $B - A \geq 0$, $1 - B \geq 0$, $A \geq 0$ and $(B - A) + (1 - B) + A = 1$, by concavity of ω , we have

$$\begin{aligned}
\omega(y - x) &= \omega((B - A)(x_q - x_p) + (1 - B)(x_{q+1} - x_p) + A(x_q - x_{p+1})) \\
&\geq (B - A)\omega(x_q - x_p) + (1 - B)\omega(x_{q+1} - x_p) + A\omega(x_q - x_{p+1}) \\
&\geq (B - A)|\Phi(x_q) - \Phi(x_p)| + (1 - B)|\Phi(x_{q+1}) - \Phi(x_p)| + A|\Phi(x_q) - \Phi(x_{p+1})| \\
&\geq (B - A)(\Phi(x_q) - \Phi(x_p)) + (1 - B)(\Phi(x_{q+1}) - \Phi(x_p)) + A(\Phi(x_q) - \Phi(x_{p+1})) \\
&= (B\Phi(x_q) + (1 - B)\Phi(x_{q+1})) - ((1 - A)\Phi(x_p) + A\Phi(x_{p+1})) \\
&= \Phi(y) - \Phi(x)
\end{aligned}$$

Similarly, we can show that $\omega(y - x) \geq \Phi(x) - \Phi(y)$, thus

$$\omega(|y - x|) = \omega(y - x) \geq |\Phi(y) - \Phi(x)|$$

Therefore, we always have $|\Phi(y) - \Phi(x)| \leq \omega(|y - x|)$ for any $x, y \in [a, b]$, thus ω is a modulus of regularity of Φ .

Now we prove the upper bound for the L^∞ error. Take any $x \in [a, b]$, then

$x \in [x_k, x_{k+1}]$ for some index k . For convenience, let $L \stackrel{\text{def.}}{=} x_{k+1} - x_k$, then we have

$$\begin{aligned}
|f(x) - \Phi(x)| &= \left| f(x) - \left(\frac{x_{k+1} - x}{L} \Phi(x_k) + \frac{x - x_k}{L} \Phi(x_{k+1}) \right) \right| \\
&= \left| f(x) - \left(\frac{x_{k+1} - x}{L} f(x_k) + \frac{x - x_k}{L} f(x_{k+1}) \right) \right| \\
&= \left| \frac{x_{k+1} - x}{L} (f(x) - f(x_k)) + \frac{x - x_k}{L} (f(x) - f(x_{k+1})) \right| \\
&\leq \frac{x_{k+1} - x}{L} |f(x) - f(x_k)| + \frac{x - x_k}{L} |f(x) - f(x_{k+1})| \\
&\leq \frac{x_{k+1} - x}{L} \omega(x - x_k) + \frac{x - x_k}{L} \omega(x_{k+1} - x) \\
&\leq \omega \left(\frac{x_{k+1} - x}{L} (x - x_k) + \frac{x - x_k}{L} (x_{k+1} - x) \right) \\
&= \omega \left(\frac{(L - (x - x_k))(x - x_k)}{L} + \frac{(x - x_k)(L - (x - x_k))}{L} \right) \\
&\leq \omega \left(\frac{\left(\frac{L^2}{4}\right)}{L} + \frac{\left(\frac{L^2}{4}\right)}{L} \right) \\
&= \omega \left(\frac{L}{2} \right) \\
&\leq \omega \left(\max_{i=0,1,\dots,n-1} \frac{|x_{i+1} - x_i|}{2} \right)
\end{aligned}$$

Since $x \in [a, b]$ was chosen arbitrarily, we conclude that

$$\|f - \Phi\|_{L^\infty([a,b])} \leq \omega \left(\max_{i=0,1,\dots,n-1} \frac{|x_{i+1} - x_i|}{2} \right).$$

This completes our proof. □

7.4 One-Dimensional Memorizers with Optimal Regularity

Proposition 7.4.1 (Efficient Universal Memorization by Two-Hidden-Layer MLPs with Optimal Regularity). *Let $M, N \in \mathbb{N}_+$. For any set of MN samples $(x_i, y_i)_{i=1}^{MN} \subseteq \mathbb{R}^2$ (where $x_1 < x_2 < \dots < x_{MN}$), there exists a ReLU MLP Φ with $\text{widthvec} = [M, 4N - 2]$ that can memorize this sample set; i.e.*

$$\Phi(x_i) = y_i \quad \text{for } i = 1, \dots, MN.$$

Furthermore, Φ is linear on the intervals $[x_i, x_{i+1}]$ for $i = 1, 2, \dots, MN - 1$, and it is constant on each of the segment $(-\infty, x_1]$ and $[x_{MN}, \infty)$. The number of nonzero parameters in Φ is at most $2MN + 2M + 8N - 4$.

Proof. Let $f : \mathbb{R} \rightarrow \mathbb{R}$ be the function to be implemented, i.e. it is the unique function on \mathbb{R} such that:

1. $f(x_i) = y_i$ for $i = 1, 2, \dots, MN$;
2. f is linear on $[x_i, x_{i+1}]$ for $i = 1, 2, \dots, MN - 1$;
3. f is constant on each of $(-\infty, x_1]$ and $[x_{MN}, \infty)$.

Then, it suffices to prove that f can be implemented by a ReLU MLP with $\text{widthvec} = [M, 4N - 2]$.

Let's relabel the samples as $(x_{ij}, y_{ij})_{i,j=1}^{M,N}$ such that x_{ij} increases as (i, j) increases lexicographically, i.e.

$$x_{11} < x_{12} < \dots < x_{1N} < x_{21} < x_{22} < \dots < x_{2N} < \dots < x_{M1} < x_{M2} < \dots < x_{MN}$$

For notational convenience, denote the $M+1$ intervals $(-\infty, x_{11}]$, $[x_{11}, x_{21}]$, $[x_{21}, x_{31}]$, \dots , $[x_{M-1,1}, x_{M,1}]$ and $[x_{M1}, \infty)$ as I_0, I_1, \dots, I_M , respectively. The notations x_{ij} and $x_{i,j}$ have the same meaning.

The idea of the proof is matching the jumps in the derivative of f step by step: for $r = 1, 2$ and $j = 2, 3, \dots, N$, let $S_r \stackrel{\text{def}}{=} \{I_{r+2k} : k \in \mathbb{N}, 1 \leq r + 2k \leq M\} = \{I_r, I_{r+2}, I_{r+4}, \dots\}$, and construct the continuous piecewise linear function $g_j^{(+r)}$ in two steps:

1. **Matching upward derivative jumps at the j th sample point in intervals in S_r :** let

$$X_j^{(+r)} \stackrel{\text{def}}{=} \{x_{r+2k,j} : k \in \mathbb{N}, f'_+(x_{r+2k,j}) - f'_-(x_{r+2k,j}) > 0\}$$

be the set of the j th sample points in intervals in S_r where f has upward derivative jumps. Suppose the elements in $X_j^{(+r)}$ are $x_{i_1,j} < x_{i_2,j} < \dots < x_{i_m,j}$, then for each $1 \leq m' \leq m$, let $g_j^{(+r)}$ be linear on $I_{i_{m'}}$ with slope $(-1)^{m'}(f'_+(x_{i_{m'},j}) - f'_-(x_{i_{m'},j}))$, and $g_j^{(+r)}(x_{i_{m'},j}) = 0$;

2. **Smooth extension to the rest of \mathbb{R} :** we have already defined $g_j^{(+r)}$ on some distinct closed intervals in Step 1 (unless $X_j^{(+r)}$ is empty, in which case simply let $g_j^{(+r)}$ be the zero function on \mathbb{R}), which separate \mathbb{R} into some segments. For each of these segments, if it is infinite then we have already defined $g_j^{(+r)}$ at its only endpoint, then let $g_j^{(+r)}$ be constant on this segment matching its value at this endpoint; if it is finite then we have already defined $g_j^{(+r)}$ at both of its endpoints, then let $g_j^{(+r)}$ be linear on this segment such that it matches both of its values at the two endpoints.

We define $X_j^{(+r)}$ and $g_j^{(-r)}$ similar to the above, except that the condition $f'_+(x_{r+2k,j}) - f'_-(x_{r+2k,j}) > 0$ in Step 1 is changed to $f'_+(x_{r+2k,j}) - f'_-(x_{r+2k,j}) < 0$. It is easy to see that all break points of $g_j^{(\pm r)}$ belong to $\{x_{11}, x_{21}, \dots, x_{M1}\}$, and it does not change sign on the intervals $\{I_0, I_1, \dots, I_M\} \setminus S_r$. Thus, all break points of $\sigma(g_j^{(\pm r)})$ belong to $\{x_{11}, x_{21}, \dots, x_{M1}\} \cup \{x_{r,j}, x_{r+2,j}, x_{r+4,j}, \dots\}$. Moreover, for any $x \in \{x_{r,j}, x_{r+2,j}, x_{r+4,j}, \dots\}$, if f' has an upward jump at x then $\sigma(g_j^{(+r)})$ also has an upward jump in the derivative with the same magnitude, if f' has a downward jump at x then $\sigma(g_j^{(-r)})$ has an upward jump in the derivative with the same magnitude, thus $\sigma(g_j^{(+r)}) - \sigma(g_j^{(-r)})$ has the same jump in derivatives at x as f . Therefore, the function

$$\sum_{\substack{1 \leq r \leq 2 \\ 2 \leq j \leq N}} \left(\sigma(g_j^{(+r)}) - \sigma(g_j^{(-r)}) \right)$$

and f has the same jumps in derivatives at $\{x_{ij} : 1 \leq i \leq M, 2 \leq j \leq N\}$, thus all break points of the function

$$g_0 \stackrel{\text{def.}}{=} f - \sum_{\substack{1 \leq r \leq 2 \\ 2 \leq j \leq N}} \sigma(g_j^{(+r)}) + \sum_{\substack{1 \leq r \leq 2 \\ 2 \leq j \leq N}} \sigma(g_j^{(-r)})$$

belong to $\{x_{11}, x_{21}, \dots, x_{M1}\}$.

We have

$$\begin{aligned} f &= \sum_{\substack{1 \leq r \leq 2 \\ 2 \leq j \leq N}} \sigma(g_j^{(+r)}) - \sum_{\substack{1 \leq r \leq 2 \\ 2 \leq j \leq N}} \sigma(g_j^{(-r)}) + g_0 \\ &= \sum_{\substack{1 \leq r \leq 2 \\ 2 \leq j \leq N}} \sigma(g_j^{(+r)}) - \sum_{\substack{1 \leq r \leq 2 \\ 2 \leq j \leq N}} \sigma(g_j^{(-r)}) + \sigma(g_0) - \sigma(-g_0) \end{aligned} \quad (7.4.1)$$

We now show that f can be represented by a ReLU MLP with $\text{widthvec} = [M, 4N - 2]$. For notational convenience, we will let $X_i = x_{i1}$ for $i = 1, \dots, M$ in the following.

We know from above that for any function $g \in \{g_j^{(\pm r)} : 1 \leq r \leq 2, 2 \leq j \leq N\} \cup \{\pm g_0\}$, the set of break points of g is contained in $\{X_1, X_2, \dots, X_M\}$, and g is constant on $(-\infty, X_1]$. Then, by comparing the jumps in derivatives as what we did above, we can see that the function

$$g(x) - g(X_1) - \sum_{i=1}^M (g'_+(X_i) - g'_-(X_i)) \sigma(x - X_i)$$

has no break points (thus is linear on all of \mathbb{R}) and is 0 for all $x \leq X_1$, hence is 0 on all of \mathbb{R} . That is,

$$\begin{aligned} g(x) &= g(X_1) + \sum_{i=1}^M (g'_+(X_i) - g'_-(X_i)) \sigma(x - X_i) \\ &= \begin{pmatrix} g'_+(X_1) - g'_-(X_1) \\ g'_+(X_2) - g'_-(X_2) \\ \vdots \\ g'_+(X_M) - g'_-(X_M) \end{pmatrix} \cdot \sigma \left(x \begin{pmatrix} 1 \\ 1 \\ \vdots \\ 1 \end{pmatrix} - \begin{pmatrix} X_1 \\ X_2 \\ \vdots \\ X_M \end{pmatrix} \right) + g(X_1), \quad \forall x \in \mathbb{R} \end{aligned} \quad (7.4.2)$$

Therefore, g can be represented by a neural network with one hidden layer consisting of M neurons, with a single input neuron and a single output neuron. Moreover, the weight matrix between the input and the hidden layer is always $\begin{bmatrix} 1 & 1 & \dots & 1 \end{bmatrix}^\top$, and the bias vector for the hidden layer is always $\begin{bmatrix} -X_1 & -X_2 & \dots & -X_M \end{bmatrix}^\top$, which is the same for all such g . Therefore, all these g 's can share their input layers and hidden layers, so we can stack their output neurons ($4N - 2$ in total, one for each g)

in the second hidden layer of the final network such that their outputs are unaffected by each other, then finally f can be implemented by adding one more neuron in the output layer based on Equation 7.4.1.

Finally, we tally the nonzero parameters in Φ . For the $4N - 2$ neurons in the second hidden layer, $4N - 4$ of them are $g_j^{(\pm r)}$ for $r = 1, 2$ and $j = 2, 3, \dots, N$. From Step 1, we can see that $g_j^{(+r)}$ has at most $2 \text{card}(X_j^{(+r)})$ break points (since Step 2 did not introduce additional ones), all of which belong to $\{X_1, X_2, \dots, X_M\}$. Thus, at most $2 \text{card}(X_j^{(+r)})$ weights between the first hidden layer and the neuron that outputs $g_j^{(+r)}$ are nonzero. Therefore, the number of nonzero parameters between the two hidden layers is at most

$$\begin{aligned}
& \sum_{\substack{1 \leq r \leq 2 \\ 2 \leq j \leq N}} 2 \text{card}(X_j^{(+r)}) + \sum_{\substack{1 \leq r \leq 2 \\ 2 \leq j \leq N}} 2 \text{card}(X_j^{(-r)}) + 2M \\
&= 2 \text{card}(\{x_{ij} : 1 \leq i \leq M, 2 \leq j \leq N, f'_+(x_{r+2k,j}) - f'_-(x_{r+2k,j}) \neq 0\}) + 2M \\
&\leq 2M(N - 1) + 2M \\
&= 2MN
\end{aligned}$$

From Equation 7.4.1, the bias for the output is 0. Therefore, Φ has at most $2MN + M + 4N - 2$ nonzero weights and at most $M + 4N - 2$ nonzero biases. Altogether, Φ has at most $2MN + 2M + 8N - 4$ nonzero parameters. \square

In Proposition 7.4.1, if there are K sample points and take

$$M = 2\lceil\sqrt{K}\rceil, \quad N = \left\lceil \frac{\lceil\sqrt{K}\rceil}{2} \right\rceil$$

then, upon noticing that

$$MN \geq K, \quad 4N - 2 = 4 \left\lceil \frac{[\sqrt{K}]}{2} \right\rceil - 2 \leq 4 \frac{[\sqrt{K}] + 1}{2} - 2 = 2[\sqrt{K}]$$

we directly deduce the following corollary.

Corollary 7.4.1. *Let $K \in \mathbb{N}_+$. For any set of K samples $(x_i, y_i)_{i=1}^K \subseteq \mathbb{R}^2$ (where $x_1 < x_2 < \dots < x_K$), there exists a ReLU MLP Φ with width at most $2[\sqrt{K}]$, depth 2, and with at most $2K + 8[\sqrt{K}]$ nonzero parameters that can memorize this sample set; i.e.*

$$\Phi(x_i) = y_i \quad \text{for } i = 1, \dots, K.$$

Furthermore, Φ is linear on the intervals $[x_i, x_{i+1}]$ for $i = 1, 2, \dots, K - 1$, and it is constant on each of the segment $(-\infty, x_1]$ and $[x_K, \infty)$.

Theorem 7.4.1 (Efficient Universal Memorization by Deep MLPs with Optimal Regularity).

Let $K, L \in \mathbb{N}_+$ with $L \geq 3$. Given any set of K samples $(x_i, y_i)_{i=1}^K \subseteq \mathbb{R}^2$ (with $x_1 < x_2 < \dots < x_K$) and any $11 < n_1, n_2, \dots, n_L \in \mathbb{N}_+$ satisfying the constraint:

$$K \leq \sum_{b=1}^{L-1} \left((n_b - 11) \left\lceil \frac{n_{b+1} - 9}{4} \right\rceil - 2 \right)$$

there exists a ReLU MLP Φ with width $\text{vec} [n_1, n_2, \dots, n_L, 8]$ that can memorize this sample set; i.e.

$$\Phi(x_i) = y_i \quad \text{for } i = 1, \dots, K.$$

Furthermore, Φ is linear on the intervals $[x_i, x_{i+1}]$ for $i = 1, 2, \dots, K - 1$, and it is constant on each of the segment $(-\infty, x_1]$ and $[x_K, \infty)$. The number of nonzero

parameters in Φ is at most

$$2K + 23 \sum_{b=1}^L n_b - 121L.$$

We note that, the factors 23 and 121 above can be further improved. For simplicity, here we only prove this weaker bound below, which already shows the significance of this theorem: as the minimum width n_{\min} (i.e. minimum number of neurons in a single hidden layer) goes to infinity, the efficiency of parameter usage

$$\text{Efficiency} \stackrel{\text{def.}}{=} \frac{\text{degree of freedom of sample set}}{\text{number of nonzero parameters}} \cdot 100\%$$

goes to 100%. This is because, assuming

$$K = \sum_{b=1}^{L-1} \left((n_b - 11) \left\lfloor \frac{n_{b+1} - 9}{4} \right\rfloor - 2 \right)$$

then we have

$$\begin{aligned} \text{Efficiency} &\geq \frac{K \dim(\mathbb{R}^2)}{2K + 23 \sum_{b=1}^L n_b - 121L} \cdot 100\% \\ &\geq \frac{2K}{2K + 23 \sum_{b=1}^L n_b} \cdot 100\% \\ &\geq \frac{2K}{2K + 23(K + 2(L - 1)) / \left\lfloor \frac{n_{\min} - 9}{4} \right\rfloor} \cdot 100\% \\ &\geq \frac{2}{2 + 23(1 + 2(L - 1)) / \left\lfloor \frac{n_{\min} - 9}{4} \right\rfloor} \cdot 100\% \\ &\rightarrow 100\% \end{aligned}$$

when the depth L is kept fixed (or $L = o(n_{\min})$) as $n_{\min} \rightarrow \infty$. The same is true for

Proposition 7.4.1.

Proof of Theorem 7.4.1. Let $f : \mathbb{R} \rightarrow \mathbb{R}$ be the function to be implemented, i.e. it is the unique function on \mathbb{R} such that:

1. $f(x_i) = y_i$ for $i = 1, 2, \dots, K$;
2. f is linear on $[x_i, x_{i+1}]$ for $i = 1, 2, \dots, K - 1$;
3. f is constant on each of $(-\infty, x_1]$ and $[x_K, \infty)$.

Then, it suffices to prove that f can be implemented by a ReLU MLP with $\text{widthvec} = [n_1, n_2, \dots, n_L, 8]$.

Clearly x_1, \dots, x_K contain all possible break points of f , so we may assume that all of them are break points of f , i.e. the left and right hand derivatives of f at any of the x_i 's are different.

The idea of the proof is the same as in the proof of Proposition 7.4.1: instead of implementing the target function f directly, we match the jumps in its derivative, which is mostly done by adding up $L - 1$ functions f_1, \dots, f_{L-1} (to be defined later), each matching those jumps in different intervals.

We divide the break points $\{x_i\}_{i=1}^K$ of f into $L - 1$ batches: for a batch number $b = 1, 2, \dots, L - 1$, let $k_b \stackrel{\text{def.}}{=} n_b - 11$ (and $k_L \stackrel{\text{def.}}{=} n_L - 11$), and denote the b th batch of break points to be

$$x_1^{(b)}, x_2^{(b)}, \dots, x_{K_b}^{(b)}, \quad K_b \stackrel{\text{def.}}{=} (n_b - 11) \left\lfloor \frac{n_{b+1} - 9}{4} \right\rfloor - 2$$

and let $X_B, Y_B, Z_B \in \mathbb{R}$ for $B = 1, 2, \dots, L$ be “separation points” between different

batches of break points such that

$$\begin{aligned}
 X_1 < Y_1 < Z_1 < x_1^{(1)} < \cdots < x_{K_1}^{(1)} < X_2 < Y_2 < Z_2 \\
 < x_2^{(2)} < \cdots < x_{K_2}^{(2)} < X_3 < Y_3 < Z_3 \\
 < \cdots \\
 < x_{L-1}^{(2)} < \cdots < x_{K_{L-1}}^{(2)} < X_L < Y_L < Z_L
 \end{aligned}$$

Let $f_b : \mathbb{R} \rightarrow \mathbb{R}$ be a function that satisfies the following conditions:

1. **Continuous piecewise linear on \mathbb{R} with few “pieces”:** all break points of f_b belong to

$$\{x_1^{(b)}, x_2^{(b)}, \dots, x_{K_b}^{(b)}, Z_b, X_{b+1}\};$$

2. **Matching derivative jumps at the b th batch of break points:**

$$(f_b)'_+(x) - (f_b)'_-(x) = f'_+(x) - f'_-(x), \quad \text{for } x = x_1^{(b)}, x_2^{(b)}, \dots, x_{K_b}^{(b)}$$

3. **Vanishing at other batches of break points:**

$$f_b(x) = 0, \quad \forall x \in (-\infty, Z_b] \cup [X_{b+1}, \infty)$$

Then, adding up all the f_b 's would recover most informations of f .

We first need to show the existence of these f_b 's: let t_b be the linear function on \mathbb{R} such that $t_b(Z_b) = f(Z_b)$ and $t_b(X_{b+1}) = f(X_{b+1})$, then it is easy to verify that the

function f_b defined by

$$f_b(x) \stackrel{\text{def.}}{=} \begin{cases} f(x) - t_b(x), & \text{if } x \in [Z_b, X_{b+1}] \\ 0, & \text{otherwise} \end{cases}$$

satisfies Conditions 1 to 3.

Now, for each f_b , we build a network Φ_b implementing f_b using the construction in the proof of Proposition 7.4.1: for $b = 1, 2, \dots, L-1$, take $M = k_b$, $N = \lfloor (k_{b+1} + 2)/4 \rfloor$ and sample set $\{(x_i^{(b)}, f_b(x_i^{(b)}))\}_{i=1}^{MN-2}$ (namely the b th batch of samples) along with the two endpoints $(Z_b, 0)$ and $(X_{b+1}, 0)$ in Proposition 7.4.1, and then denote the structure of the resulting network as

$$x \implies \mathbf{h}_1^{(b)}(x) \implies \mathbf{h}_2^{(b)}(x) \implies \Phi_b(x)$$

where $\mathbf{h}_1^{(b)}(x), \mathbf{h}_2^{(b)}(x)$ are column vectors representing the two hidden layers, and $\mathbf{h}_1^{(b)}, \mathbf{h}_2^{(b)}$ are considered as functions of $x \in \mathbb{R}$. By Proposition 7.4.1, the length of $\mathbf{h}_1^{(b)}(x)$ is $M = k_b$, and the length of $\mathbf{h}_2^{(b)}(x)$ is $4N - 2 \leq k_{b+1}$, which we treat as exactly k_{b+1} for simplicity. Moreover, we know that

$$\mathbf{h}_1^{(b)}(x) = \begin{bmatrix} \sigma(x - B_1^{(b)}) \\ \vdots \\ \sigma(x - B_{k_b}^{(b)}) \end{bmatrix}, \quad \forall x \in \mathbb{R}$$

Fix a batch number $1 \leq b \leq L - 2$. Define $I_b \stackrel{\text{def.}}{=} [Z_b, X_{b+1}]$, as Φ_b only matches the jumps in f' inside this interval. Next, we modify Φ_b so that its second hidden layer $\mathbf{h}_2^{(b)}$ can also serve as the first hidden layer in Φ_{b+1} on the next interval I_{b+1} .

There are k_{b+1} functions in $\mathbf{h}_2^{(b)}$, we denote them as $\sigma(g_1^{(b)}), \sigma(g_2^{(b)}), \dots, \sigma(g_{k_{b+1}}^{(b)})$, i.e.

$$\mathbf{h}_2^{(b)} = \begin{bmatrix} \sigma(g_1^{(b)}) \\ \vdots \\ \sigma(g_{k_{b+1}}^{(b)}) \end{bmatrix}$$

Take any $g_j^{(b)}$ with $1 \leq j \leq k_{b+1}$, which is an affine transformation of $\mathbf{h}_1^{(b)}$, and it is linear on $[X_{b+1}, \infty)$. Let

$$\tilde{g}_j^{(b)}(x) \stackrel{\text{def.}}{=} g_j^{(b)}(x) + p_j^{(b)}\sigma(x - X_{b+1}) + q_j^{(b)}\sigma(x - Y_{b+1}), \quad \forall x \in \mathbb{R}$$

for suitable choices of $p_j^{(b)}, q_j^{(b)} \in \mathbb{R}$ such that:

1. $\tilde{g}_j^{(b)}(Y_{b+1}) \neq 0$;
2. $\tilde{g}_j^{(b)}(B_j^{(b+1)}) = 0$;
3. $\tilde{g}_j^{(b)}(X_{b+1})\tilde{g}_j^{(b)}(Y_{b+1}) \geq 0$, so that $\sigma(\tilde{g}_j^{(b)})$ doesn't have any additional break points between X_{b+1} and $B_j^{(b+1)}$.

Since $B_j^{(b+1)} \geq Z_{b+1} > Y_{b+1}$, there exist such choices for $p_j^{(b)}, q_j^{(b)} \in \mathbb{R}$. Also, we note that the $\tilde{g}_j^{(b)}$ is a modification of the $g_j^{(b)}$ outside I_b , i.e.

$$\tilde{g}_j^{(b)}(x) = g_j^{(b)}(x), \quad \forall x \in I_b$$

thus we define the following modification layer:

$$\bar{\mathbf{h}}_2^{(b)} \stackrel{\text{def.}}{=} \begin{bmatrix} \sigma(\tilde{g}_1^{(b)}) \\ \vdots \\ \sigma(\tilde{g}_{k_{b+1}}^{(b)}) \end{bmatrix} \quad (7.4.3)$$

and we have that

$$\bar{\mathbf{h}}_2^{(b)}(x) = \mathbf{h}_2^{(b)}(x), \quad \forall x \in I_b$$

Let $\mathcal{L}_{\mathbf{h}_2^{(b)}, \Phi_b}$ be the affine transformation that maps the second hidden layer $\mathbf{h}_2^{(b)}$ of Φ_b to its output. Then, the function

$$\tilde{\Phi}_b \stackrel{\text{def.}}{=} \mathcal{L}_{\mathbf{h}_2^{(b)}, \Phi_b}(\bar{\mathbf{h}}_2^{(b)})$$

“mimics” the output of Φ_b on I_b , in the sense that

$$\tilde{\Phi}_b(x) = \mathcal{L}_{\mathbf{h}_2^{(b)}, \Phi_b}(\bar{\mathbf{h}}_2^{(b)}(x)) = \mathcal{L}_{\mathbf{h}_2^{(b)}, \Phi_b}(\mathbf{h}_2^{(b)}(x)) = \Phi_b(x), \quad \forall x \in I_b$$

Here, we emphasize that $\tilde{\Phi}_b$ is treated as a function, not a network.

Now, $\tilde{g}_j^{(b)}$ is linear on $[Z_{b+1}, \infty)$ and crosses the x -axis at $B_j^{(b+1)}$ with slope $s_j^{(b)} \neq 0$, so we have

$$\sigma(\tilde{g}_j^{(b)}(x)) = \begin{cases} s_j^{(b)} \sigma(x - B_j^{(b+1)}), & \text{if } s_j^{(b)} > 0 \\ s_j^{(b)} x - s_j^{(b)} \sigma(x - B_j^{(b+1)}), & \text{if } s_j^{(b)} < 0 \end{cases}, \quad \forall x \in [Z_{b+1}, \infty)$$

Correspondingly, we define affine transformations

$$\mathcal{L}_j^{(b)}(x, y) = \begin{cases} y/s_j^{(b)}, & \text{if } s_j^{(b)} > 0 \\ -y/s_j^{(b)} + x, & \text{if } s_j^{(b)} < 0 \end{cases}, \quad \forall x, y \in \mathbb{R}$$

which recover $\sigma(x - B_j)$ on $[Z_{b+1}, \infty)$ upon composing with $\sigma(\tilde{g}_j^{(b)})$, i.e.

$$\mathcal{L}_j^{(b)}\left(x, \sigma(\tilde{g}_j^{(b)}(x))\right) = \sigma(x - B_j^{(b+1)}), \quad \forall x \in [Z_{b+1}, \infty)$$

We will see that, in the final construction of Φ , we will keep a copy of the input x in every hidden layer. Now, we have the following “mimic layer” having exactly the same behavior as $\mathbf{h}_1^{(b+1)}$ on $I_{b+1} \subseteq [Z_{b+1}, \infty)$:

$$\tilde{\mathbf{h}}_1^{(b+1)}(x) \stackrel{\text{def.}}{=} \begin{bmatrix} \mathcal{L}_1^{(b)}\left(x, \sigma(\tilde{g}_1^{(b)}(x))\right) \\ \vdots \\ \mathcal{L}_{k_{b+1}}^{(b)}\left(x, \sigma(\tilde{g}_{k_{b+1}}^{(b)}(x))\right) \end{bmatrix}, \quad \forall x \in \mathbb{R} \quad (7.4.4)$$

In other words,

$$\tilde{\mathbf{h}}_1^{(b+1)}(x) = \mathbf{h}_1^{(b+1)}(x), \quad \forall x \in I_{b+1}$$

Notice that, from Equations 7.4.4 and 7.4.3, we can see that the function $\tilde{\mathbf{h}}_1^{(b+1)}(x)$ is simply an affine transformation of $\left[\bar{\mathbf{h}}_2^{(b)}(x) \quad x\right]^\top$, and thus does not need an extra layer to process.

So far, we have tweaked the hidden layers $\mathbf{h}_2^{(b)}$ and $\mathbf{h}_1^{(b+1)}$ into $\bar{\mathbf{h}}_2^{(b)}$ and $\tilde{\mathbf{h}}_1^{(b+1)}$ respectively, so that $\tilde{\mathbf{h}}_1^{(b+1)}$ does not need an extra layer to process. At the same time, we have a collection of the resulting tweaked versions of $\Phi_1, \dots, \Phi_{L-1}$, namely

$\tilde{\Phi}_1, \dots, \tilde{\Phi}_{L-1}$, with $\tilde{\Phi}_b$ having the same output as Φ_b on I_b , for all batch number b . However, these $\tilde{\Phi}_b$'s cannot be used directly, as they have other derivative jumps outside I_b , which we now eliminate by post-processing the data $\tilde{\Phi}_b(x)$ together with x . For each $b = 1, 2, \dots, L-1$, let $w_b > 0$ be large enough so that the function

$$q_b(x) \stackrel{\text{def.}}{=} \tilde{\Phi}_b(x) + w_b x, \quad \forall x \in \mathbb{R}$$

is strictly increasing on \mathbb{R} . Since $\tilde{\Phi}_b$ is continuous piecewise linear with finite number of “pieces”, its Lipschitz constant is finite, then w_b only need to be larger than this constant, so such choice of w_b exists. Define the 1D map projecting \mathbb{R} to $q_b(I_b) = [q_b(Z_b), q_b(X_{b+1})]$:

$$p_b(y) \stackrel{\text{def.}}{=} \sigma(y - q_b(Z_b)) - \sigma(y - q_b(X_{b+1})) + q_b(Z_b) = \begin{cases} q_b(Z_b), & \text{if } y < q_b(Z_b) \\ y, & \text{if } q_b(Z_b) \leq y \leq q_b(X_{b+1}) \\ q_b(X_{b+1}), & \text{if } y > q_b(X_{b+1}) \end{cases}, \quad \forall y \in \mathbb{R}$$

and then let

$$Q_b(x) \stackrel{\text{def.}}{=} p_b(q_b(x)) - w_b x, \quad \forall x \in \mathbb{R}$$

We claim that $Q_b(x) = \tilde{\Phi}_b(x)$ for all $x \in I_b$, and Q_b is linear on each of the two segments of $\mathbb{R} \setminus I_b$. Indeed, if $x \in I_b = [Z_b, X_{b+1}]$, then $q_b(Z_b) \leq q_b(x) \leq q_b(X_{b+1})$ as q_b is strictly increasing, thus $Q_b(x) = p_b(q_b(x)) - w_b x = q_b(x) - w_b x = \tilde{\Phi}_b(x)$. If $x < Z_b$, then $q_b(x) < q_b(Z_b)$, thus $Q_b(x) = p_b(q_b(x)) - w_b x = q_b(Z_b) - w_b x$, which is linear. If $x > X_{b+1}$, then $q_b(x) > q_b(X_{b+1})$, thus $Q_b(x) = p_b(q_b(x)) - w_b x = q_b(X_{b+1}) - w_b x$, which is linear.

Therefore, $Q_b = \tilde{\Phi}_b = \Phi_b$ on I_b , which has break points $\{x_i^{(b)}\}_{i=1}^{MN-2}$ along with Z_b, X_{b+1} . Moreover, Q_b has the same derivative jumps at $\{x_i^{(b)}\}_{i=1}^{MN-2}$ as the target function f . Thus, the function $f - \sum_{b=1}^{L-1} Q_b$ only has derivative jumps at X_1, \dots, X_L and Z_1, \dots, Z_L . Then, by choosing appropriate real coefficients u_b 's and v_b 's, the function

$$f(x) - \sum_{b=1}^{L-1} Q_b(x) - \sum_{b=1}^L u_b \sigma(x - X_b) - \sum_{b=1}^L v_b \sigma(x - Z_b)$$

will have no break points on \mathbb{R} , which makes it a linear function, say $ux + v$ for some $u, v \in \mathbb{R}$. Then,

$$f(x) = \sum_{b=1}^{L-1} Q_b(x) + \sum_{b=1}^L u_b \sigma(x - X_b) + \sum_{b=1}^L v_b \sigma(x - Z_b) + ux + v, \quad \forall x \in \mathbb{R} \quad (7.4.5)$$

We can now implement f based on Equation 7.4.5. We will keep a copy of the original input x via the identity $\sigma(x) - \sigma(-x) = x$ ($\forall x \in \mathbb{R}$) in every hidden layer of the final network Φ , but for notational convenience we will omit the two neurons $\sigma(\pm x)$ in the following representations. Also, we define

$$UV_{A,B}(x) \stackrel{\text{def.}}{=} \sum_{b=A}^B u_b \sigma(x - X_b) + \sum_{b=A}^B v_b \sigma(x - Z_b), \quad \forall x \in \mathbb{R}$$

to simplify notations later.

Now, we start by constructing the first few layers of Φ computing $\sum_{b=1}^{L-1} Q_b(x)$ as

follows:

$$\begin{array}{ccccccc}
 x \implies & \begin{bmatrix} \mathbf{h}_1^{(1)}(x) \\ \sigma(x - X_2) \\ \sigma(x - Y_2) \\ \sigma(x - Z_2) \end{bmatrix} & \implies & \begin{bmatrix} g_1^{(1)}(x) \\ \vdots \\ g_{k_2}^{(1)}(x) \\ \sigma(x - X_2) \\ \sigma(x - Y_2) \\ \sigma(x - Z_2) \end{bmatrix} & \implies & \begin{bmatrix} \tilde{g}_1^{(1)}(x) \\ \vdots \\ \tilde{g}_{k_2}^{(1)}(x) \\ UV_{2,2}(x) \end{bmatrix} & \implies & \begin{bmatrix} \bar{\mathbf{h}}_2^{(1)}(x) \\ \sigma(x - X_3) \\ \sigma(x - Y_3) \\ \sigma(x - Z_3) \\ \sigma(\pm UV_{2,2}(x)) \end{bmatrix} \\
 & & & & & & & \\
 \implies & \begin{bmatrix} \tilde{\mathbf{h}}_1^{(2)} \\ \tilde{\Phi}_1(x) \\ \sigma(x - X_3) \\ \sigma(x - Y_3) \\ \sigma(x - Z_3) \\ UV_{2,2}(x) \end{bmatrix} & \implies & \begin{bmatrix} g_1^{(2)}(x) \\ \vdots \\ g_{k_3}^{(2)}(x) \\ q_1(x) \\ \sigma(x - X_3) \\ \sigma(x - Y_3) \\ \sigma(x - Z_3) \\ UV_{2,2}(x) \end{bmatrix} & \implies & \begin{bmatrix} \tilde{g}_1^{(2)}(x) \\ \vdots \\ \tilde{g}_{k_3}^{(2)}(x) \\ q_1(x) - q_1(Z_1) \\ q_1(x) - q_1(X_2) \\ UV_{2,3}(x) \end{bmatrix} & \implies & \begin{bmatrix} \bar{\mathbf{h}}_2^{(2)}(x) \\ \sigma(x - X_4) \\ \sigma(x - Y_4) \\ \sigma(x - Z_4) \\ \sigma(q_1(x) - q_1(Z_1)) \\ \sigma(q_1(x) - q_1(X_2)) \\ \sigma(\pm UV_{2,3}(x)) \end{bmatrix}
 \end{array}$$

$$\begin{aligned}
&\Rightarrow \begin{bmatrix} \tilde{\mathbf{h}}_1^{(3)} \\ \tilde{\Phi}_2(x) \\ \sigma(x - X_4) \\ \sigma(x - Y_4) \\ \sigma(x - Z_4) \\ Q_1(x) \\ UV_{2,3}(x) \end{bmatrix} \Rightarrow \begin{bmatrix} g_1^{(3)}(x) \\ \vdots \\ g_{k_4}^{(3)}(x) \\ q_2(x) \\ \sigma(x - X_4) \\ \sigma(x - Y_4) \\ \sigma(x - Z_4) \\ Q_1(x) \\ UV_{2,3}(x) \end{bmatrix} \Rightarrow \begin{bmatrix} \tilde{g}_1^{(3)}(x) \\ \vdots \\ \tilde{g}_{k_4}^{(3)}(x) \\ q_2(x) - q_2(Z_2) \\ q_2(x) - q_2(X_3) \\ Q_1(x) \\ UV_{2,4}(x) \end{bmatrix} \Rightarrow \begin{bmatrix} \bar{\mathbf{h}}_2^{(3)}(x) \\ \sigma(x - X_5) \\ \sigma(x - Y_5) \\ \sigma(x - Z_5) \\ \sigma(q_2(x) - q_2(Z_2)) \\ \sigma(q_2(x) - q_2(X_3)) \\ \sigma(\pm Q_1(x)) \\ \sigma(\pm UV_{2,4}(x)) \end{bmatrix} \\
&\hspace{20em} (7.4.6)
\end{aligned}$$

As always, auxiliary layers are colored in green to indicate that we can ignore these layers when estimating the size of the final network, as they are affine transformations that do not use any activation functions, thus each of them can be integrated with the layer after it and does not require an extra layer to process. Removing the auxiliary layers and then extending this pattern, we obtain the final construction of the network Φ :

$$x \Rightarrow \begin{bmatrix} \mathbf{h}_1^{(1)}(x) \\ \sigma(x - X_2) \\ \sigma(x - Y_2) \\ \sigma(x - Z_2) \end{bmatrix} \Rightarrow \begin{bmatrix} \bar{\mathbf{h}}_2^{(1)}(x) \\ \sigma(x - X_3) \\ \sigma(x - Y_3) \\ \sigma(x - Z_3) \\ \sigma(\pm UV_{2,2}(x)) \end{bmatrix} \Rightarrow \begin{bmatrix} \bar{\mathbf{h}}_2^{(2)}(x) \\ \sigma(x - X_4) \\ \sigma(x - Y_4) \\ \sigma(x - Z_4) \\ \sigma(q_1(x) - q_1(Z_1)) \\ \sigma(q_1(x) - q_1(X_2)) \\ \sigma(\pm UV_{2,3}(x)) \end{bmatrix}$$

$$\begin{aligned}
&\Rightarrow \begin{bmatrix} \bar{\mathbf{h}}_2^{(3)}(x) \\ \sigma(x - X_5) \\ \sigma(x - Y_5) \\ \sigma(x - Z_5) \\ \sigma(q_2(x) - q_2(Z_2)) \\ \sigma(q_2(x) - q_2(X_3)) \\ \sigma(\pm Q_1(x)) \\ \sigma(\pm UV_{2,4}(x)) \end{bmatrix} \Rightarrow \begin{bmatrix} \bar{\mathbf{h}}_2^{(4)}(x) \\ \sigma(x - X_6) \\ \sigma(x - Y_6) \\ \sigma(x - Z_6) \\ \sigma(q_3(x) - q_3(Z_3)) \\ \sigma(q_3(x) - q_3(X_4)) \\ \sigma(\pm(Q_1(x) + Q_2(x))) \\ \sigma(\pm UV_{2,5}(x)) \end{bmatrix} \Rightarrow \dots \\
&\Rightarrow \begin{bmatrix} \bar{\mathbf{h}}_2^{(L-2)}(x) \\ \sigma(x - X_L) \\ \sigma(x - Y_L) \\ \sigma(x - Z_L) \\ \sigma(q_{L-3}(x) - q_{L-3}(Z_{L-3})) \\ \sigma(q_{L-3}(x) - q_{L-3}(X_{L-2})) \\ \sigma(\pm \sum_{b=1}^{L-4} Q_b(x)) \\ \sigma(\pm UV_{2,L-1}(x)) \end{bmatrix} \Rightarrow \begin{bmatrix} \bar{\mathbf{h}}_2^{(L-1)}(x) \\ \sigma(x - X_1) \\ \sigma(x - Y_1) \\ \sigma(x - Z_1) \\ \sigma(q_{L-2}(x) - q_{L-2}(Z_{L-2})) \\ \sigma(q_{L-2}(x) - q_{L-2}(X_{L-1})) \\ \sigma(\pm \sum_{b=1}^{L-3} Q_b(x)) \\ \sigma(\pm UV_{2,L}(x)) \end{bmatrix} \\
&\Rightarrow \begin{bmatrix} \sigma(q_{L-1}(x) - q_{L-1}(Z_{L-1})) \\ \sigma(q_{L-1}(x) - q_{L-1}(X_L)) \\ \sigma(\pm \sum_{b=1}^{L-2} Q_b(x)) \\ \sigma(\pm UV_{1,L}(x)) \end{bmatrix} \\
&\Rightarrow \begin{bmatrix} \sum_{b=1}^{L-1} Q_b(x) \\ UV_{1,L}(x) \\ ux + v \end{bmatrix} \Rightarrow \sum_{b=1}^{L-1} Q_b(x) + UV_{1,L}(x) + ux + v = f(x)
\end{aligned}$$

Together with the omitted two neurons $\sigma(\pm x)$, each hidden layer in Φ has at most 11 additional neurons other than those of the \mathbf{h} 's. Therefore, it has widthvec at most $[k_1 + 11, k_2 + 11, \dots, k_L + 11, 8] = [n_1, n_2, \dots, n_L, 8]$.

Finally, we tally the nonzero parameters in Φ . From the detailed structure of Φ shown in Formula 7.4.6, we can see that most connections (weights) between pairs of hidden layers are the same as their counterparts in the network Φ_b : the neurons in $\mathbf{h}_1^{(1)}$ and $\bar{\mathbf{h}}_2^{(1)}$ are connected in the same way as the neurons in $\mathbf{h}_1^{(1)}$ and $\mathbf{h}_2^{(1)}$ are connected; for $b = 1, 2, \dots, L - 1$, the neurons in $\bar{\mathbf{h}}_2^{(b)}$ and $\bar{\mathbf{h}}_2^{(b+1)}$ are connected in the same way as the neurons in $\mathbf{h}_1^{(b+1)}$ and $\mathbf{h}_2^{(b+1)}$ are connected. From the proof of Proposition 7.4.1, we know that there are at most $2k_b \lfloor (k_{b+1} + 2)/4 \rfloor$ nonzero weights between $\mathbf{h}_1^{(b)}$ and $\mathbf{h}_2^{(b)}$. Therefore, assuming the worst case that each hidden layer has 11 additional neurons other than those in the \mathbf{h} 's, we conclude that the number of nonzero weights in Φ is at most

$$\begin{aligned}
W &\stackrel{\text{def.}}{=} \sum_{b=1}^{L-1} \left(2k_b \left\lfloor \frac{k_{b+1} + 2}{4} \right\rfloor + 11k_b + 11k_{b+1} + 11 \cdot 11 \right) + k_1 + 8(k_L + 11) + 8 \\
&\leq 2K + 22 \sum_{b=1}^L k_b + 121(L - 1) + 88 + 8 \\
&\leq 2K + 22 \sum_{b=1}^L (n_b - 11) + 121L - 25 \\
&= 2K + 22 \sum_{b=1}^L n_b - 121L - 25
\end{aligned}$$

and the number of biases in Φ is at most

$$B \stackrel{\text{def.}}{=} \sum_{b=1}^L n_b + 8 + 1$$

Altogether, the number of nonzero parameters in Φ is at most

$$W + B \leq 2K + 23 \sum_{b=1}^L n_b - 121L.$$

This completes our proof. □

The following example elucidates Proposition 7.4.1, and frames it in a way which is comparable to the memorization guarantees for MLPs derived in [80, 78]. Though those networks utilize fewer neurons when performing their memorization, they exhibit a much larger Lipschitz constant due to their highly irregular structure; since they are only designed to optimize bit-extraction of [73, 67](see [5] for estimates on the maximal bit-extraction “capacity” of MLPs).

Example 7.4.1 (A 1D MLP with width W and depth D can memorize $\mathcal{O}(W^2D)$ data points with optimal Lipschitz constant). Given any set of K samples $(x_i, y_i)_{i=1}^K \subseteq \mathbb{R}^2$, where $x_1 < x_2 < \dots < x_K$, and

$$K = \left((W - 11) \left\lfloor \frac{W - 9}{4} \right\rfloor - 2 \right) (D - 2)$$

there exists an MLP with width at most W and depth at most D that can memorize this sample set in such a way that it is linear on $[x_i, x_{i+1}]$ for $i = 1, 2, \dots, K - 1$, and is constant on each of $(-\infty, x_1]$ and $[x_K, \infty)$. The number of nonzero parameters in Φ is at most

$$2K + (23W - 121)D.$$

7.5 ReLU Approximations on Polytopes via the Kuhn Triangulation

In this section, we generalize the concept of continuous piecewise linear approximators and the methods for constructing them to higher dimensions. As we will see, compared to the ones in Theorem 6.2.1, these high dimensional continuous piecewise linear approximators require more width, but they can achieve about the same level of global error with much higher regularity (lower Lipschitz constant).

Definition 7.5.1 (Simplex). A d -dimensional simplex (or d -simplex) S is the convex hull of $d + 1$ points $\mathbf{v}_1, \dots, \mathbf{v}_{d+1}$ in \mathbb{R}^d that are affinely independent, i.e. the vectors

$$\mathbf{v}_{d+1} - \mathbf{v}_1, \mathbf{v}_{d+1} - \mathbf{v}_2, \dots, \mathbf{v}_{d+1} - \mathbf{v}_d$$

are linearly independent in \mathbb{R}^d , and

$$S \stackrel{\text{def.}}{=} \text{conv}(\{\mathbf{v}_1, \dots, \mathbf{v}_{d+1}\}) = \left\{ \sum_{i=1}^{d+1} \lambda_i \mathbf{v}_i : \lambda_1, \dots, \lambda_{d+1} \geq 0, \sum_{i=1}^{d+1} \lambda_i = 1 \right\}.$$

Furthermore, $\mathbf{v}_1, \dots, \mathbf{v}_{d+1}$ are called the vertices of the simplex S .

Definition 7.5.2 (Triangulation of polytopes). A triangulation of a polytope P in \mathbb{R}^d is a finite collection of d -simplices $\{S_i\}_{i=1}^n$ such that

1. $S_1 \cup S_2 \cup \dots \cup S_n = P$;
2. The set of all their vertices is the vertex set of P ¹;

¹Some literatures allow the existence of additional vertices inside P , but here we require that there be no additional vertices other than those of P to simplify later arguments.

3. Their interiors are pairwise disjoint.

Triangulations are generally not unique. For example, Figure 7.1 shows two different triangulations of a hexagon.




Figure 7.1: Two different triangulations of a hexagon.

The next lemma shows how to triangulate all unit cubes $\{[i, i + 1] : i \in \mathbb{Z}\}^d$ of the entire \mathbb{R}^d simultaneously using hyperplanes. The triangulation of each of these unit cubes will be a Kuhn triangulation.

Lemma 7.5.1. *The set of hyperplanes in \mathbb{R}^d*

$$H \stackrel{\text{def}}{=} \{x_i \pm x_j = 2k : 1 \leq i < j \leq d, k \in \mathbb{Z}\} \cup \{x_i = k : 1 \leq i \leq d, k \in \mathbb{Z}\}$$

triangulates each of the unit cubes in $\{[i, i + 1] : i \in \mathbb{Z}\}^d$.

Proof. Let $\cup H \stackrel{\text{def}}{=} \bigcup_{h \in H} h$ be the union of the hyperplanes (when considered as subsets of \mathbb{R}^d) in H . By the definition of H , $\cup H$ has the following two properties:

1. $\cup H$ has period 2 along each axis:

$$\mathbf{z} \in \cup H \iff (\mathbf{z} + 2\mathbf{e}_i) \in \cup H, \quad i = 1, 2, \dots, d$$

2. $\cup H$ is symmetric about each axis:

$$\mathbf{z} \in \cup H \iff (-\mathbf{z}) \in \cup H, \quad i = 1, 2, \dots, d$$

By property 1, we only need to show that $\cup H$ triangulates $[-1, 1]^d$. By property 2, we only need to show that $\cup H$ triangulates $[0, 1]^d$.

By Definition 6.1.1, the $d!$ simplices

$$\text{conv} \left(\left\{ \sum_{i=1}^k \mathbf{e}_{\tau(i)} : 0 \leq k \leq d \right\} \right), \quad \tau \in S_d \quad (7.5.1)$$

form the Kuhn triangulation of $[0, 1]^d$. Fix any permutation $\tau \in S_d$. The simplex

$$\text{conv} \left(\left\{ \sum_{i=1}^k \mathbf{e}_{\tau(i)} : 0 \leq k \leq d \right\} \right)$$

has $d + 1$ faces

$$\text{conv} \left(\left\{ \sum_{i=1}^k \mathbf{e}_{\tau(i)} : 0 \leq k \leq d, k \neq j \right\} \right), \quad j = 0, 1, \dots, d.$$

For $j = 0$, the corresponding face is contained in the hyperplane $x_{\tau(1)} = 1$ in H . For $j = d$, the corresponding face is contained in the hyperplane $x_{\tau(d)} = 0$ in H . For $1 \leq j \leq d - 1$, the corresponding face is contained in the hyperplane $x_{\tau(j)} - x_{\tau(j+1)} = 0$ in H . Therefore, all the faces of all $d!$ simplices are contained in some hyperplanes in H . It remains to show that the intersection point of any d hyperplanes in H (if any) is a lattice point.

Let h_1, h_2, \dots, h_d be any d hyperplanes in H that intersect at a single point

(x'_1, \dots, x'_d) , i.e.

$$h_1 \cap h_2 \cap \dots \cap h_d = \{(x'_1, \dots, x'_d)\}$$

We will proceed by induction to show that this is a lattice point, and the statement is clearly true for $d = 1, 2$. Thus, we assume that $d \geq 3$.

If one of these hyperplanes has the form $x_{i_0} = k$, then the remaining $d - 1$ hyperplanes, when treated as equations and substituted $x_{i_0} = k$, are $d - 1$ equations on $d - 1$ variables $\{x_1, \dots, x_d\} \setminus \{x_{i_0}\}$, which have the form $x_i \pm x_j = 2k$ (if $i \neq i_0 \neq j$) or $x_{i_0} \pm x_j = 2k$ and have a unique solution $\{x'_1, \dots, x'_d\} \setminus \{x'_{i_0}\}$ by our assumption above. By induction, $\{x'_1, \dots, x'_d\} \setminus \{x'_{i_0}\} \subset \mathbb{Z}$, and we also have $x'_{i_0} = k \in \mathbb{Z}$, so $(x'_1, \dots, x'_d) \in \mathbb{Z}^d$.

If any two of these hyperplanes have the form $x_{i_0} + x_{j_0} = 2k$ and $x_{i_0} - x_{j_0} = 2l$, then $x'_{i_0} = k + l \in \mathbb{Z}, x'_{j_0} = k - l \in \mathbb{Z}$. Similar to above, the remaining $d - 2$ hyperplanes, when treated as equations and substituted $x_{i_0} = k + l$ and $x_{j_0} = k - l$, are $d - 2$ equations on $d - 2$ variables $\{x_1, \dots, x_d\} \setminus \{x_{i_0}, x_{j_0}\}$, which have the form $x_i \pm x_j = 2k$ (if $i, j \notin \{i_0, j_0\}$) or $x_{i_0} \pm x_j = 2k$ or $x_{j_0} \pm x_j = 2k$ and have a unique solution $\{x'_1, \dots, x'_d\} \setminus \{x'_{i_0}, x'_{j_0}\}$ by our assumption above. By induction, $\{x'_1, \dots, x'_d\} \setminus \{x'_{i_0}, x'_{j_0}\} \subset \mathbb{Z}$, and we also have $x'_{i_0} = k + l \in \mathbb{Z}, x'_{j_0} = k - l \in \mathbb{Z}$, so $(x'_1, \dots, x'_d) \in \mathbb{Z}^d$.

Finally, if h_1, h_2, \dots, h_d all have the form $x_i \pm x_j = 2k$ and they don't contain two hyperplanes of the form $x_i + x_j = 2k$ and $x_i - x_j = 2k$ for any $1 \leq i < j \leq d$, then consider the (undirected simple) graph $G = (\mathcal{V}, \mathcal{E})$, where $\mathcal{V} = \{x_1, x_2, \dots, x_d\}$ is the vertex set, and edge set

$$\mathcal{E} = \{\{x_i, x_j\} : \text{there is a hyperplane of the form } x_i \pm x_j = 2k\}$$

Since G has d vertices and d edges, it has a cycle. Let $\{x_{i_1}, x_{i_2}\}, \{x_{i_2}, x_{i_3}\}, \dots, \{x_{i_{m-1}}, x_{i_m}\}, \{x_{i_m}, x_{i_1}\}$ be a shortest cycle, so that x_{i_1}, \dots, x_{i_m} are distinct. Then, h_1, h_2, \dots, h_d contain the following $m - 1$ hyperplanes:

$$x_{i_s} \pm x_{i_{s+1}} = 2k_s, \quad s = 1, 2, \dots, m - 1$$

By adding or subtracting these equations, we obtain an equation of the form

$$x_{i_1} \pm x_{i_m} = 2(k_1 \pm k_2 \pm \dots \pm k_m)$$

Without loss of generality, assume that the above equation is of the form $x_{i_1} + x_{i_m} = 2k_0$ for some $k_0 \in \mathbb{Z}$. If the hyperplane h_{i_m, i_1} corresponding to the edge $\{x_{i_m}, x_{i_1}\}$ has the form $x_{i_1} + x_{i_m} = 2l_0$, then $h_1 \cap h_2 \cap \dots \cap h_d$ is either empty (if $l_0 \neq k_0$) or infinite (if $l_0 = k_0$ and the intersection is not empty, since the equation for h_{i_m, i_1} would be redundant and there are essentially only $d - 1$ equations for d variables), which contradicts our assumption that h_1, h_2, \dots, h_d intersect at a single point. Therefore, the hyperplane corresponding to the edge $\{x_{i_m}, x_{i_1}\}$ has the form $x_{i_1} - x_{i_m} = 2l_0$, so $x'_{i_1} = k_0 + l_0 \in \mathbb{Z}$, $x'_{i_m} = k_0 - l_0 \in \mathbb{Z}$. Proceeding by induction as above, we conclude that $(x'_1, \dots, x'_d) \in \mathbb{Z}^d$ is a lattice point.

□

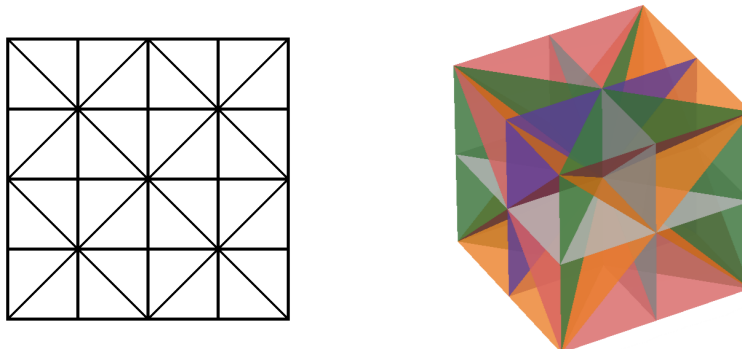


Figure 7.2: Finer Kuhn Triangulation (Definition 6.1.1) of a 2-dimensional square (left), and a 3-dimensional cube (right), than in Figure 6.1.

Definition 7.5.3 (Continuous piecewise linear functions on polytopes). Let f be a function from some polytope $P \subseteq \mathbb{R}^d$ to \mathbb{R} . We say that f is continuous piecewise linear on P if f is continuous on P , and there exists a triangulation of P such that f is affine on each of its simplices.

We first show how to construct continuous piecewise linear approximators on $[0, n]^d$ for some positive integer n .

Lemma 7.5.2. *Let $\Omega \subseteq \mathbb{R}^d$ be a connected open set, and $f : \Omega \rightarrow \mathbb{R}$ be a function such that for every point in Ω , it has a neighborhood on which f is affine. Then, f is affine on Ω .*

Proof. Fix any $\mathbf{x} \in \Omega$. Then, there exists a neighborhood of \mathbf{x} on which $f \equiv g$, for some affine function g defined on all of \mathbb{R}^d .

Assume for contradiction that there is some $\mathbf{y} \in \Omega$ such that $f(\mathbf{y}) \neq g(\mathbf{y})$, so $\mathbf{y} \neq \mathbf{x}$. Since Ω is connected, there exists a path γ that connects \mathbf{x} and \mathbf{y} inside Ω , i.e. $\gamma : [0, 1] \rightarrow \Omega$ is continuous and $\gamma(0) = \mathbf{x}$, $\gamma(1) = \mathbf{y}$. Let

$$T \stackrel{\text{def.}}{=} \{t \in [0, 1] : f(\gamma(t)) \neq g(\gamma(t))\}, \quad t_0 \stackrel{\text{def.}}{=} \inf T$$

Since $f(\mathbf{y}) \neq g(\mathbf{y})$, $f(\gamma(1)) \neq g(\gamma(1))$, so $t_0 \leq 1$ exists. Since $\gamma(0) = \mathbf{x}$ and $f \equiv g$ on some neighborhood of \mathbf{x} , $t_0 > 0$. By assumption, there exists an open neighborhood $U \subseteq \Omega$ of $\gamma(t_0)$ on which f is affine. Since $\gamma(t_0)$ either belongs to $\gamma(T)$ or is a limit point of $\gamma(T)$, and $\gamma(T) \subseteq \mathbb{R}^d \setminus U$ which is closed, $\gamma(t_0) \in \mathbb{R}^d \setminus U$, so $\gamma(t_0) \neq \mathbf{x}$. As $t \rightarrow t_0$, $\gamma(t) \rightarrow \gamma(t_0)$, and since $\gamma(t_0) \neq \mathbf{x} = \gamma(0)$, there are infinite points in $\gamma([0, t_0])$ that are inside U . Since both f and g are affine on U and $f \equiv g$ on $\gamma([0, t_0]) \cap U$ which is infinite, $f \equiv g$ on U .

However, there exists a sequence $(t_i)_{i=1}^\infty$ in T that converges to t_0 , so $\gamma(t_i) \rightarrow \gamma(t_0) \in U$ as $i \rightarrow \infty$, then there exists some i_0 large enough so that $\gamma(t_{i_0}) \in U$, but $f(\gamma(t_{i_0})) \neq g(\gamma(t_{i_0}))$, which contradicts our conclusion above. Therefore, $f \equiv g$ on Ω , thus f is affine on Ω . \square

The following lemma provides an explicit formula for the “hat functions” for the Kuhn triangulation. Figure 7.3 shows two hat functions for the triangulation in Lemma 7.5.1.




Figure 7.3: The supports of two hat functions in 3D view (left) and 2d view (right) with respect to the triangulation in Lemma 7.5.1. Each of them is affine on any of the smallest triangles, so any linear combinations of them is continuous piecewise linear with respect to the triangulation in Lemma 7.5.1.

Lemma 7.5.3 (Constructing continuous piecewise linear approximators in \mathbb{R}^d). *Let*

f be a function from $[0, n]^d$ to \mathbb{R} for some $n \in \mathbb{N}_+$, then the function Φ defined by

$$\Phi(\mathbf{x}) \stackrel{\text{def.}}{=} \sum_{\mathbf{y} \in \{0, 1, \dots, n\}^d} f(\mathbf{y}) \sigma(1 - \max\{|x_i - y_i| : y_i \text{ is even}\} - \max\{|x_i - y_i| : y_i \text{ is odd}\})$$

for $\mathbf{x} = (x_1, \dots, x_d) \in [0, n]^d$ is continuous piecewise linear on each of the n^d unit cubes $\{[0, 1], [1, 2], \dots, [n-1, n]\}^d$ of $[0, n]^d$ with respect to the triangulation in Lemma 7.5.1, and

$$\Phi(\mathbf{y}) = f(\mathbf{y}), \quad \forall \mathbf{y} \in \{0, 1, \dots, n\}^d$$

Proof. For convenience, we first define

$$T_{\mathbf{y}}(\mathbf{x}) \stackrel{\text{def.}}{=} \sigma(1 - \max\{|x_i - y_i| : y_i \text{ is even}\} - \max\{|x_i - y_i| : y_i \text{ is odd}\}), \quad \forall \mathbf{x} \in \mathbb{R}^d$$

to be the *hat function*² centred at \mathbf{y} . Then,

$$\Phi(\mathbf{x}) = \sum_{\mathbf{y} \in \{0, 1, \dots, n\}^d} f(\mathbf{y}) T_{\mathbf{y}}(\mathbf{x})$$

For any $\mathbf{y}, \mathbf{z} \in \{0, 1, \dots, n\}^d$ with $\mathbf{y} \neq \mathbf{z}$, there is some i such that $|y_i - z_i| \geq 1$, so $T_{\mathbf{z}}(\mathbf{y}) = 0$, thus

$$\Phi(\mathbf{y}) = f(\mathbf{y}) T_{\mathbf{y}}(\mathbf{y}) + \sum_{\substack{\mathbf{z} \in \{0, 1, \dots, n\}^d \\ \mathbf{z} \neq \mathbf{y}}} f(\mathbf{z}) T_{\mathbf{z}}(\mathbf{y}) = f(\mathbf{y}) T_{\mathbf{y}}(\mathbf{y}) = f(\mathbf{y})$$

Now we show that g is continuous piecewise linear with respect to the triangulations in Lemma 7.5.1. We only need to show that the hat functions $\{T_{\mathbf{y}}\}_{\mathbf{y} \in \{0, 1, \dots, n\}^d}$

²This is the same as the hat function (or nodal basis function) discussed in [24], and here we have provided its formula explicitly when the underlying triangulation is the one in Lemma 7.5.1.

are continuous piecewise linear. Since the hat functions are obviously continuous, by Lemma 7.5.2, we only need to show that each hat function is locally affine at every point outside the union of those hyperplanes.

Take any $\mathbf{x} \in \mathbb{R}^d$ that is outside the union of these hyperplanes, and fix an arbitrary $\mathbf{y} \in \{0, 1, \dots, n\}^d$. We need to find a neighborhood of \mathbf{x} on which $T_{\mathbf{y}}$ is affine. For notational convenience, let

$$S_{\mathbf{y}}(\mathbf{z}) \stackrel{\text{def.}}{=} 1 - \max\{|z_i - y_i| : y_i \text{ is even}\} - \max\{|z_i - y_i| : y_i \text{ is odd}\}, \quad \forall \mathbf{z} \in \mathbb{R}^d$$

If $S_{\mathbf{y}}(\mathbf{x}) = 0$, then there are some i, j with $i \neq j$ such that

$$1 - |x_i - y_i| - |x_j - y_j| = 0, \quad y_i \text{ is even, } y_j \text{ is odd}$$

removing the absolute values results in an equation of the form

$$\pm x_i \pm x_j = \pm y_i \pm y_j + 1$$

where $\pm y_i \pm y_j + 1$ is even, contradicting the assumption that \mathbf{x} is outside those hyperplanes. If $S_{\mathbf{y}}(\mathbf{x}) < 0$, then $T_{\mathbf{y}} \equiv \sigma(S_{\mathbf{y}}) \equiv 0$ on some neighborhood of \mathbf{x} , and we are done. Therefore, in the following we will assume that $S_{\mathbf{y}}(\mathbf{x}) > 0$, so there is some neighborhood U_1 of \mathbf{x} on which $S_{\mathbf{y}} > 0$.

If $|x_i - y_i| = 0$ for some i , then $x_i = y_i \in \mathbb{Z}$, contradicting the assumption that \mathbf{x} is outside those hyperplanes. Then, $|x_i - y_i| \neq 0$ for all i .

If there are some i, j with $i \neq j$ such that y_i and y_j are both even, and $|x_i - y_i| = |x_j - y_j|$, then removing the absolute values results in an equation of the form

$\pm x_i \pm x_j = \pm y_i \pm y_j$ where $\pm y_i \pm y_j$ is even, contradicting the assumption that \mathbf{x} is outside those hyperplanes. Then, there is a unique i_0 such that

$$\max \{|x_i - y_i| : y_i \text{ is even}\} = |x_{i_0} - y_{i_0}|, \quad y_{i_0} \text{ is even}$$

thus $|x_i - y_i| < |x_{i_0} - y_{i_0}|$ for all $i \neq i_0$ such that y_i is even. Therefore, there exists a neighborhood U_2 of \mathbf{x} such that for any $\mathbf{z} \in U_2$, $z_{i_0} - y_{i_0}$ does not change sign and $\max \{|z_i - y_i| : y_i \text{ is even}\} = |z_{i_0} - y_{i_0}|$, so the function $\max \{|z_i - y_i| : y_i \text{ is even}\}$ of \mathbf{z} is affine on U_2 . Similarly, there exists a neighborhood U_3 of \mathbf{x} such that the function $\max \{|z_i - y_i| : y_i \text{ is odd}\}$ of \mathbf{z} is affine on U_3 .

Now take $U = U_1 \cap U_2 \cap U_3$, so U is also a neighborhood of \mathbf{x} . Moreover, $S_{\mathbf{y}} > 0$ on U , so $T_{\mathbf{y}} \equiv S_{\mathbf{y}}$ on U . Since $U \subseteq U_2$ and $U \subseteq U_3$, the functions $\max \{|z_i - y_i| : y_i \text{ is even}\}$ and $\max \{|z_i - y_i| : y_i \text{ is odd}\}$ of \mathbf{z} are affine on U , so $S_{\mathbf{y}}$ is affine on U , hence $T_{\mathbf{y}}$ is also affine on U . \square

Lemma 7.5.4 (Continuous piecewise linear approximators with respect to the Kuhn triangulation preserve regularity under ℓ^1 norm). *Let f be a function from $[0, n]^d$ to \mathbb{R} for some $n \in \mathbb{N}_+$. Let $\omega : [0, nd] \rightarrow \mathbb{R}$ be a modulus of regularity of f under ℓ^1 norm of \mathbb{R}^d . Let $\Phi : [0, n]^d \rightarrow \mathbb{R}$ be continuous piecewise linear on each of the n^d unit cubes $\{[0, 1], [1, 2], \dots, [n-1, n]\}^d$ with respect to the triangulation in Lemma 7.5.1, and*

$$\Phi(\mathbf{y}) = f(\mathbf{y}), \quad \forall \mathbf{y} \in \{0, 1, \dots, n\}^d$$

Then, Φ satisfies the approximation guarantee

$$\|f - \Phi\|_{L^\infty([0, n]^d)} \leq \omega\left(\frac{d}{2}\right)$$

Furthermore, ω is a modulus of regularity of Φ under ℓ^1 norm of \mathbb{R}^d .

During the proof of Lemma 7.5.4 we will make the following notational and terminological conventions. For any vector $\mathbf{z} \in \mathbb{R}^d$, we denote its i -th coordinate by $(\mathbf{z})_i$. The notation \mathbf{z}_i without brackets will have other contextual meanings specified when appropriate. We say that a function $g : \mathbb{R}^d \rightarrow \mathbb{R}$ has ω -regularity at a pair of points $\mathbf{u}, \mathbf{v} \in \mathbb{R}^d$ if $|g(\mathbf{u}) - g(\mathbf{v})| \leq \omega(\|\mathbf{u} - \mathbf{v}\|_1)$.

Proof. We turn the problem into a path-finding problem: observe that, for any pair of $\mathbf{x}, \mathbf{y} \in [0, n]^d$, if we can find a pair of corresponding finite linear paths $\gamma_{\mathbf{x}}, \gamma_{\mathbf{y}} : [0, T] \rightarrow [0, n]^d$ and a time $t_0 \in [0, T]$ such that the following conditions hold:

1. $\gamma_{\mathbf{x}}(0), \gamma_{\mathbf{x}}(1)$ are contained in the same simplex;
2. $\gamma_{\mathbf{y}}(0), \gamma_{\mathbf{y}}(1)$ are contained in the same simplex;
3. $\mathbf{x} = \gamma_{\mathbf{x}}(t_0)$;
4. $\mathbf{y} = \gamma_{\mathbf{y}}(t_0)$;
5. As t changes linearly from 0 to 1, $\|\gamma_{\mathbf{x}}(t) - \gamma_{\mathbf{y}}(t)\|_1$ changes linearly from $\|\gamma_{\mathbf{x}}(0) - \gamma_{\mathbf{y}}(0)\|_1$ to $\|\gamma_{\mathbf{x}}(1) - \gamma_{\mathbf{y}}(1)\|_1$;
6. Φ has ω -regularity at $\gamma_{\mathbf{x}}(0), \gamma_{\mathbf{y}}(0)$ and $\gamma_{\mathbf{x}}(1), \gamma_{\mathbf{y}}(1)$.

then we can conclude that Φ also has ω -regularity at \mathbf{x}, \mathbf{y} ; here, by $\gamma_{\mathbf{x}}, \gamma_{\mathbf{y}}$ being linear we mean that each of their coordinates are linear functions of t .

This is because, by conditions 1 and 2, Φ is linear on the line segment connecting $\gamma_{\mathbf{x}}(0), \gamma_{\mathbf{x}}(1)$ as well as the line segment connecting $\gamma_{\mathbf{y}}(0), \gamma_{\mathbf{y}}(1)$, so

$$\Phi(\mathbf{x}) = \frac{t_0}{T}\Phi(\gamma_{\mathbf{x}}(1)) + \frac{T - t_0}{T}\Phi(\gamma_{\mathbf{x}}(0)), \quad \Phi(\mathbf{y}) = \frac{t_0}{T}\Phi(\gamma_{\mathbf{y}}(1)) + \frac{T - t_0}{T}\Phi(\gamma_{\mathbf{y}}(0))$$

by conditions 3,4 and 5, we have that

$$\|\mathbf{x} - \mathbf{y}\|_1 = \frac{t_0}{T} \|\gamma_{\mathbf{x}}(1) - \gamma_{\mathbf{y}}(1)\|_1 + \frac{T - t_0}{T} \|\gamma_{\mathbf{x}}(0) - \gamma_{\mathbf{y}}(0)\|_1$$

thus

$$\begin{aligned} |\Phi(\mathbf{x}) - \Phi(\mathbf{y})| &= \left| \frac{t_0}{T} (\Phi(\gamma_{\mathbf{x}}(1)) - \Phi(\gamma_{\mathbf{y}}(1))) + \frac{T - t_0}{T} (\Phi(\gamma_{\mathbf{x}}(0)) - \Phi(\gamma_{\mathbf{y}}(0))) \right| \\ &\leq \frac{t_0}{T} |\Phi(\gamma_{\mathbf{x}}(1)) - \Phi(\gamma_{\mathbf{y}}(1))| + \frac{T - t_0}{T} |\Phi(\gamma_{\mathbf{x}}(0)) - \Phi(\gamma_{\mathbf{y}}(0))| \\ &\leq \frac{t_0}{T} \omega(\|\gamma_{\mathbf{x}}(1) - \gamma_{\mathbf{y}}(1)\|_1) + \frac{T - t_0}{T} \omega(\|\gamma_{\mathbf{x}}(0) - \gamma_{\mathbf{y}}(0)\|_1) \\ &\hspace{20em} \text{(by condition 6)} \\ &\leq \omega \left(\frac{t_0}{T} \|\gamma_{\mathbf{x}}(1) - \gamma_{\mathbf{y}}(1)\|_1 + \frac{T - t_0}{T} \|\gamma_{\mathbf{x}}(0) - \gamma_{\mathbf{y}}(0)\|_1 \right) \\ &\hspace{20em} \text{(since } \omega \text{ is concave)} \\ &= \omega(\|\mathbf{x} - \mathbf{y}\|_1) \end{aligned}$$

In other words, for any pair of $\mathbf{x}, \mathbf{y} \in [0, n]^d$, if we can find two finite linear paths such that:

- 1' Each path is entirely contained in some simplex;
- 2' ω -regularity of Φ is assured at the pair of starting points and the pair of ending points;
- 3' As two points travel along the two paths, the ℓ^1 distance between them changes linearly;
- 4' As two points travel along the two paths, they reach \mathbf{x}, \mathbf{y} respectively at the same time,

then the concavity of ω would guarantee the ω -regularity of Φ at any intermediate pair of points, including \mathbf{x}, \mathbf{y} .

Now, we fix any pair of $\mathbf{x}, \mathbf{y} \in [0, n]^d$. For any $\mathbf{z} \in [0, n]^d$, suppose it is contained in some simplex $P_{\mathbf{z}}$ with vertices $\mathbf{u}_1, \dots, \mathbf{u}_{d+1} \in \{0, 1, \dots, n\}^d$, then there exists $\lambda_1, \dots, \lambda_{d+1} \in [0, 1]$ such that

$$\mathbf{z} = \lambda_1 \mathbf{u}_1 + \dots + \lambda_{d+1} \mathbf{u}_{d+1}, \quad \lambda_1 + \dots + \lambda_{d+1} = 1$$

We say that \mathbf{z} is contained in the interior of a k -dimensional face if exactly $k + 1$ of $\lambda_1, \dots, \lambda_{d+1}$ are nonzero, and define

$$d_{\mathbf{z}} \stackrel{\text{def.}}{=} \min\{k \in \mathbb{N}_+ : \mathbf{z} \text{ is contained in the interior of a } k\text{-dimensional face}\}$$

Using the idea above, we can find some linear paths that “push” \mathbf{x}, \mathbf{y} to some boundary points \mathbf{x}', \mathbf{y}' of some $(d_{\mathbf{x}} - 1), (d_{\mathbf{y}} - 1)$ -dimensional faces which contain \mathbf{x}, \mathbf{y} respectively, where we can use induction on $d_{\mathbf{x}} + d_{\mathbf{y}}$ to conclude ω -regularity at \mathbf{x}, \mathbf{y} .

Without loss of generality, assume $\mathbf{x} \in [0, 1]^d$. We first assume that $\mathbf{y} \in [0, 1]^d$ and demonstrate how to find such paths for \mathbf{x} and \mathbf{y} , then we generalize the result. As argued above, let $F_{\mathbf{x}}, F_{\mathbf{y}} \subseteq [0, 1]^d$ be $(d_{\mathbf{x}} - 1), (d_{\mathbf{y}} - 1)$ -dimensional faces which contain \mathbf{x}, \mathbf{y} , and $P_{\mathbf{x}}, P_{\mathbf{y}}$ be some simplices in $[0, 1]^d$ which contain $F_{\mathbf{x}}, F_{\mathbf{y}}$, respectively.

First, we define two infinite linear paths $\gamma_{\mathbf{x}}, \gamma_{\mathbf{y}} : [0, \infty) \rightarrow \mathbb{R}^d$ by:

$$\gamma_{\mathbf{x}}(t) \stackrel{\text{def.}}{=} t\mathbf{x}, \quad \gamma_{\mathbf{y}}(t) \stackrel{\text{def.}}{=} t\mathbf{y}, \quad \forall t \in [0, \infty)$$

Note that, $\|\gamma_{\mathbf{x}}(t) - \gamma_{\mathbf{y}}(t)\|_1 = t\|\mathbf{x} - \mathbf{y}\|_1$ changes linearly with time, and both \mathbf{x}, \mathbf{y}

are reached at the same time $t = 1$, thus conditions 3' and 4' are satisfied. Next, we modify $\gamma_{\mathbf{x}}$ and $\gamma_{\mathbf{y}}$ so that conditions 1' and 2' are satisfied, while ensuring that both conditions 3' and 4' remain satisfied.

From Equation 6.1.1, let $\tau_{\mathbf{x}}, \tau_{\mathbf{y}} \in S_d$ be the permutations corresponding to the simplices $P_{\mathbf{x}}, P_{\mathbf{y}}$. Let $\{\mathbf{v}_1, \dots, \mathbf{v}_{d_{\mathbf{x}}}\}, \{\mathbf{w}_1, \dots, \mathbf{w}_{d_{\mathbf{y}}}\} \subseteq \{0, 1\}^d$ be the vertices of $F_{\mathbf{x}}, F_{\mathbf{y}}$. Then, \mathbf{v}_i 's and \mathbf{w}_i 's have the form

$$\{\mathbf{v}_1, \dots, \mathbf{v}_{d_{\mathbf{x}}}\} = \left\{ \sum_{i=1}^{a_1} \mathbf{e}_{\tau_{\mathbf{x}}(i)}, \dots, \sum_{i=1}^{a_{d_{\mathbf{x}}}} \mathbf{e}_{\tau_{\mathbf{x}}(i)} \right\}, \quad 0 \leq a_1 < a_2 < \dots < a_{d_{\mathbf{x}}} \leq d$$

$$\{\mathbf{w}_1, \dots, \mathbf{w}_{d_{\mathbf{y}}}\} = \left\{ \sum_{i=1}^{b_1} \mathbf{e}_{\tau_{\mathbf{y}}(i)}, \dots, \sum_{i=1}^{b_{d_{\mathbf{y}}}} \mathbf{e}_{\tau_{\mathbf{y}}(i)} \right\}, \quad 0 \leq b_1 < b_2 < \dots < b_{d_{\mathbf{y}}} \leq d$$

Define new paths $\gamma_{\mathbf{x}}^*$ and $\gamma_{\mathbf{y}}^*$ as follows: let $(\gamma_{\mathbf{x}}^*)_i = (\gamma_{\mathbf{x}})_i$ for all coordinate components for which $i \neq \tau_{\mathbf{x}}(1), \dots, \tau_{\mathbf{x}}(a_1)$, otherwise let $(\gamma_{\mathbf{x}}^*)_i \equiv 1$; similarly, let $(\gamma_{\mathbf{y}}^*)_i = (\gamma_{\mathbf{y}})_i$ for all coordinate components for which $i \neq \tau_{\mathbf{y}}(1), \dots, \tau_{\mathbf{y}}(b_1)$, otherwise let $(\gamma_{\mathbf{y}}^*)_i \equiv 1$. Finally, let

$$T \stackrel{\text{def.}}{=} \max\{t \in [0, \infty) : \gamma_{\mathbf{x}}^*([0, t]) \subseteq F_{\mathbf{x}}, \gamma_{\mathbf{y}}^*([0, t]) \subseteq F_{\mathbf{y}}\}$$

For any $i = \tau_{\mathbf{x}}(1), \dots, \tau_{\mathbf{x}}(a_1)$, since $(\mathbf{v}_1)_i, \dots, (\mathbf{v}_{d_{\mathbf{x}}})_i = 1$ and \mathbf{x} is contained in the convex hull of $\{\mathbf{v}_1, \dots, \mathbf{v}_{d_{\mathbf{x}}}\}$, $(\mathbf{x})_i = 1$, thus $\gamma_{\mathbf{x}}^*(1) = \gamma_{\mathbf{x}}(1) = \mathbf{x} \in F_{\mathbf{x}}$, and since

$$\gamma_{\mathbf{x}}^*(0) = \sum_{i=1}^{a_1} \mathbf{e}_{\tau_{\mathbf{x}}(i)} \in \{\mathbf{v}_1, \dots, \mathbf{v}_{d_{\mathbf{x}}}\} \subseteq F_{\mathbf{x}},$$

where $F_{\mathbf{x}}$ is convex, $\gamma_{\mathbf{x}}^*([0, 1]) \subseteq F_{\mathbf{x}}$. Similarly, $\gamma_{\mathbf{y}}^*([0, 1]) \subseteq F_{\mathbf{y}}$. Thus, $T \geq 1$. If T is infinite, then both paths $\gamma_{\mathbf{x}}^*, \gamma_{\mathbf{y}}^*$ are constant, in this case both \mathbf{x}, \mathbf{y} must be lattice points, at which Φ has ω -regularity since the values of Φ and f are the same at \mathbf{x}, \mathbf{y} .

Therefore, we may assume that $T \geq 1$ is finite.

We now show that the finite linear paths $\gamma_{\mathbf{x}}^*(t), \gamma_{\mathbf{y}}^*(t)$ ($t \in [0, T]$) satisfy conditions 1' to 4'. Condition 1' is automatically satisfied by the definition of T . We have proved above that $\gamma_{\mathbf{x}}^*(1) = \mathbf{x}$, and $\gamma_{\mathbf{y}}^*(1) = \mathbf{y}$ can be proved similarly, also $T \geq 1$, thus condition 4' is satisfied.

For condition 2', since

$$\gamma_{\mathbf{x}}^*(0) = \sum_{i=1}^{a_1} \mathbf{e}_{\tau_{\mathbf{x}}(i)} \in \{0, 1\}^d, \quad \gamma_{\mathbf{y}}^*(0) = \sum_{i=1}^{b_1} \mathbf{e}_{\tau_{\mathbf{y}}(i)} \in \{0, 1\}^d$$

are lattice points on which the values of Φ and f are the same, we have ω -regularity of Φ at $\gamma_{\mathbf{x}}^*(0)$ and $\gamma_{\mathbf{y}}^*(0)$.

Since the faces $F_{\mathbf{x}}, F_{\mathbf{y}}$ of the convex polytopes $S_{\mathbf{x}}, S_{\mathbf{y}}$ are themselves convex polytopes, they are the convex hull of their vertices:

$$F_{\mathbf{x}} = \{p'_1 \mathbf{v}_1 + \cdots + p'_{d_{\mathbf{x}}} \mathbf{v}_{d_{\mathbf{x}}} : p'_1, \dots, p'_{d_{\mathbf{x}}} \geq 0, p'_1 + \cdots + p'_{d_{\mathbf{x}}} = 1\}$$

$$F_{\mathbf{y}} = \{q'_1 \mathbf{w}_1 + \cdots + q'_{d_{\mathbf{y}}} \mathbf{w}_{d_{\mathbf{y}}} : q'_1, \dots, q'_{d_{\mathbf{y}}} \geq 0, q'_1 + \cdots + q'_{d_{\mathbf{y}}} = 1\}$$

By the definition of T , at least one of the ending points $\gamma_{\mathbf{x}}^*(T), \gamma_{\mathbf{y}}^*(T)$ lies on the boundary of its face. Thus, without loss of generality, we may assume that $\gamma_{\mathbf{x}}^*(T)$ lies on the boundary of $F_{\mathbf{x}}$. Then, when $\gamma_{\mathbf{x}}^*(T)$ is expressed in the form $p'_1 \mathbf{v}_1 + \cdots + p'_{d_{\mathbf{x}}} \mathbf{v}_{d_{\mathbf{x}}}$, some of the coefficients p'_i will be 0 (unless $d_{\gamma_{\mathbf{x}}^*(T)} = 1$, which is the base case), so $\gamma_{\mathbf{x}}^*(T)$ lies in the interior of some $d_{\gamma_{\mathbf{x}}^*(T)}$ -dimensional face with $d_{\gamma_{\mathbf{x}}^*(T)} < d_{\mathbf{x}}$, while $\gamma_{\mathbf{y}}^*(T)$ is still contained in the $d_{\mathbf{y}}$ -dimensional face $F_{\mathbf{y}}$, so $d_{\gamma_{\mathbf{y}}^*(T)} \leq d_{\mathbf{y}}$. By the induction hypothesis, since $d_{\gamma_{\mathbf{x}}^*(T)} + d_{\gamma_{\mathbf{y}}^*(T)} < d_{\mathbf{x}} + d_{\mathbf{y}}$, Φ has ω -regularity at $\gamma_{\mathbf{x}}^*(T), \gamma_{\mathbf{y}}^*(T)$. In the base case where $d_{\gamma_{\mathbf{x}}^*(T)} = d_{\gamma_{\mathbf{y}}^*(T)} = 1$, \mathbf{x} and \mathbf{y} are lattice points on which the values of Φ and

f are the same, so we have ω -regularity of Φ at \mathbf{x} and \mathbf{y} . Therefore, condition 2' is satisfied.

Finally, for condition 3', let $i \in \{1, 2, \dots, d\}$ be arbitrary. As t changes linearly from 0 to T , $(\gamma_{\mathbf{x}}(t))_i - (\gamma_{\mathbf{y}}(t))_i = t((\mathbf{x})_i - (\mathbf{y})_i)$ changes linearly and does not change sign. By definitions of $\gamma_{\mathbf{x}}^*$ and $\gamma_{\mathbf{y}}^*$, compared to $\gamma_{\mathbf{x}}$ and $\gamma_{\mathbf{y}}$, $(\gamma_{\mathbf{x}}^*)_i$ and $(\gamma_{\mathbf{y}}^*)_i$ are either the same or always equal to 1. Then, there are three possibilities:

1. If both $(\gamma_{\mathbf{x}}^*(t))_i \equiv (\gamma_{\mathbf{x}}(t))_i$ and $(\gamma_{\mathbf{y}}^*(t))_i \equiv (\gamma_{\mathbf{y}}(t))_i$ for all $t \in [0, T]$, then as t changes linearly from 0 to T , $(\gamma_{\mathbf{x}}^*(t))_i - (\gamma_{\mathbf{y}}^*(t))_i = (\gamma_{\mathbf{x}}(t))_i - (\gamma_{\mathbf{y}}(t))_i = t((\mathbf{x})_i - (\mathbf{y})_i)$ changes linearly and does not change sign;
2. If exactly one of $(\gamma_{\mathbf{x}}^*(t))_i$ and $(\gamma_{\mathbf{y}}^*(t))_i$ always equal to 1, without loss of generality assume $(\gamma_{\mathbf{x}}^*(t))_i \equiv 1$, then $(\gamma_{\mathbf{x}}^*(t))_i - (\gamma_{\mathbf{y}}^*(t))_i = 1 - (\gamma_{\mathbf{y}}(t))_i \geq 1 - 1 = 0$ changes linearly with t and does not change sign;
3. If both of $(\gamma_{\mathbf{x}}^*(t))_i$ and $(\gamma_{\mathbf{y}}^*(t))_i$ always equal to 1, then $(\gamma_{\mathbf{x}}^*(t))_i - (\gamma_{\mathbf{y}}^*(t))_i \equiv 0$ changes linearly with t and does not change sign.

Therefore, $(\gamma_{\mathbf{x}}^*(t))_i - (\gamma_{\mathbf{y}}^*(t))_i$ always changes linearly with t and does not change sign, thus

$$\|\gamma_{\mathbf{x}}^*(t) - \gamma_{\mathbf{y}}^*(t)\|_1 = \sum_{i=1}^d |(\gamma_{\mathbf{x}}^*(t))_i - (\gamma_{\mathbf{y}}^*(t))_i|$$

changes linearly with t , hence condition 3' is satisfied.

Now, we have found suitable paths for $\mathbf{y} \in [0, 1]^d$, next we show how to find such paths for an arbitrary $\mathbf{y}' \in [0, n]^d$. Suppose $\mathbf{y}' \in Q_{\mathbf{y}'} \stackrel{\text{def.}}{=} [m_1, m_1 + 1] \times [m_2, m_2 + 1] \times \dots \times [m_d, m_d + 1]$ for some integers m_1, \dots, m_d . Define the function $T_{\mathbf{y}'} : [0, 1]^d \rightarrow Q_{\mathbf{y}'}$

by:

$$(T_{\mathbf{y}'}(\mathbf{z}))_i \stackrel{\text{def.}}{=} \begin{cases} (\mathbf{z})_i + 2k_i, & \text{if } m_i = 2k_i \text{ for some } k_i \in \mathbb{N} \\ 2k_i - (\mathbf{z})_i, & \text{if } m_i = 2k_i - 1 \text{ for some } k_i \in \mathbb{N} \end{cases}, \quad i = 1, 2, \dots, d$$

Then $T_{\mathbf{y}'}$ is a composition of reflections and translations, thus it is a linear bijection between $[0, 1]^d$ and $Q_{\mathbf{y}'}$.

Let $\cup H$ be the union of the hyperplanes (which are considered as subsets of \mathbb{R}^d) in H , which is a subset of \mathbb{R}^d that triangulates every unit cube in $\{[i, i + 1] : i \in \mathbb{Z}\}^d$. By the definition of H , $\cup H$ has the following two properties:

1. Translation invariance:

$$\mathbf{z} \in \cup H \iff (\mathbf{z} + 2\mathbf{e}_i) \in \cup H, \quad i = 1, 2, \dots, d$$

2. Reflection invariance:

$$\mathbf{z} \in \cup H \iff (2\mathbf{e}_i - \mathbf{z}) \in \cup H, \quad i = 1, 2, \dots, d$$

Then, by the definition of $T_{\mathbf{y}'}$, the part of $\cup H$ inside $Q_{\mathbf{y}'}$ is exactly the image of the part of $\cup H$ inside $[0, 1]^d$ under $T_{\mathbf{y}'}$, and vice versa: the part of $\cup H$ inside $[0, 1]^d$ is exactly the image of the part of $\cup H$ inside $Q_{\mathbf{y}'}$ under $T_{\mathbf{y}'}^{-1}$. Therefore, the simplices in $[0, 1]^d$ and the simplices in $Q_{\mathbf{y}'}$ are in one-to-one correspondence via $T_{\mathbf{y}'}$.

Let $\mathbf{y} \stackrel{\text{def.}}{=} T_{\mathbf{y}'}^{-1}(\mathbf{y}')$. Since $\mathbf{x}, \mathbf{y} \in [0, 1]^d$, by the arguments above, there exist finite linear paths $\gamma_{\mathbf{x}}^*(t), \gamma_{\mathbf{y}}^*(t)$ ($t \in [0, T]$) that satisfy conditions 1' to 4'. Using the same notations as above, let $F_{\mathbf{x}}, F_{\mathbf{y}} \subseteq [0, 1]^d$ be $(d_{\mathbf{x}} - 1), (d_{\mathbf{y}} - 1)$ -dimensional faces which

contain \mathbf{x} , \mathbf{y} , and $P_{\mathbf{x}}, P_{\mathbf{y}}$ be some simplices in $[0, 1]^d$ which contain $F_{\mathbf{x}}, F_{\mathbf{y}}$, respectively.

Define the path for \mathbf{y}' by:

$$\gamma_{\mathbf{y}'}^*(t) \stackrel{\text{def.}}{=} T_{\mathbf{y}'}(\gamma_{\mathbf{y}}^*(t)), \quad t \in [0, T]$$

We now show that the finite linear paths $\gamma_{\mathbf{x}}^*(t), \gamma_{\mathbf{y}'}^*(t)$ ($t \in [0, T]$) satisfy condition 1' to 4'. Since $\gamma_{\mathbf{y}}^*([0, T]) \subseteq P_{\mathbf{y}}$, $\gamma_{\mathbf{y}'}^*([0, T]) \subseteq T_{\mathbf{y}'}(P_{\mathbf{y}})$ and $T_{\mathbf{y}'}(P_{\mathbf{y}})$ is a simplex, so condition 1' is satisfied. Similar as above, since $\gamma_{\mathbf{x}}^*(0), \gamma_{\mathbf{y}'}^*(0) = T_{\mathbf{y}'}(\gamma_{\mathbf{y}}^*(0))$ are lattice points and $\gamma_{\mathbf{x}}^*(T), \gamma_{\mathbf{y}'}^*(T) = T_{\mathbf{y}'}(\gamma_{\mathbf{y}}^*(T))$ are in the boundary of their faces $F_{\mathbf{x}}, T_{\mathbf{y}'}(F_{\mathbf{y}})$, by doing induction on $d_{\mathbf{x}} + d_{\mathbf{y}'}$, we conclude that Φ has ω -regularity at $\gamma_{\mathbf{x}}^*(0), \gamma_{\mathbf{y}'}^*(0)$ and $\gamma_{\mathbf{x}}^*(T), \gamma_{\mathbf{y}'}^*(T)$, so condition 2' is satisfied. Since $\gamma_{\mathbf{x}}^*(1) = \mathbf{x}$ and $\gamma_{\mathbf{y}'}^*(1) = T_{\mathbf{y}'}(\gamma_{\mathbf{y}}^*(1)) = T_{\mathbf{y}'}(\mathbf{y}) = \mathbf{y}'$, condition 4' is satisfied.

Finally, for condition 3', if $(\mathbf{y})_i \in [0, 1]$, then $m_i = 0$, $k_i = 0$, $(\gamma_{\mathbf{y}'}^*)_i = 2k_i + (\gamma_{\mathbf{y}}^*)_i = (\gamma_{\mathbf{y}}^*)_i$, and since $(\gamma_{\mathbf{x}}^*(t))_i - (\gamma_{\mathbf{y}}^*(t))_i$ always changes linearly with t and does not change sign, so does $(\gamma_{\mathbf{x}}^*(t))_i - (\gamma_{\mathbf{y}'}^*(t))_i$. Otherwise, if $(\mathbf{y})_i > 1$, then $m_i \geq 1$, $k_i \geq 1$, $(\gamma_{\mathbf{y}'}^*)_i \geq 2k_i - (\gamma_{\mathbf{y}}^*)_i \geq 2 - 1 = 1 \geq (\gamma_{\mathbf{x}}^*)_i$, so $(\gamma_{\mathbf{x}}^*)_i - (\gamma_{\mathbf{y}'}^*)_i$ does not change sign, and obviously it changes linearly with t . Thus,

$$\|\gamma_{\mathbf{x}}^*(t) - \gamma_{\mathbf{y}'}^*(t)\|_1 = \sum_{i=1}^d |(\gamma_{\mathbf{x}}^*(t))_i - (\gamma_{\mathbf{y}'}^*(t))_i|$$

changes linearly with t , hence condition 3' is satisfied.

Therefore, the paths $\gamma_{\mathbf{x}}^*, \gamma_{\mathbf{y}'}^*$ for \mathbf{x}, \mathbf{y}' indeed satisfy conditions 1' to 4'. By the argument above, Φ has ω -regularity at \mathbf{x}, \mathbf{y}' . Since \mathbf{x}, \mathbf{y}' were chosen arbitrarily, we conclude that ω is also a modulus of regularity of Φ under ℓ^1 norm of \mathbb{R}^d .

For the error estimate, fix any $\mathbf{x} \in [0, n]^d$. Suppose $\mathbf{x} \in Q_{\mathbf{x}} \stackrel{\text{def.}}{=} [n_1, n_1 + 1] \times$

$[n_2, n_2 + 1] \times \cdots \times [n_d, n_d + 1]$ for some integers n_1, \dots, n_d . Let $\mathbf{u}_1, \dots, \mathbf{u}_{d+1}$ be the vertices of a simplex $P_{\mathbf{x}} \subseteq Q_{\mathbf{x}}$ which contains \mathbf{x} , then there exist $\lambda_1, \dots, \lambda_{d+1} \geq 0$ such that

$$\mathbf{x} = \lambda_1 \mathbf{u}_1 + \cdots + \lambda_{d+1} \mathbf{u}_{d+1}, \quad \lambda_1 + \cdots + \lambda_{d+1} = 1$$

Since Φ is affine on this simplex,

$$\Phi(\mathbf{x}) = \sum_{i=1}^{d+1} \lambda_i \Phi(\mathbf{u}_i) = \sum_{i=1}^{d+1} \lambda_i f(\mathbf{u}_i)$$

thus

$$\begin{aligned} |\Phi(\mathbf{x}) - f(\mathbf{x})| &= |\lambda_1(f(\mathbf{u}_1) - f(\mathbf{x})) + \cdots + \lambda_{d+1}(f(\mathbf{u}_{d+1}) - f(\mathbf{x}))| \\ &\leq \lambda_1 |f(\mathbf{u}_1) - f(\mathbf{x})| + \cdots + \lambda_{d+1} |f(\mathbf{u}_{d+1}) - f(\mathbf{x})| \\ &\leq \lambda_1 \omega(\|\mathbf{u}_1 - \mathbf{x}\|_1) + \cdots + \lambda_{d+1} \omega(\|\mathbf{u}_{d+1} - \mathbf{x}\|_1) \\ &\leq \omega(\lambda_1 \|\mathbf{u}_1 - \mathbf{x}\|_1 + \cdots + \lambda_{d+1} \|\mathbf{u}_{d+1} - \mathbf{x}\|_1) \end{aligned} \tag{7.5.2}$$

where (7.5.2) follows by the concavity of the modulus ω , and

$$\begin{aligned} \sum_{j=1}^{d+1} \lambda_j \|\mathbf{u}_j - \mathbf{x}\|_1 &= \sum_{j=1}^{d+1} \lambda_j \sum_{i=1}^d |(\mathbf{u}_j)_i - (\mathbf{x})_i| \\ &= \sum_{i=1}^d \sum_{j=1}^{d+1} \lambda_j |(\mathbf{u}_j)_i - (\mathbf{x})_i| \\ &= \sum_{i=1}^d \left(\sum_{\substack{j=1 \\ (\mathbf{u}_j)_i = n_i}}^{d+1} \lambda_j |(\mathbf{u}_j)_i - (\mathbf{x})_i| + \sum_{\substack{j=1 \\ (\mathbf{u}_j)_i = n_i+1}}^{d+1} \lambda_j |(\mathbf{u}_j)_i - (\mathbf{x})_i| \right) \end{aligned}$$

$$\begin{aligned}
&= \sum_{i=1}^d \left(\sum_{\substack{j=1 \\ (\mathbf{u}_j)_i = n_i}}^{d+1} \lambda_j ((\mathbf{x})_i - n_i) + \sum_{\substack{j=1 \\ (\mathbf{u}_j)_i = n_i + 1}}^{d+1} \lambda_j (n_i + 1 - (\mathbf{x})_i) \right) \\
&= \sum_{i=1}^d \left(((\mathbf{x})_i - n_i) \sum_{\substack{j=1 \\ (\mathbf{u}_j)_i = n_i}}^{d+1} \lambda_j + (1 - ((\mathbf{x})_i - n_i)) \sum_{\substack{j=1 \\ (\mathbf{u}_j)_i = n_i + 1}}^{d+1} \lambda_j \right)
\end{aligned}$$

for which we have

$$\begin{aligned}
(\mathbf{x})_i - n_i &= \left(\sum_{j=1}^{d+1} \lambda_j (\mathbf{u}_j)_i \right) - n_i \\
&= \left(\sum_{\substack{j=1 \\ (\mathbf{u}_j)_i = n_i}}^{d+1} \lambda_j n_i + \sum_{\substack{j=1 \\ (\mathbf{u}_j)_i = n_i + 1}}^{d+1} \lambda_j (n_i + 1) \right) - n_i \\
&= \sum_{\substack{j=1 \\ (\mathbf{u}_j)_i = n_i + 1}}^{d+1} \lambda_j \\
&= 1 - \sum_{\substack{j=1 \\ (\mathbf{u}_j)_i = n_i}}^{d+1} \lambda_j
\end{aligned}$$

thus

$$\begin{aligned}
((\mathbf{x})_i - n_i) \sum_{\substack{j=1 \\ (\mathbf{u}_j)_i = n_i}}^{d+1} \lambda_j &= \left(1 - \sum_{\substack{j=1 \\ (\mathbf{u}_j)_i = n_i}}^{d+1} \lambda_j \right) \sum_{\substack{j=1 \\ (\mathbf{u}_j)_i = n_i}}^{d+1} \lambda_j \leq \frac{1}{4} \\
(1 - ((\mathbf{x})_i - n_i)) \sum_{\substack{j=1 \\ (\mathbf{u}_j)_i = n_i + 1}}^{d+1} \lambda_j &= \left(1 - \sum_{\substack{j=1 \\ (\mathbf{u}_j)_i = n_i + 1}}^{d+1} \lambda_j \right) \sum_{\substack{j=1 \\ (\mathbf{u}_j)_i = n_i + 1}}^{d+1} \lambda_j \leq \frac{1}{4}
\end{aligned}$$

Combining the results above, we have that

$$\begin{aligned} \sum_{j=1}^{d+1} \lambda_j \|\mathbf{u}_j - \mathbf{x}\|_1 &= \sum_{i=1}^d \left(((\mathbf{x})_i - n_i) \sum_{\substack{j=1 \\ (\mathbf{u}_j)_i = n_i}}^{d+1} \lambda_j + (1 - ((\mathbf{x})_i - n_i)) \sum_{\substack{j=1 \\ (\mathbf{u}_j)_i = n_i + 1}}^{d+1} \lambda_j \right) \\ &\leq \sum_{i=1}^d \left(\frac{1}{4} + \frac{1}{4} \right) \\ &= \frac{d}{2} \end{aligned}$$

Therefore,

$$|\Phi(\mathbf{x}) - f(\mathbf{x})| \leq \omega \left(\sum_{j=1}^{d+1} \lambda_j \|\mathbf{u}_j - \mathbf{x}\|_1 \right) \leq \omega \left(\frac{d}{2} \right)$$

Since $\mathbf{x} \in [0, n]^d$ was chosen arbitrarily, we conclude that

$$\|f - \Phi\|_{L^\infty([0, n]^d)} \leq \omega \left(\frac{d}{2} \right).$$

This concludes our proof. □

Indeed, we may deduce an even stronger conclusion. Namely, we only used information of f at the lattice points, so the modulus of regularity ω_Φ of Φ can be the minimum concave function that satisfies

$$\omega_\Phi(i) = \max_{\substack{\mathbf{x}, \mathbf{y} \in \{0, 1, \dots, n\}^d \\ \|\mathbf{x} - \mathbf{y}\|_1 = i}} |f(\mathbf{x}) - f(\mathbf{y})|, \quad \text{for } i = 0, 1, \dots, nd$$

which is a polygonal function. Therefore, we have the following corollary.

Corollary 7.5.1. *Let f be a function from $[0, n]^d$ to \mathbb{R} for some $n \in \mathbb{N}_+$. Let ω_f be a modulus of regularity of f under ℓ^1 norm of \mathbb{R}^d . Let $\Phi : [0, n]^d \rightarrow \mathbb{R}$ be continuous*

piecewise linear on each of the n^d unit cubes $\{[0, 1], [1, 2], \dots, [n-1, n]\}^d$ with respect to the triangulation in Lemma 7.5.1, and

$$\Phi(\mathbf{y}) = f(\mathbf{y}), \quad \forall \mathbf{y} \in \{0, 1, \dots, n\}^d$$

Then, any monotone increasing and concave function ω_Φ that satisfies the following condition is a modulus of regularity of Φ under ℓ^1 norm of \mathbb{R}^d :

$$\omega_\Phi(i) = \max_{\substack{\mathbf{x}, \mathbf{y} \in \{0, 1, \dots, n\}^d \\ \|\mathbf{x} - \mathbf{y}\|_1 = i}} |f(\mathbf{x}) - f(\mathbf{y})|, \quad \text{for } i = 0, 1, \dots, nd$$

Moreover,

$$\|f - \Phi\|_{L^\infty([0, n]^d)} \leq \omega_f\left(\frac{d}{2}\right)$$

7.6 Proof of the Main Result (Theorem 4.1.1)

Using the results above, we are able to derive our first main theorem, namely Theorem 4.1.1.

Proof of Theorem 4.1.1. To apply the previous results, we first scale f into a function \tilde{f} defined on $[0, n]^d$:

$$\tilde{f}(\mathbf{x}) \stackrel{\text{def.}}{=} f\left(\frac{\mathbf{x}}{n}\right), \quad \forall \mathbf{x} \in [0, n]^d \tag{7.6.1}$$

Since ω is a modulus of regularity of f , for all $\mathbf{x}, \mathbf{y} \in [0, 1]^d$, we have

$$|\tilde{f}(\mathbf{x}) - \tilde{f}(\mathbf{y})| = \left| f\left(\frac{\mathbf{x}}{n}\right) - f\left(\frac{\mathbf{y}}{n}\right) \right| \leq \omega\left(\left\|\frac{\mathbf{x}}{n} - \frac{\mathbf{y}}{n}\right\|_1\right) = \omega\left(\frac{\|\mathbf{x} - \mathbf{y}\|_1}{n}\right)$$

then $\tilde{\omega}$ defined below is a modulus of regularity of \tilde{f} :

$$\tilde{\omega}(x) \stackrel{\text{def.}}{=} \omega\left(\frac{x}{n}\right), \quad \forall x \in [0, nd]$$

Let

$$\tilde{\Phi}(\mathbf{x}) \stackrel{\text{def.}}{=} \sum_{\mathbf{y} \in \{0,1,\dots,n\}^d} \tilde{f}(\mathbf{y}) T_{\mathbf{y}}(\mathbf{x})$$

where $T_{\mathbf{y}}$ be the hat function

$$T_{\mathbf{y}}(\mathbf{x}) \stackrel{\text{def.}}{=} \sigma\left(1 - \max\{|x_i - y_i| : y_i \text{ is even}\} - \max\{|x_i - y_i| : y_i \text{ is odd}\}\right)$$

for $\mathbf{x} = (x_1, \dots, x_d) \in [0, n]^d$. By Lemma 7.5.3, $\tilde{\Phi}$ is continuous piecewise linear on each of the n^d unit cubes $\{[0, 1], [1, 2], \dots, [n-1, n]\}^d$ of $[0, n]^d$ with respect to the triangulation in Lemma 7.5.1, and

$$\tilde{\Phi}(\mathbf{y}) = \tilde{f}(\mathbf{y}), \quad \forall \mathbf{y} \in \{0, 1, \dots, n\}^d \tag{7.6.2}$$

Then, by Lemma 7.5.4, $\tilde{\omega}$ is also a modulus of regularity of $\tilde{\Phi}$ under ℓ^1 norm of \mathbb{R}^d .

Moreover,

$$\|\tilde{f} - \tilde{\Phi}\|_{L^\infty([0,n]^d)} \leq \tilde{\omega}\left(\frac{d}{2}\right)$$

Now let

$$\Phi(\mathbf{x}) \stackrel{\text{def.}}{=} \tilde{\Phi}(n\mathbf{x}), \quad \forall \mathbf{x} \in [0, 1]^d \tag{7.6.3}$$

Then, we deduce can Φ uniformly approximates f on $[0, 1]^d$ since

$$\|f - \Phi\|_{L^\infty([0,1]^d)} = \|\tilde{f} - \tilde{\Phi}\|_{L^\infty([0,n]^d)} \leq \tilde{\omega}\left(\frac{d}{2}\right) = \omega\left(\frac{d}{2n}\right)$$

Define the number of samples point $N(n) \stackrel{\text{def.}}{=} (1+n)^d$ and the grid $\{\mathbf{x}_i\}_{i=1}^{N(n)} \stackrel{\text{def.}}{=} n^{-1} \cdot \mathbb{X}_n$. Note that, together (7.6.1) and (7.6.2) along with the definition of Φ given in (7.6.3) implies that

$$\Phi(\mathbf{x}_i) = f(\mathbf{x}_i) \quad \forall i = 1, \dots, N(n).$$

In other words, Φ interpolates f on the grid $\{\mathbf{x}_i\}_{i=1}^{N(n)}$. Thus, $\sum_{i=1}^{N(n)} |\Phi(\mathbf{x}_i) - f(\mathbf{x}_i)| = 0$.

Since $\tilde{\omega}$ is a modulus of regularity of $\tilde{\Phi}$, for all $\mathbf{x}, \mathbf{y} \in [0, 1]^d$,

$$|\Phi(\mathbf{x}) - \Phi(\mathbf{y})| = |\tilde{\Phi}(n\mathbf{x}) - \tilde{\Phi}(n\mathbf{y})| \leq \tilde{\omega}(\|n\mathbf{x} - n\mathbf{y}\|_1) = \omega(\|\mathbf{x} - \mathbf{y}\|_1)$$

Therefore, ω is a modulus of regularity of Φ . Now we remain to show that Φ can be implemented by a ReLU MLP with width at most $8d(n+1)^d$ and depth at most $\lceil \log_2 d \rceil + 4$.

It is easy to see that

$$x = \sigma(x) - \sigma(-x), \quad |x| = \sigma(x) + \sigma(-x), \quad \forall x \in \mathbb{R}$$

then

$$\max\{x, y\} = \frac{x+y}{2} + \frac{|x-y|}{2} = \frac{\sigma(x+y)}{2} - \frac{\sigma(-x-y)}{2} + \frac{\sigma(x-y)}{2} + \frac{\sigma(-x+y)}{2}, \quad \forall x, y \in \mathbb{R}$$

Using this formula, for any $k \in \mathbb{N}_+$, we may compute the maximum value function on

2^k arguments via the following network structure:

$$\begin{array}{c}
 \left[\begin{array}{c} x_1^{(1)} \\ x_2^{(1)} \\ x_3^{(1)} \\ x_4^{(1)} \\ \vdots \\ x_{2^{k-1}}^{(1)} \\ x_{2^k}^{(1)} \end{array} \right] \Rightarrow \left[\begin{array}{c} \sigma \left(\frac{\pm x_1^{(1)} \pm x_2^{(1)}}{2} \right) \\ \sigma \left(\frac{\pm x_3^{(1)} \pm x_4^{(1)}}{2} \right) \\ \vdots \\ \sigma \left(\frac{\pm x_{2^{k-1}}^{(1)} \pm x_{2^k}^{(1)}}{2} \right) \end{array} \right] \Rightarrow \left[\begin{array}{c} \max\{x_1^{(1)}, x_2^{(1)}\} \stackrel{\text{def.}}{=} x_1^{(2)} \\ \max\{x_3^{(1)}, x_4^{(1)}\} \stackrel{\text{def.}}{=} x_2^{(2)} \\ \vdots \\ \max\{x_{2^{k-1}}^{(1)}, x_{2^k}^{(1)}\} \stackrel{\text{def.}}{=} x_{2^{k-1}}^{(2)} \end{array} \right] \\
 \\
 \Rightarrow \left[\begin{array}{c} \sigma \left(\frac{\pm x_1^{(2)} \pm x_2^{(2)}}{2} \right) \\ \sigma \left(\frac{\pm x_3^{(2)} \pm x_4^{(2)}}{2} \right) \\ \vdots \\ \sigma \left(\frac{\pm x_{2^{k-1}-1}^{(2)} \pm x_{2^k-1}^{(2)}}{2} \right) \end{array} \right] \Rightarrow \left[\begin{array}{c} \max\{x_1^{(2)}, x_2^{(2)}\} \stackrel{\text{def.}}{=} x_1^{(3)} \\ \max\{x_3^{(2)}, x_4^{(2)}\} \stackrel{\text{def.}}{=} x_2^{(3)} \\ \vdots \\ \max\{x_{2^{k-1}-1}^{(2)}, x_{2^k-1}^{(2)}\} \stackrel{\text{def.}}{=} x_{2^{k-2}}^{(3)} \end{array} \right]
 \end{array}$$

$$\begin{aligned}
&\implies \dots \implies \begin{bmatrix} \max\{x_1^{(k-1)}, x_2^{(k-1)}\} \stackrel{\text{def.}}{=} x_1^{(k)} \\ \max\{x_3^{(k-1)}, x_4^{(k-1)}\} \stackrel{\text{def.}}{=} x_2^{(k)} \end{bmatrix} \\
&\implies \left[\sigma \left(\frac{\pm x_1^{(k)} \pm x_2^{(k)}}{2} \right) \right] \implies \left[\max\{x_1^{(k)}, x_2^{(k)}\} \right] = \left[\max\{x_1^{(1)}, x_2^{(1)}, \dots, x_{2^k}^{(1)}\} \right]
\end{aligned} \tag{7.6.4}$$

where $\sigma(\pm a \pm b)$ abbreviates the four neurons $\sigma(a+b)$, $\sigma(a-b)$, $\sigma(-a+b)$ and $\sigma(-a-b)$, and the column vectors in green represent the “auxiliary” hidden layers which are simply affine transformations that don’t use any activation functions, thus each of them can be integrated with the layer after it, and does not require an extra hidden layer to process. Therefore, we will ignore these layers. We denote this network as \mathbf{M}_k , thus $\mathbf{M}_k : \mathbb{R}^{2^k} \rightarrow \mathbb{R}$ outputs the maximum value of its 2^k input arguments, and it has width 2^{k+1} and depth $k + 1$. We note that, after deleting an appropriate set of neurons in \mathbf{M}_k (which we will still call \mathbf{M}_k), it can be applied to k' arguments for any $1 \leq k' \leq 2^k$ and output their maximum.

Let

$$d_0 \stackrel{\text{def.}}{=} \lceil \log_2 d \rceil$$

For any lattice point $\mathbf{y} = (y_1, y_2, \dots, y_d) \in \{0, 1, \dots, n\}^d$, suppose $y_{i_1}, y_{i_2}, \dots, y_{i_p}$ are even and $y_{j_1}, y_{j_2}, \dots, y_{j_q}$ are odd, for some $0 \leq p, q \leq d$, then we can implement the

hat function $T_{\mathbf{y}}$ using the network \mathbf{M}_{d_0} as follows:

$$\begin{aligned}
\begin{bmatrix} x_1 \\ x_2 \\ \vdots \\ x_d \end{bmatrix} &\Rightarrow \begin{bmatrix} \sigma(\pm(x_1 - y_1)) \\ \sigma(\pm(x_2 - y_2)) \\ \vdots \\ \sigma(\pm(x_d - y_d)) \end{bmatrix} \Rightarrow \begin{bmatrix} \sigma(x_1 - y_1) + \sigma(-(x_1 - y_1)) = |x_1 - y_1| \\ \sigma(x_2 - y_2) + \sigma(-(x_2 - y_2)) = |x_2 - y_2| \\ \vdots \\ \sigma(x_d - y_d) + \sigma(-(x_d - y_d)) = |x_d - y_d| \end{bmatrix} \\
&\Rightarrow \begin{bmatrix} \mathbf{M}_{d_0} \left(\begin{bmatrix} |x_{i_1} - y_{i_1}| & |x_{i_2} - y_{i_2}| & \cdots & |x_{i_p} - y_{i_p}| \end{bmatrix}^\top \right) \\ \mathbf{M}_{d_0} \left(\begin{bmatrix} |x_{j_1} - y_{j_1}| & |x_{j_2} - y_{j_2}| & \cdots & |x_{j_q} - y_{j_q}| \end{bmatrix}^\top \right) \end{bmatrix} \\
&= \begin{bmatrix} \max \{|x_i - y_i| : y_i \text{ is even}\} \\ \max \{|x_i - y_i| : y_i \text{ is odd}\} \end{bmatrix} \\
&\Rightarrow \left[\sigma(1 - \max \{|x_i - y_i| : y_i \text{ is even}\}) - \max \{|x_i - y_i| : y_i \text{ is odd}\} \right] \\
&= T_{\mathbf{y}} \left(\begin{bmatrix} x_1 & x_2 & \cdots & x_d \end{bmatrix}^\top \right) \tag{7.6.5}
\end{aligned}$$

As before, $\sigma(\pm a)$ abbreviates the two neurons $\sigma(a), \sigma(-a)$, and the column vectors in green represent the “auxiliary” hidden layer which we can ignore. We call this network $\mathbf{T}_{\mathbf{y}}$, which has width at most

$$\text{width}(\mathbf{T}_{\mathbf{y}}) = \text{width}(\mathbf{M}_{d_0}) + \text{width}(\mathbf{M}_{d_0}) \leq 2^{d_0+1} + 2^{d_0+1} = 2^{d_0+2} \leq 2^{(\log_2 d+1)+2} = 8d \tag{7.6.6}$$

and depth at most

$$\text{depth}(\mathbf{T}_{\mathbf{y}}) = 1 + \text{depth}(\mathbf{M}_{d_0}) + 1 \leq d_0 + 3 \tag{7.6.7}$$

Let $\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_{(n+1)^d}$ be the $(n+1)^d$ lattice points inside the cube $[0, n]^d$. Since

$$\Phi(\mathbf{x}) = \tilde{\Phi}(n\mathbf{x}) = \sum_{\mathbf{y} \in \{0,1,\dots,n\}^d} \tilde{f}(\mathbf{y}) T_{\mathbf{y}}(n\mathbf{x}), \quad \forall \mathbf{x} \in [0, n]^d$$

then Φ can be implemented as follows:

$$\mathbf{x} \implies n\mathbf{x} \implies \begin{bmatrix} \mathbf{T}_{\mathbf{y}_1}(n\mathbf{x}) \\ \mathbf{T}_{\mathbf{y}_2}(n\mathbf{x}) \\ \vdots \\ \mathbf{T}_{\mathbf{y}_{(n+1)^d}}(n\mathbf{x}) \end{bmatrix} \implies \left[\sum_{i=1}^{(n+1)^d} \tilde{f}(\mathbf{y}_i) T_{\mathbf{y}_i}(n\mathbf{x}) \right] = \Phi(\mathbf{x}) \quad (7.6.8)$$

This network has width

$$\text{width}(\Phi) = \sum_{i=1}^{(n+1)^d} \text{width}(\mathbf{T}_{\mathbf{y}_i}) \leq \sum_{i=1}^{(n+1)^d} 8d = 8d(n+1)^d$$

and depth

$$\text{depth}(\Phi) = \max_{1 \leq i \leq (n+1)^d} \text{depth}(\mathbf{T}_{\mathbf{y}_i}) + 1 \leq (d_0 + 3) + 1 = \lceil \log_2 d \rceil + 4.$$

Finally, we tally the nonzero parameters in Φ . From (7.6.4), the network \mathbf{M}_k has no nonzero biases, and has nonzero weights at most

$$M \stackrel{\text{def.}}{=} 4 \cdot 2^{k-1} + 8 \cdot \sum_{i=1}^{k-2} 2^i \leq 2 \cdot 2^k + 8 \cdot 2^{k-1} = 6 \cdot 2^k$$

where the factor 4 colored in blue is the number of nonzero connections between the second layer in (7.6.4) and the second layer in (7.6.5), after “integrating” the green

auxiliary layer in (7.6.5) with the layer after it. In other words, this is because in (7.6.5), each of the neurons in the input layer of \mathbf{M}_{d_0} is the sum of two neurons in the second layer, so since in (7.6.4) each neuron in the second layer has 2 nonzero connection with the first layer, in (7.6.5) this number will have to be doubled, becoming 4.

From (7.6.5), the network \mathbf{T}_y has at most $T^{(b)} \stackrel{\text{def.}}{=} 2d + 1$ nonzero biases, and has nonzero weights at most

$$T^{(w)} \stackrel{\text{def.}}{=} 2d + M + 2 = 2d + 6 \cdot 2^{d_0} + 2 \leq 2d + 6(2d - 1) + 2 \leq 14d - 4$$

In total, the network \mathbf{T}_y has nonzero parameters at most

$$T \stackrel{\text{def.}}{=} T^{(w)} + T^{(b)} = 16d - 3 \tag{7.6.9}$$

From (7.6.8), Φ has no additional nonzero biases, thus it has nonzero parameters at most

$$(n + 1)^d T + (n + 1)^d \leq 16d(n + 1)^d. \tag{7.6.10}$$

□

Remark 7.6.1 (Proof of Proposition 4.1.1 is given in the Proof of Theorem 4.1.1). The proof of Proposition 4.1.1 directly follows from the proof of Theorem 4.1.1, upon noting that all parameters except for the ones between the input layer and the first hidden layer and the ones between the output layer and the last hidden layer come from $\{0, \pm 1/2\}$, and all parameters except for the ones between the output layer and the last hidden layer are independent of the “sample values” $\{f(\mathbf{x}_i)\}_{i=1}^{N(n)}$; i.e. the value

of the encoder $\mathcal{E}_n(f)$.

7.7 Proof of Secondary Approximation Results

Proof of Proposition 4.1.2. In the proof of Theorem 4.1.1, we constructed final network in (7.6.8). Here, we only need to distribute the $(n+1)^d$ sub-networks $\mathbf{T}_{\mathbf{y}_1}, \dots, \mathbf{T}_{\mathbf{y}_{(n+1)^d}}$ to L different layers.

We relabel the $(n+1)^d$ sub-network and function value pairs $(\mathbf{T}_{\mathbf{y}_j}, f(\mathbf{y}_j))$ for $j = 1, \dots, (n+1)^d$ as $(\mathbf{T}_1^{(i)}, f_1^{(i)}), \dots, (\mathbf{T}_{m_i}^{(i)}, f_{m_i}^{(i)})$ for $i = 1, \dots, L$. To simplify notations, for $k = 1, \dots, L$, let

$$T_k(\mathbf{x}) \stackrel{\text{def.}}{=} \sum_{i=1}^k \sum_{j=1}^{m_i} f_j^{(i)} \mathbf{T}_j^{(i)}(n\mathbf{x}), \quad \forall \mathbf{x} \in [0, 1]^d,$$

Then, $\Phi(\mathbf{x}) = T_L(\mathbf{x})$. Now, we can construct Φ as follows:

$$\begin{aligned} \mathbf{x} &\implies \begin{bmatrix} \mathbf{T}_1^{(1)}(n\mathbf{x}) \\ \vdots \\ \mathbf{T}_{m_1}^{(1)}(n\mathbf{x}) \\ n\mathbf{x} \end{bmatrix} \implies \begin{bmatrix} \mathbf{T}_1^{(2)}(n\mathbf{x}) \\ \vdots \\ \mathbf{T}_{m_2}^{(2)}(n\mathbf{x}) \\ \sigma(\pm T_1(\mathbf{x})) \\ n\mathbf{x} \end{bmatrix} \implies \begin{bmatrix} \mathbf{T}_1^{(3)}(n\mathbf{x}) \\ \vdots \\ \mathbf{T}_{m_3}^{(3)}(n\mathbf{x}) \\ \sigma(\pm T_2(\mathbf{x})) \\ n\mathbf{x} \end{bmatrix} \implies \dots \\ &\implies \begin{bmatrix} \mathbf{T}_1^{(L)}(n\mathbf{x}) \\ \vdots \\ \mathbf{T}_{m_L}^{(L)}(n\mathbf{x}) \\ \sigma(\pm T_{L-1}(\mathbf{x})) \\ n\mathbf{x} \end{bmatrix} \implies T_L(\mathbf{x}) = \Phi(\mathbf{x}) \end{aligned}$$

where we used the fact that $\sigma(n\mathbf{x}) = n\mathbf{x}$ for any $\mathbf{x} \in [0, 1]^d$. We know from (7.6.6) and (7.6.7) that the sub-networks $\mathbf{T}_j^{(i)}$ has width $W \leq 8d$, depth $D \leq \lceil \log_2 d \rceil + 3$, and nonzero parameters $T \leq 16d - 3$. Then, Φ has width at most

$$\text{width}(\Phi) \leq W \max\{m_1, \dots, m_L\} + 2 + d = 8d \max\{m_1, \dots, m_L\} + d + 2$$

depth at most

$$\text{depth}(\Phi) \leq L(D + 1) = L(\lceil \log_2 d \rceil + 4)$$

and nonzero parameters at most

$$\begin{aligned} \text{par}(\Phi) &\leq m_1 T + d + \sum_{i=2}^L (m_i T + 2m_{i-1} + 2 + d) + m_L + 2 \\ &\leq (T + 2) \sum_{i=1}^L m_i + L(d + 2) \\ &\leq 16d(n + 1)^d + L(d + 2). \end{aligned}$$

□

Proof of Corollary 5.1.1. Since the support of Φ in Theorem 4.1.1 is contained in $[-1/n, 1 + 1/n]^d$ which is just slightly larger than $[0, 1]^d$, the original construction of Φ could lead to a steep descent near the boundary of $[0, 1]^d$. To resolve this issue, we avoid taking values of Φ outside $[0, 1]^d$ by first projecting \mathbb{R}^d to $[0, 1]^d$ by applying

the following 1D projection coordinate-wise:

$$p(x) \stackrel{\text{def.}}{=} \sigma(x) - \sigma(x - 1) = \begin{cases} 0, & \text{if } x < 0 \\ x, & \text{if } x \in [0, 1], \\ 1, & \text{if } x > 1 \end{cases} \quad \forall x \in \mathbb{R} \quad (7.7.1)$$

For any $\mathbf{x}, \mathbf{y} \in \mathbb{R}^d$, $p(\mathbf{x}), p(\mathbf{y}) \in [0, 1]^d$, $\|p(\mathbf{x}) - p(\mathbf{y})\|_1 \leq \|\mathbf{x} - \mathbf{y}\|_1$, thus

$$|\Phi(p(\mathbf{x})) - \Phi(p(\mathbf{y}))| \leq \omega(\|p(\mathbf{x}) - p(\mathbf{y})\|_1) = \bar{\omega}(\|p(\mathbf{x}) - p(\mathbf{y})\|_1) \leq \bar{\omega}(\|\mathbf{x} - \mathbf{y}\|_1)$$

Therefore, $\bar{\omega}$ is a modulus of regularity of $\Phi \circ p$, whose restriction on $[0, 1]^d$ is the same as Φ . Thus, the function $\Phi \circ p$ satisfies our requirement, which is implemented by the following network:

$$\mathbf{x} \implies \begin{bmatrix} \sigma(\mathbf{x}) \\ \sigma(\mathbf{x} - 1) \end{bmatrix} \implies \begin{bmatrix} \Phi(\sigma(\mathbf{x}) - \sigma(\mathbf{x} - 1)) \end{bmatrix} = \Phi(p(\mathbf{x})) \quad (7.7.2)$$

which has width

$$\text{width}(\Phi \circ p) = \max\{2d, \text{width}(\Phi)\} \leq \max\{2d, 8d(n + 1)^d\} = 8d(n + 1)^d$$

and depth

$$\text{depth}(\Phi \circ p) = \text{depth}(\Phi) + 1 \leq \lceil \log_2 d \rceil + 5.$$

Finally, we tally the nonzero parameters in this network. From the proof of Theorem 4.1.1 (see the left-hand side of (7.6.10)), we see that the network Φ has at most $(16d - 2)(n + 1)^d$ nonzero parameters (we omitted the (-2) in the result to simplify

the expression). Now, in the original construction of Φ (see (7.6.5) and (7.6.8)), there are $2d(n+1)^d$ neurons in its first hidden layer, each connecting to $m = 1$ neuron in its input layer. Here, in the network (7.7.2) implementing $\Phi \circ p$, this number m is doubled, resulting in an additional

$$W_1 \stackrel{\text{def.}}{=} (2m - m) \cdot 2d(n+1)^d = 2d(n+1)^d$$

nonzero weights. For the biases, it is clear that only $B \stackrel{\text{def.}}{=} d$ additional biases (namely -1) were introduced in the first hidden layer. Lastly, the only additional nonzero parameters in (7.7.2) are the weights between the input layer and the first hidden layer, which are $W_2 \stackrel{\text{def.}}{=} 2d$ in total. Therefore, the number of nonzero parameters in (7.7.2) is at most

$$\begin{aligned} (16d - 2)(n+1)^d + W_1 + B + W_2 &= 18d(n+1)^d - 2(n+1)^d + 3d \\ &\leq 18d(n+1)^d - 2 \cdot 2^d + 3d \\ &\leq 18d(n+1)^d. \end{aligned}$$

□

Proof of Corollary 5.1.2. Fix $n \in \mathbb{N}_+$. First note that the metric projection onto the unit cube in \mathbb{R}^d , given for each $\mathbf{x} \in \mathbb{R}^d$ by $\Pi(\mathbf{x}) \in \operatorname{argmin}_{\mathbf{z} \in [0,1]^d} \|\mathbf{z} - \mathbf{x}\|_1$ is well-defined since $[0,1]^d$ is closed and convex. One can readily verify that $\Pi = p$, where p is defined in (7.7.1).

Now, by Corollary 5.1.1, there exists a ReLU MLP $\Phi : \mathbb{R}^d \rightarrow \mathbb{R}$ which is uniformly continuous with modulus of regularity $\bar{\omega}$, width at most $8d(n+1)^d$, and depth at most

$\lceil \log_2 d \rceil + 5$, satisfying the uniform estimate

$$\|f - \Phi\|_{L^\infty([0,1]^d)} \leq \bar{\omega} \left(\frac{d}{2n} \right). \quad (7.7.3)$$

Consequentially, (7.7.3) implies that: for each $\mathbf{x} \in \mathbb{R}^d$ the following holds

$$\begin{aligned} |f(\mathbf{x}) - \Phi(\mathbf{x})| &\leq |f(p(\mathbf{x})) - \Phi(p(\mathbf{x}))| + |f(p(\mathbf{x})) - f(\mathbf{x})| + |\Phi(p(\mathbf{x})) - \Phi(\mathbf{x})| \\ &= |f(p(\mathbf{x})) - \Phi(p(\mathbf{x}))| + |f(p(\mathbf{x})) - f(\mathbf{x})| + |\Phi(p(\mathbf{x})) - \Phi(p(\mathbf{x}))| \end{aligned} \quad (7.7.4)$$

$$\begin{aligned} &= |f(p(\mathbf{x})) - \Phi(p(\mathbf{x}))| + |f(p(\mathbf{x})) - f(\mathbf{x})| \\ &\leq |f(p(\mathbf{x})) - \Phi(p(\mathbf{x}))| + \bar{\omega}(\|p(\mathbf{x}) - \mathbf{x}\|_1) \\ &= |f(p(\mathbf{x})) - \Phi(p(\mathbf{x}))| + \bar{\omega}(\|p(\mathbf{x}) - \mathbf{x}\|_1) \\ &\leq \sup_{\mathbf{u} \in [0,1]^d} |f(\mathbf{u}) - \Phi(\mathbf{u})| + \bar{\omega}(\|p(\mathbf{x}) - \mathbf{x}\|_1) \\ &\leq \bar{\omega} \left(\frac{d}{2n} \right) + \bar{\omega}(\|p(\mathbf{x}) - \mathbf{x}\|_1), \end{aligned} \quad (7.7.5)$$

where (7.7.4) followed by the proof of by Corollary 5.1.1, specifically in (7.7.2), we see that the first layer of Φ is given by pre-composition with p .

Since $[0, 1]^d$ is compact then the minimal distance to any given point is realized. Also, since p is the metric projection of \mathbb{R}^d onto $[0, 1]^d$, then

$$d_{[0,1]^d}(\mathbf{x}) = \|p(\mathbf{x}) - \mathbf{x}\|_1$$

for every $\mathbf{x} \in \mathbb{R}^d$. Consequentially, the right-hand side of (7.7.5) can be re-expressed

as

$$\bar{\omega}\left(\frac{d}{2n}\right) + \bar{\omega}(\|p(\mathbf{x}) - \mathbf{x}\|_1) = \bar{\omega}\left(\frac{d}{2n}\right) + \bar{\omega}(\text{dist}_{[0,1]^d}(\mathbf{x})),$$

which concludes our proof. \square

7.8 Proof of Statistical Result - Theorem 4.2.1

Let \mathcal{F} be a non-empty set of functions from \mathbb{R}^d to \mathbb{R} , Z_1, \dots, Z_N be random vectors in \mathbb{R}^d , and let $\mathcal{Z} \stackrel{\text{def.}}{=} (Z_n)_{n=1}^N$. Its *empirical Rademacher complexity*, denoted by $\mathfrak{R}_{\mathcal{Z}}(\mathcal{F})$, is the (random) quantity

$$\mathfrak{R}_{\mathcal{Z}}(\mathcal{F}) \stackrel{\text{def.}}{=} \frac{1}{N} \mathbb{E}_{\sigma} \left[\sup_{g \in \mathcal{F}} \sum_{n=1}^N \sigma_n g(Z_n) \right] \quad (7.8.1)$$

where $\sigma = (\sigma_n)_{n=1}^N$ and $\sigma_1, \dots, \sigma_N$ are i.i.d. Rademacher random variables; i.e. $\mathbb{P}(\sigma_1 = 1) = \mathbb{P}(\sigma_1 = -1) = 1/2$.

Proof of Theorem 4.2.1. Let $\delta_1, \delta_2 \in [0, 1)$, to be fixed retroactively. Let $\mathcal{Z} \stackrel{\text{def.}}{=} \{(X_n, Y_n)\}_{n=1}^N$. By [8, Theorems 8 and 12], the following holds with probability at-least $1 - \delta_1$

$$\sup_{\Phi \in \mathcal{NN}_{\Delta, W}^{L, n}} |\mathcal{R}(\Phi) - \mathcal{R}_N(\Phi)| \leq 2L_{\ell} \mathfrak{R}_{\mathcal{Z}}(\mathcal{NN}_{\Delta, W}^{L, n}) + \frac{\sqrt{8 \log(2/\delta_1)}}{\sqrt{N}}. \quad (7.8.2)$$

Since $\mathcal{NN}_{\Delta, W}^{L, n} \subset \mathcal{NN}_{\Delta, W} \cap \text{Lip}(\mathbb{R}^d, [0, 1], L)$ then, by the definition of the empirical Rademacher complexity of a class in (7.8.1), we have that

$$\mathfrak{R}_{\mathcal{Z}}(\mathcal{NN}_{\Delta, W}^{L, n}) \leq \min \left\{ \underbrace{\mathfrak{R}_{\mathcal{Z}}(\mathcal{NN}_{\Delta, W})}_{\text{(I)}}, \underbrace{\mathfrak{R}_{\mathcal{Z}}(\text{Lip}(\mathbb{R}^d, [0, 1], L))}_{\text{(II)}} \right\}. \quad (7.8.3)$$

Step 1 - Bounding Term (I):

Let \mathbf{X} be the $N \times d$ random matrix with rows given by X_1, \dots, X_N . Term (I) can be bounded using the empirical Rademacher complexity bound derived in [6, Theorem 3.3] which implies that

$$(I) \leq \frac{4}{N^{3/2}} + \frac{26 \log(N) \log(2W)}{N} \left(\|\mathbf{X}\|_F \prod_{l=1}^{\Delta} s_l \left(\sum_{l=1}^{\Delta} \left(\frac{b_l}{s_l} \right)^{2/3} \right)^{3/2} \right) \quad (7.8.4)$$

where $\|\cdot\|_F$ and $\|\cdot\|_2$ respectively denote the Fröbenius and the spectral matrix norms, and s_l denotes the maximum spectral norm of the l^{th} matrix $\mathbf{W}^{(l)}$ for any ReLU MLP Φ with representation (3.3.1), and b_l denotes the maximum $\|\cdot\|_{1,2}$ matrix norm thereof (defined by the sum of the Euclidean norm of its columns³).

Recall that, if $\Phi \in \mathcal{NN}_{\Delta, W}^{L, n}$ then by Proposition 4.1.1 the entries of all its weight matrices have absolute value at-most $1/2$ if $l \in \{2, \dots, \Delta-1\}$ and $\max\{n, 1\}$ otherwise. These observations, together with the elementary matrix norm inequalities (all found on [22, page 56]), and the fact that each $\Phi \in \mathcal{NN}_{\Delta, W}$ has width at-most W , and that by enlarging the class a bit, we may assume that $s_1 = s_{\Delta} = n W^{3/2}$ and that $b_l = s_l = \frac{1}{2} W^{3/2}$ for $l = 1, \dots, \Delta$. Thus, the right-hand side of (7.8.4) simplifies to

$$(I) \leq \frac{4}{N^{3/2}} + \frac{26 \log(N) \log(2W)}{2^{\max\{0, \Delta-2\}} N} \|\mathbf{X}\|_F W^{3\Delta/2}. \quad (7.8.5)$$

Using the standard matrix-norm bounds $\|\cdot\|_F \leq \sqrt{N} \|\cdot\|_{\text{op}}$ (where $\|\cdot\|_{\text{op}}$ denotes the operator norm) and the fact that the operator norm of a matrix \mathbf{A} equals to its

³The authors take a transpose in their main result; since this matrix norm is typically defined by summing over the Euclidean norms of the rows of a matrix.

largest singular value $\sigma_{max}(\mathbf{A})$, we reduce the right-hand side of (7.8.5) to

$$\begin{aligned}
\text{(I)} &\leq \frac{4}{N^{3/2}} + \frac{26 \log(N) \log(2W)}{2^{\max\{0, \Delta-2\}} N} \sqrt{N} \|\mathbf{X}\|_{\text{op}} W^{3\Delta/2} \\
&= \frac{4}{N^{3/2}} + \frac{26 \log(N) \log(2W)}{2^{\max\{0, \Delta-2\}} \sqrt{N}} \sigma_{max}(\mathbf{X}) W^{3\Delta/2}.
\end{aligned} \tag{7.8.6}$$

Define $\bar{\mathbf{X}} \stackrel{\text{def.}}{=} \sqrt{N} \cdot \mathbf{X}$. By the min-max characterization of singular values, see e.g. [31, Theorem 4.2.11], we have that $\sigma_{max}(\bar{\mathbf{X}}) = \sqrt{N} \sigma_{max}(\mathbf{X})$. Observe also that $\bar{\mathbf{X}}$ is isotropic since

$$\mathbb{E}[\bar{\mathbf{X}}\bar{\mathbf{X}}^\top] = \mathbb{E}[(\sqrt{N})\mathbf{X}(\sqrt{N}\mathbf{X})^\top] = N\mathbb{E}[\mathbf{X}\mathbf{X}^\top] = N\frac{1}{N}\mathbf{I}_d = \mathbf{I}_d.$$

Therefore, a consequence to a version of Gordon’s Theorem given in [79, Theorem 4.6.1] applies to the random matrix $\bar{\mathbf{X}}$ since it has independent, sub-Gaussian, and isotropic rows; thus: with probability at-least $1 - \delta_2$ we have that

$$\sigma_{max}(\bar{\mathbf{X}}) \leq \sqrt{d} + C(\sqrt{N} + \sqrt{\ln(2/\delta_2)}).$$

Therefore, with probability at-least $1 - \delta_2$, the maximal singular value of \mathbf{X} is bounded above by ⁴

$$\sigma_{max}(\mathbf{X}) \leq \frac{1}{\sqrt{N}}(\sqrt{d} + C(\sqrt{N} + \sqrt{\ln(2/\delta_2)})) = \frac{\sqrt{d}}{\sqrt{N}} + C\frac{\sqrt{\ln(2/\delta_2)}}{\sqrt{N}} + C. \tag{7.8.7}$$

Incorporating (7.8.7) into the right-hand side of (7.8.6), implies that: the following

⁴Remark that: without the correct scaling of $\frac{1}{N} I_d$ the covariance of each X_1 the third term on the right-hand side of (7.8.7) would not tend to 0 as N tends to infinity.

holds with probability at least $1 - \delta_2$

$$\begin{aligned}
\text{(I)} &\leq \frac{4}{N^{3/2}} + \frac{26 \log(N) \log(2W)}{2^{\max\{0, \Delta-2\}} \sqrt{N}} \sigma_{\max}(\mathbf{X}) W^{3\Delta/2} \\
&\leq \frac{4}{N^{3/2}} + \frac{26 \log(N) \log(2W) W^{3\Delta/2}}{2^{\max\{0, \Delta-2\}} N} (\sqrt{d} + C \sqrt{N} + C \sqrt{\ln(2/\delta_2)}).
\end{aligned} \tag{7.8.8}$$

Fix $\delta \in (0, 1)$ and set $\delta_1 \stackrel{\text{def.}}{=} \delta_2 \stackrel{\text{def.}}{=} \delta/2$. By a union bound together with (7.8.8) and (7.8.3) we deduce that: the following holds with probability at-least $1 - \delta$

$$\begin{aligned}
\sup_{\Phi \in \mathcal{NN}_{\Delta, W}^{L, n}} |\mathcal{R}(\Phi) - \mathcal{R}_N(\Phi)| &\leq 2 L_\ell \mathfrak{R}_{\mathcal{Z}}(\mathcal{NN}_{\Delta, W}^{L, n}) + \frac{\sqrt{8 \log(4/\delta)}}{\sqrt{N}} \\
\mathfrak{R}_{\mathcal{Z}}(\mathcal{NN}_{\Delta, W}^{L, n}) &\leq \min \left\{ \frac{4}{N^{3/2}} + \frac{26 \log(N) \log(2W) W^{3\Delta/2}}{2^{\max\{0, \Delta-2\}} N} (\sqrt{d} + C \sqrt{N} + C \sqrt{\ln(4/\delta)}), \right. \\
&\quad \left. \underbrace{\mathfrak{R}_{\mathcal{Z}}(\text{Lip}(\mathbb{R}^d, [0, 1], L))}_{\text{(III)}} \right\}.
\end{aligned} \tag{7.8.9}$$

Step 2 - Bounding Term (II):

By [33, Lemma 25], we have that

$$\text{(III)} \leq C_d \frac{L^{\frac{d}{d+3}}}{N^{\frac{1}{d+3}}} = \left((8(d+1)^2 16^d)^{\frac{1}{d+3}} + 4\sqrt{2} \frac{16^{\frac{d}{d+3}}}{(18(d+1))^{\frac{d+1}{d+3}}} \right) \frac{L^{\frac{d}{d+3}}}{N^{\frac{1}{d+3}}} \tag{7.8.10}$$

where $C_d \stackrel{\text{def.}}{=} (8(d+1)^2 16^d)^{\frac{1}{d+3}} + 4\sqrt{2} \frac{16^{\frac{d}{d+3}}}{(18(d+1))^{\frac{d+1}{d+3}}}$. Incorporating (7.8.10) into (7.8.9) yields: the following holds with probability at-least $1 - \delta$

$$\begin{aligned}
\sup_{\Phi \in \mathcal{NN}_{\Delta, W}^{L, n}} |\mathcal{R}(\Phi) - \mathcal{R}_N(\Phi)| &\leq 2 L_\ell \mathfrak{R}_{\mathcal{Z}}(\mathcal{NN}_{\Delta, W}^{L, n}) + \frac{\sqrt{8 \log(4/\delta)}}{\sqrt{N}} \\
\mathfrak{R}_{\mathcal{Z}}(\mathcal{NN}_{\Delta, W}^{L, n}) &\leq \min \left\{ \frac{4}{N^{3/2}} + \frac{26 \log(N) \log(2W) W^{3\Delta/2}}{2^{\max\{0, \Delta-2\}} N} (\sqrt{d} + C \sqrt{N} + C \sqrt{\ln(4/\delta)}), C_d \frac{L^{\frac{d}{d+3}}}{N^{\frac{1}{d+3}}} \right\}.
\end{aligned}$$

□

7.9 Proof of Results in the Discussion Section

Proof of Theorem 6.1.1. We will prove the following stronger statement: in the context of Lemma 7.5.4, for $n = 1$, the Kuhn triangulation is the only triangulation (up to reflections) such that for all Lipschitz functions f with modulus of regularity ω which is a linear function, ω is also a modulus of regularity of Φ . By a reflection, we mean a transformation of the form $x_i \mapsto 1 - x_i$, for some $i = 1, \dots, d$. In other words, we will only consider Lipschitz functions and linear moduli of regularity.

Note that, all reflections are involutions and they are commutative with each other, so we will refer to a transformation as a reflection of the axes x_{i_1}, \dots, x_{i_m} if it is a composition of the reflections $x_j \mapsto 1 - x_j$ for $j = i_1, \dots, i_m$, in any order; in particular, we say that it reflects the x_i -axis if $i \in \{i_1, \dots, i_m\}$.

For $d = 1, 2$, the Kuhn triangulation is the only triangulation of $[0, 1]^d$ up to reflections, so we will assume that $d \geq 3$ throughout the rest of the proof.

Let $S \stackrel{\text{def.}}{=} \{S_k\}_{k=1}^m$ be any triangulation of $[0, 1]^d$ that makes Lemma 7.5.4 true. We will prove that S is some reflected version of the Kuhn triangulation.

For $i = 1, \dots, d$ and $\delta = 0, 1$, let

$$F_{i,\delta} \stackrel{\text{def.}}{=} [0, 1]^d \cap \{(x_1, \dots, x_d) \in \mathbb{R}^d : x_i = \delta\}$$

be some face of the unit cube $[0, 1]^d$, and

$$S_{i,\delta} \stackrel{\text{def.}}{=} \{S_k \cap F_{i,\delta} : S_k \cap F_{i,\delta} \text{ is a } (d-1)\text{-dimensional simplex, } k = 1, \dots, m\}$$

be the “restricted triangulation” of S on $F_{i,\delta}$. For any such $(d - 1)$ -dimensional face $F_{i,\delta}$, we show that the finite collection $S_{i,\delta}$ of $(d - 1)$ -dimensional simplices is a triangulation of $F_{i,\delta}$:

1. $\cup S_{i,\delta} = F_{i,\delta}$: assume for contradiction that $\cup S_{i,\delta}$ is a proper subset of $F_{i,\delta}$. Denote $F_{i,\delta}^o$ as the interior of $F_{i,\delta}$, then $F_{i,\delta}^o \setminus (\cup S_{i,\delta})$ is open and non-empty, since otherwise $(\cup S_{i,\delta}) \supseteq F_{i,\delta}^o$, which implies that $(\cup S_{i,\delta}) \supseteq F_{i,\delta}$ as $\cup S_{i,\delta}$ is a closed set. Consider the Lebesgue measure on the $(d - 1)$ -dimensional hyperplane $x_i = \delta$ that contains $F_{i,\delta}$. The measure of $F_{i,\delta}^o \setminus (\cup S_{i,\delta})$ is positive as it is non-empty and open. If $S_k \cap F_{i,\delta}$ is not $(d - 1)$ -dimensional, then it is at most $(d - 2)$ -dimensional, which has measure 0. There are only a finite number of these $(d - 2)$ -dimensional faces with 0 measure, so their total measure is 0, which cannot fill up $F_{i,\delta}^o \setminus (\cup S_{i,\delta})$ as it has positive measure, so we get a contradiction as desired;
2. The set of vertices of the simplices in $S_{i,\delta}$ is $\{0, 1\}^d \cap F_{i,\delta}$, which is the set of vertices of $F_{i,\delta}$;
3. The interiors of the simplices in $S_{i,\delta}$ are pairwise disjoint: without loss of generality, consider the case $i = 1$ and $\delta = 0$. Assume for contradiction that \mathbf{z} is contained in the interiors of S_p and S_q , for two different simplices in $S_{1,0}$. Assume $S_p = S_{p'} \cap F_{1,0}$ for some $S_{p'} \in S$, then compared to S_p , $S_{p'}$ has an additional vertex in the “upper face” $F_{1,1}$ (otherwise $S_{p'}$ would not be a simplex in \mathbb{R}^d as all its vertices lie in the $(d - 1)$ -dimensional hyperplane $x_1 = 0$), so $\mathbf{z} + t\mathbf{e}_1$ is contained in the interior of $S_{p'}$, for $t > 0$ small enough. Similarly, assume $S_q = S_{q'} \cap F_{1,0}$ for some $S_{q'} \in S$, then $\mathbf{z} + t\mathbf{e}_1$ is also contained in the interior of

S_{q^t} for $t > 0$ small enough, which is impossible as S is a triangulation and its simplices has pairwise disjoint interiors, so we get a contradiction as desired.

Therefore, $S_{i,\delta}$ is indeed a triangulation of $F_{i,\delta}$, which is a translated and reflected $(d-1)$ -dimensional unit cube $[0, 1]^{d-1}$.

Assume for contradiction that $S_{i,\delta}$ is not a reflected version of the Kuhn triangulation. Without loss of generality, consider the case $i = 1$ and $\delta = 0$. By induction, there exists a function $f : F_{1,0} \rightarrow \mathbb{R}$ and a linear modulus of regularity ω_f of f which is not a modulus of regularity of the continuous piecewise linear approximator $\Phi : F_{1,0} \rightarrow \mathbb{R}$ of f with respect to the triangulation $S_{1,0}$. Consider the extension f^* of f to $[0, 1]^d$ defined by:

$$f^*(x_1, x_2, x_3, \dots, x_d) \stackrel{\text{def.}}{=} f(0, x_2, x_3, \dots, x_d), \quad \forall (x_1, x_2, x_3, \dots, x_d) \in [0, 1]^d$$

Clearly ω_f is also a modulus of regularity of f^* as it is non-decreasing. Let Φ^* be the continuous piecewise linear approximator of f on $[0, 1]^d$ with respect to the triangulation S , then the restriction of Φ^* on $F_{1,0}$ is Φ , but this implies that ω_f is not a modulus of regularity of Φ^* , as this is not so even just on the subset $F_{1,0}$ of $[0, 1]^d$. Thus, we get a contradiction as desired.

Therefore, $S_{i,\delta}$ is indeed some reflected version of the Kuhn triangulation, for $i = 1, \dots, d$ and $\delta = 0, 1$. Since $F_{d,0}$ is the unit cube $[0, 1]^{d-1}$ in the $(d-1)$ -dimensional hyperplane $x_d = 0$, without loss of generality, we may assume that $S_{d,0}$ is the “original” Kuhn triangulation of $F_{d,0}$ without any reflections. Then, by definition of the Kuhn triangulation, any simplex in $S_{d,0}$ has an edge connecting $\mathbf{u}_1 \stackrel{\text{def.}}{=} (0, 0, \dots, 0, 0)$ and $\mathbf{u}_2 \stackrel{\text{def.}}{=} (1, 1, \dots, 1, 0)$. Since $S_{d,0}$ is the “restricted” triangulation of S on $F_{d,0}$, there is some corresponding simplex in S that has an edge connecting \mathbf{u}_1 and \mathbf{u}_2 .

Next, we identify the “separating hyperplanes” (i.e. each simplex lies on the same side of it) in S and show that they are the same as the ones given in Lemma 7.5.1. We first show that the hyperplanes $x_i - x_j = 0$ for all $1 \leq i < j \leq d-1$ are separating hyperplanes of S .

Assume for contradiction that there is some $1 \leq i < j \leq d-1$ such that the hyperplane $x_i - x_j = 0$ is not a separating hyperplane. Without loss of generality, assume that $i = 1, j = 2$. Then, there is some simplex in S that has two vertices $\mathbf{v}_1, \mathbf{v}_2$ on different sides of the hyperplane $x_1 - x_2 = 0$, so it has an edge connecting $\mathbf{v}_1, \mathbf{v}_2$ as it is a simplex. Since $\mathbf{v}_1, \mathbf{v}_2 \in \{0, 1\}^d$ and they lie on different sides of the hyperplane $x_1 - x_2 = 0$, $\mathbf{v}_1 = (0, 1, \delta_3^{(1)}, \dots, \delta_d^{(1)})$, $\mathbf{v}_2 = (1, 0, \delta_3^{(2)}, \dots, \delta_d^{(2)})$, where the $\delta_a^{(b)}$'s are either 0 or 1.

Now, consider the following “counterexample” function:

$$f(x_1, x_2, \dots, x_d) \stackrel{\text{def.}}{=} \sigma(x_1 - x_2) + \sum_{i=3}^d x_i, \quad \forall (x_1, x_2, \dots, x_d) \in [0, 1]^d$$

Clearly f is Lipschitz continuous with Lipschitz constant 1 (under ℓ^1 norm of \mathbb{R}^d), so $\omega(x) \stackrel{\text{def.}}{=} x (\forall x \geq 0)$ is a modulus of regularity of f . However, for the continuous piecewise linear approximator Φ of f on $[0, 1]^d$ with respect to the triangulation S : consider the pair of points

$$\mathbf{u} \stackrel{\text{def.}}{=} \frac{\mathbf{u}_1 + \mathbf{u}_2}{2} = \left(\frac{1}{2}, \frac{1}{2}, 0, 0, \dots, 0, 0 \right)$$

$$\mathbf{v} \stackrel{\text{def.}}{=} \frac{\mathbf{v}_1 + \mathbf{v}_2}{2} = \left(\frac{1}{2}, \frac{1}{2}, \frac{\delta_3^{(1)} + \delta_3^{(2)}}{2}, \dots, \frac{\delta_{d-1}^{(1)} + \delta_{d-1}^{(2)}}{2}, \frac{\delta_d^{(1)} + \delta_d^{(2)}}{2} \right)$$

Since there is some simplex in S that has an edge connecting \mathbf{u}_1 and \mathbf{u}_2 ,

$$\Phi(\mathbf{u}) = \Phi\left(\frac{\mathbf{u}_1 + \mathbf{u}_2}{2}\right) = \frac{\Phi(\mathbf{u}_1) + \Phi(\mathbf{u}_2)}{2} = \frac{f(\mathbf{u}_1) + f(\mathbf{u}_2)}{2} = 0$$

Since there is some simplex in S that has an edge connecting \mathbf{v}_1 and \mathbf{v}_2 ,

$$\Phi(\mathbf{v}) = \Phi\left(\frac{\mathbf{v}_1 + \mathbf{v}_2}{2}\right) = \frac{\Phi(\mathbf{v}_1) + \Phi(\mathbf{v}_2)}{2} = \frac{f(\mathbf{v}_1) + f(\mathbf{v}_2)}{2} = \frac{1}{2} + \sum_{i=3}^d \frac{\delta_i^{(1)} + \delta_i^{(2)}}{2}$$

Then, ω is not a modulus of regularity of Φ because

$$|\Phi(\mathbf{u}) - \Phi(\mathbf{v})| = \frac{1}{2} + \sum_{i=3}^d \frac{\delta_i^{(1)} + \delta_i^{(2)}}{2} > \sum_{i=3}^d \frac{\delta_i^{(1)} + \delta_i^{(2)}}{2} = \omega(\|\mathbf{u} - \mathbf{v}\|_1)$$

This contradicts our assumption that the continuous piecewise linear approximators of any Lipschitz continuous functions on $[0, 1]^d$ with respect to the triangulation S always preserve regularity under ℓ^1 norm of \mathbb{R}^d . Thus, we get a contradiction as desired.

Therefore, the hyperplanes $x_i - x_j = 0$ for all $1 \leq i < j \leq d - 1$ are indeed separating hyperplanes of S .

By similar arguments as above, switching to other “pivot axis” than x_d , we can conclude that for any $1 \leq i \leq j \leq d$, either $x_i - x_j = 0$ or $x_i + x_j = 1$ is a separating hyperplane of S , taking possible reflections into account. We have shown above that for all $1 \leq i \leq j < d$, the separating hyperplane is of the form $x_i - x_j = 0$. Thus, we remain to show that all remaining hyperplanes either all have the form $x_i - x_d = 0$ or all have the form $x_i + x_d = 1$ for $1 \leq i \leq d - 1$, and then the result follows from Lemma 7.5.1.

Assume for contradiction that there are some distinct i, j with $1 \leq i, j \leq d - 1$ such that both $x_i - x_d = 0$ and $x_j + x_d = 1$ are separating hyperplanes of S . Without loss of generality, assume that $i = 1, j = 2$. Then, consider the following d separating hyperplanes of S :

$$x_1 - x_d = 0, \quad x_2 + x_d = 1, \quad x_1 - x_2 = 0, \quad x_k = 0 \text{ for } 3 \leq k \leq d - 1$$

They intersect at a single point $(\frac{1}{2}, \frac{1}{2}, 0, 0, \dots, 0, \frac{1}{2})$, which must be a vertex of some simplex in S , but this is impossible since it is not a lattice point. Therefore, we get a contradiction as desired. \square

Proof of Theorem 6.2.1. The proof is a mild modification of the one in [70], which approximates the target function by assigning constant values (sample values) on some pairwise disjoint cubes inside $[0, 1]^d$, and the rest of $[0, 1]^d$ is called “trifling regions” where errors are controlled by making these regions extremely small and taking the median of 3^d such MLPs with different trifling regions. The problem with this approach is that the Lipschitz constant explodes in these extremely small trifling regions. Here, we partially resolve this issue by taking the median of only $2d + 1$ of these MLPs, thus enabling these trifling regions to become as large as possible, which leads to a lower Lipschitz constant.

For the case $d = 1$: Let $x_i = \frac{i}{n^2}$ for $i = 0, 1, \dots, n^2$. Since

$$\left\lfloor \frac{2n + 2}{4} \right\rfloor = \left\lfloor \frac{n + 1}{2} \right\rfloor \geq \frac{n + 1}{2} - \frac{1}{2} = \frac{n}{2}$$

then by Proposition 7.4.1 (taking $M = 2n + 1, N = \lfloor \frac{2n+2}{4} \rfloor$), given the set of no more than MN samples $\{(x_i, f(x_i)) : i = 0, 1, \dots, n^2\}$, there exists $\Phi \in \mathcal{NN}(\#input =$

1; $\text{widthvec} = [2n + 1, 2n]$) such that $\Phi(x_i) = f(x_i)$ for $i = 0, 1, \dots, n^2$, and Φ is linear on $[x_i, x_{i+1}]$ for $i = 0, 1, \dots, n^2 - 1$. Then, the result follows from Lemma 7.3.4.

Therefore, in the rest of this proof, we will assume that $d \geq 2$. We will also assume that f is non-negative, since otherwise we can first construct Φ for the non-negative function $f - \min f$ and then add back $\min f$ (which exists since f is a continuous function on a compact domain) in the output layer.

For $k = 1, 2, \dots, 2d + 1$ we define some modifications of these “trifling region” in [70]:

$$\Omega_k \stackrel{\text{def.}}{=} \bigcup_{i=1}^d \left\{ \mathbf{x} = (x_1, \dots, x_d) \in [0, 1]^d : x_i \in \bigcup_{j=0}^{n-1} \left(\frac{1}{n} \left(\frac{k-1}{2d+1} + j \right), \frac{1}{n} \left(\frac{k}{2d+1} + j \right) \right) \right\} \quad (7.9.1)$$

Ω_k is the union of nd “thick hyperplanes”, where for each of the d axes, there are n of these parallel and equispaced “thick hyperplanes” that are perpendicular to it. Then, Ω_k separates $[0, 1]^d$ into n^d (if $k = 1$ or $2d + 1$) or $(n + 1)^d$ (if $2 \leq k \leq 2d$) non-overlapping equispaced cuboids:

$$Q_k \stackrel{\text{def.}}{=} [0, 1]^d \setminus \Omega_k = \left[[0, 1] \setminus \bigcup_{j=0}^{n-1} \left(\frac{1}{n} \left(\frac{k-1}{2d+1} + j \right), \frac{1}{n} \left(\frac{k}{2d+1} + j \right) \right) \right]^d \quad (7.9.2)$$

Most of these cuboids are d -dimensional cubes, except for the ones that touches some face of $[0, 1]^d$ for $2 \leq k \leq 2d$. See Figure 7.4 for an illustration of $\Omega_1, \dots, \Omega_5$ and Q_1, \dots, Q_5 for the case $n = d = 2$.

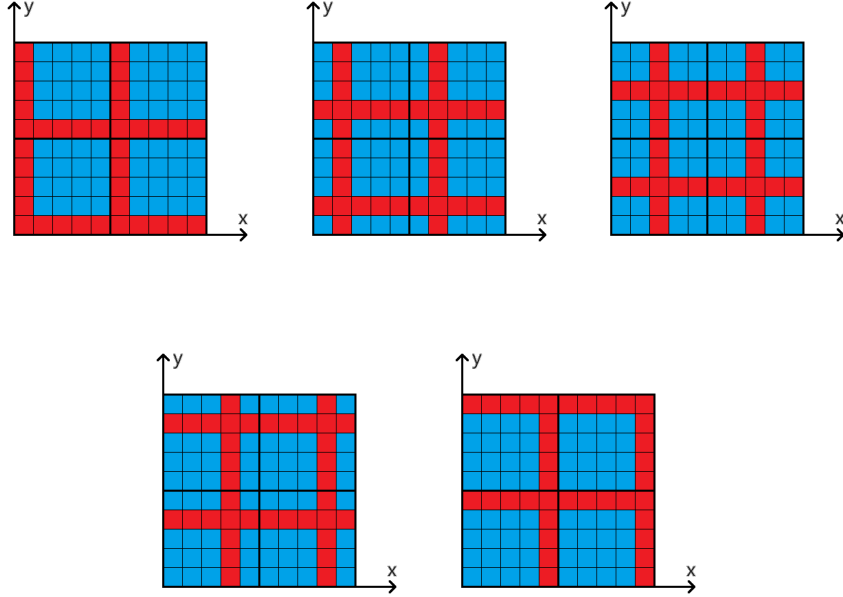


Figure 7.4: An illustration of the trifling regions $\Omega_1, \dots, \Omega_5$ (colored in red) and Q_1, \dots, Q_5 (colored in blue) for the case $n = d = 2$. Each of the smallest cubes appears in blue regions (where approximation errors are controlled) for at least 3 times. By Lemma 7.3.1, the errors of their medians are also controlled.

Step 1: Construct the projection maps $\pi_k : [0, 1]^d \rightarrow \mathbb{R}$ for each Q_k

Fix $2 \leq k \leq 2d$. The cases $k = 1$ or $2d + 1$ can be treated similarly and require less neurons, but we will assume that they require the same number of neurons as other cases for convenience. We will define π_k in such a way that it is constant on each cuboid of Q_k , and different cuboids correspond to different constants. Notice that

$$[0, 1] \setminus \bigcup_{j=0}^{n-1} \left(\frac{1}{n} \left(\frac{k-1}{2d+1} + j \right), \frac{1}{n} \left(\frac{k}{2d+1} + j \right) \right)$$

is the union of $n + 1$ equi-spaced closed intervals, so we define a map $p_k : [0, 1] \rightarrow \mathbb{R}$ which maps the first of them (counting from the left) to 0, the second of them to 1, \dots , the last one of them to n , and let p_k be linear on each of the remaining intervals such

that it is continuous. Then, p_k has $2n + 2$ break points (including 0 and 1), so by [69, Lemma 2.1], p_k can be implemented by a ReLU MLP with width $2n + 1$ and depth 1. Then, for $\mathbf{x} = (x_1, \dots, x_d) \in \mathbb{R}^d$, we define

$$\pi_k(\mathbf{x}) = \sum_{i=1}^d (n + 1)^{i-1} p_k(x_i)$$

which can be seen as the number $\overline{(p_k(x_d)p_k(x_{d-1}) \cdots p_k(x_1))}$ in base $n + 1$, so it is easy to see that π_k indeed maps each of the $(n + 1)^d$ cuboids to different constants, namely $0, 1, 2, \dots, (n + 1)^d - 1$. Then, we implement π_k as follows:

$$\begin{bmatrix} x_1 \\ x_2 \\ \vdots \\ x_d \end{bmatrix} \implies \begin{bmatrix} p_k(x_1) \\ p_k(x_2) \\ \vdots \\ p_k(x_d) \end{bmatrix} \implies \sum_{i=1}^d (n + 1)^{i-1} p_k(x_i)$$

so π_k has width $(2n + 1)d$ and depth 3 (counting the output layer, which is still a hidden layer of the final network).

Step 2: Construct a 1-dimensional memorizer $\phi_k : \mathbb{R} \rightarrow \mathbb{R}$ for each Q_k

Fix $2 \leq k \leq 2d$. As in Step 1, the cases $k = 1$ or $2d + 1$ can be treated similarly and require less neurons, but we will assume that they require the same number of neurons as other cases for convenience.

For $j = 0, 1, 2, \dots, (n + 1)^d - 1$, let y_j be the value of f at the center of the cuboid that is mapped to j under π_k . Note that the range of π_k is $[0, (n + 1)^d - 1]$, so we don't need to care about the behavior of our memorizer outside this interval.

Let $\{(j, y_j) : j = 0, 1, 2, \dots, (n + 1)^d - 1\}$ be a sample set of size $(n + 1)^d$ for our

memorizer. Since

$$\left\lfloor \frac{2\lceil(n+1)^{d/2}\rceil + 2}{4} \right\rfloor = \left\lfloor \frac{\lceil(n+1)^{d/2}\rceil + 1}{2} \right\rfloor \geq \frac{\lceil(n+1)^{d/2}\rceil + 1}{2} - \frac{1}{2} \geq \frac{(n+1)^{d/2}}{2}$$

then by Proposition 7.4.1 (taking $M = 2\lceil(n+1)^{d/2}\rceil$, $N = \lfloor \frac{2\lceil(n+1)^{d/2}\rceil + 2}{4} \rfloor$), there is a ReLU MLP ϕ_k with width $2\lceil(n+1)^{d/2}\rceil$ and depth 3 (counting the output layer, which is still a hidden layer of the final network) that memorizes this sample set, and ϕ_k is linear on $[j, j+1]$ for $j = 0, 1, \dots, n^d - 2$.

Step 3: Combine Step 1, Step 2 and Lemma 7.3.2 to construct the final network Φ

We denote the ReLU MLP implementing the median function on $2d + 1$ non-negative arguments as \mathbf{M}_{2d+1} , which has width $6d + 3$ and depth $11d + 3$ by Lemma

7.3.2, then we construct Φ as follows:

$$\begin{aligned}
 \mathbf{x} &\Rightarrow \begin{bmatrix} \mathbf{x} \\ \pi_1(\mathbf{x}) \end{bmatrix} \Rightarrow \begin{bmatrix} \mathbf{x} \\ \pi_1(\mathbf{x}) \\ \pi_2(\mathbf{x}) \end{bmatrix} \Rightarrow \cdots \Rightarrow \begin{bmatrix} \mathbf{x} \\ \pi_1(\mathbf{x}) \\ \pi_2(\mathbf{x}) \\ \vdots \\ \pi_{2d}(\mathbf{x}) \end{bmatrix} \Rightarrow \begin{bmatrix} \pi_1(\mathbf{x}) \\ \pi_2(\mathbf{x}) \\ \vdots \\ \pi_{2d}(\mathbf{x}) \\ \pi_{2d+1}(\mathbf{x}) \end{bmatrix} \\
 &\Rightarrow \begin{bmatrix} \phi_1 \circ \pi_1(\mathbf{x}) \\ \pi_2(\mathbf{x}) \\ \vdots \\ \pi_{2d+1}(\mathbf{x}) \end{bmatrix} \Rightarrow \begin{bmatrix} \phi_1 \circ \pi_1(\mathbf{x}) \\ \phi_2 \circ \pi_2(\mathbf{x}) \\ \vdots \\ \pi_{2d+1}(\mathbf{x}) \end{bmatrix} \Rightarrow \cdots \Rightarrow \begin{bmatrix} \phi_1 \circ \pi_1(\mathbf{x}) \\ \phi_2 \circ \pi_2(\mathbf{x}) \\ \vdots \\ \phi_{2d+1} \circ \pi_{2d+1}(\mathbf{x}) \end{bmatrix} \\
 &\Rightarrow \mathbf{M}_{2d+1} \left(\begin{bmatrix} \phi_1 \circ \pi_1(\mathbf{x}) \\ \phi_2 \circ \pi_2(\mathbf{x}) \\ \vdots \\ \phi_{2d+1} \circ \pi_{2d+1}(\mathbf{x}) \end{bmatrix} \right) \Rightarrow \Phi(\mathbf{x})
 \end{aligned}$$

Its width and depth are

$$\begin{aligned}
 \text{width}(\Phi) &= \max \{ \text{width}(\pi_{2d+1}) + 2d, \text{width}(\phi_{2d+1}) + 2d, \text{width}(M_{2d+1}) \} \\
 &= \max \{ (2n + 3)d, 6d + 3, 2\lceil (n + 1)^{d/2} \rceil + 2d \}
 \end{aligned}$$

$$\begin{aligned}
\text{depth}(\Phi) &= \sum_{k=1}^{2d+1} \text{depth}(\pi_k) + \sum_{k=1}^{2d+1} \text{depth}(\phi_k) + \text{depth}(M_{2d+1}) \\
&= 3(2d+1) + 3(2d+1) + (11d+3) \\
&= 23d+9
\end{aligned}$$

Step 4: Estimate the L^∞ error

We first show that the L^∞ error of $\phi_k \circ \pi_k$ is controlled on Q_k for all k : pick any $\mathbf{w} \in Q_k$, then since π_k maps the cuboid in Q_k containing \mathbf{w} (whose center we denote by \mathbf{w}_0) to some constant j , which is further mapped by ϕ_k to $y_j \stackrel{\text{def.}}{=} f(\mathbf{w}_0)$, we have $\phi_k \circ \pi_k(\mathbf{w}) = f(\mathbf{w}_0)$, so

$$|\phi_k \circ \pi_k(\mathbf{w}) - f(\mathbf{w})| = |f(\mathbf{w}_0) - f(\mathbf{w})| \leq \nu \|\mathbf{w}_0 - \mathbf{w}\|_1^\alpha$$

Since the edge length of any cuboid in Q_k is at most $l \stackrel{\text{def.}}{=} \frac{2d}{n(2d+1)}$ (this can be seen from Equation 7.9.2), $\|\mathbf{w}_0 - \mathbf{w}\|_1 \leq \frac{d}{2}l = \frac{d^2}{n(2d+1)}$, then

$$|\phi_k \circ \pi_k(\mathbf{w}) - f(\mathbf{w})| \leq \nu \|\mathbf{w}_0 - \mathbf{w}\|_1^\alpha \leq \nu \left(\frac{d^2}{n(2d+1)} \right)^\alpha$$

Now, for any $\mathbf{x} = (x_1, x_2, \dots, x_d) \in [0, 1]^d$, we show that it belongs to at least $d+1$ out of $2d+1$ Q_k 's: notice that in the definition of the trifling regions Ω_k (Equation 7.9.1), each x_i belongs to at most one of the following:

$$\bigcup_{j=0}^{n-1} \left(\frac{1}{n} \left(\frac{k-1}{2d+1} + j \right), \frac{1}{n} \left(\frac{k}{2d+1} + j \right) \right), \quad \text{for } k = 1, 2, \dots, 2d+1$$

so \mathbf{x} belongs to at most d trifling regions, thus \mathbf{x} belongs to at least $d+1$ of the Q_k 's.

Therefore, for any $\mathbf{x} = (x_1, x_2, \dots, x_d) \in [0, 1]^d$, at least $d+1$ of $\phi_k \circ \pi_k(\mathbf{x})$ (for

$k = 1, 2, \dots, 2d + 1$ belong to the interval $\left[f(\mathbf{x}) - \nu \left(\frac{d^2}{n(2d+1)} \right)^\alpha, f(\mathbf{x}) + \nu \left(\frac{d^2}{n(2d+1)} \right)^\alpha \right]$, then by Lemma 7.3.1,

$$\Phi(\mathbf{x}) = \text{median} \{ \phi_k \circ \pi_k(\mathbf{x}) : 1 \leq k \leq 2d + 1 \} \in \left[f(\mathbf{x}) - \nu \left(\frac{d^2}{n(2d+1)} \right)^\alpha, f(\mathbf{x}) + \nu \left(\frac{d^2}{n(2d+1)} \right)^\alpha \right]$$

hence

$$\|f - \Phi\|_{L^\infty([0,1]^d)} \leq \nu \left(\frac{d^2}{n(2d+1)} \right)^\alpha$$

Step 5: Estimate the Lipschitz constants

To compute the Lipschitz constants, we first fix $2 \leq k \leq 2d$, and compute the Lipschitz constant of $\phi_k \circ \pi_k$. It is easy to see that $\phi_1 \circ \pi_1$ and $\phi_{2d+1} \circ \pi_{2d+1}$ have at least the same regularity since there are less cuboids in Q_1 and Q_{2d+1} . Finally, by Lemma 7.3.3, their median has the same regularity.

We first compute the Lipschitz constant of π_k . Let $\mathbf{x} \in [0, 1]^d$ and $\Delta \mathbf{x}$ small, we want to find an upper bound for $|\pi_k(\mathbf{x} + \Delta \mathbf{x}) - \pi_k(\mathbf{x})|$. By definition of p_k , it only increases linearly by 1 inside intervals of length $\frac{1}{n(2d+1)}$, so $|p'_k| \leq n(2d+1)$, thus

$$\|p_k(\mathbf{x} + \Delta \mathbf{x}) - p_k(\mathbf{x})\|_1 \leq n(2d+1) \|\Delta \mathbf{x}\|_1$$

where p_k applies to \mathbf{x} elementwise. Let $\mathbf{n} \stackrel{\text{def.}}{=} \left[(n+1)^0 \quad (n+1)^1 \quad \dots \quad (n+1)^{d-1} \right]^\top$,

then

$$\begin{aligned}
|\pi_k(\mathbf{x} + \Delta\mathbf{x}) - \pi_k(\mathbf{x})| &= |\mathbf{n} \cdot p_k(\mathbf{x} + \Delta\mathbf{x}) - \mathbf{n} \cdot p_k(\mathbf{x})| \\
&= |\mathbf{n} \cdot (p_k(\mathbf{x} + \Delta\mathbf{x}) - p_k(\mathbf{x}))| \\
&\leq (n + 1)^{d-1} \cdot \|p_k(\mathbf{x} + \Delta\mathbf{x}) - p_k(\mathbf{x})\|_1 \\
&\leq (n + 1)^d (2d + 1) \|\Delta\mathbf{x}\|_1
\end{aligned}$$

therefore

$$\text{Lip}(\pi_k) \leq (n + 1)^d (2d + 1)$$

Now we compute the Lipschitz constant of ϕ_k . As mentioned above, let

$$\overline{(a_d a_{d-1} \cdots a_1)}_{n+1} \stackrel{\text{def.}}{=} \sum_{i=1}^d (n + 1)^{i-1} a_i$$

denote the number $\overline{a_1 a_2 \cdots a_d}$ in base $n + 1$. Let $P_{\overline{(a_d a_{d-1} \cdots a_1)}_{n+1}}^{(k)}$ be the center of the unique cuboid in Q_k such that $p_i \left(P_{\overline{(a_d a_{d-1} \cdots a_1)}_{n+1}}^{(k)} \right) = a_i$ for $i = 1, 2, \dots, d$. Note that $a_i \in \{0, 1, 2, \dots, n\}$ for all i . Then, ϕ_k is the unique continuous piecewise linear function that goes through $\{(j, f(P_j^{(k)})) : j = 0, 1, 2, \dots, (n + 1)^d - 1\}$ and is linear between adjacent points in this set. Then,

$$\text{Lip}(\phi_k) = \max_{j=0,1,2,\dots,(n+1)^d-2} \left| f(P_{j+1}^{(k)}) - f(P_j^{(k)}) \right| \leq \max_{j=0,1,2,\dots,(n+1)^d-2} \nu \left\| P_{j+1}^{(k)} - P_j^{(k)} \right\|_1^\alpha \leq \nu d^\alpha$$

Finally, by Lemma 7.3.3, the median function preserves regularity, hence

$$\begin{aligned}\text{Lip}(\Phi) &\leq \max_{k=1,2,\dots,2d+1} \text{Lip}(\phi_k \circ \pi_k) \\ &\leq \max_{k=1,2,\dots,2d+1} \text{Lip}(\phi_k) \cdot \text{Lip}(\pi_k) \\ &\leq \nu(n+1)^d (2d+1) d^\alpha.\end{aligned}$$

This concludes our proof.

□

Bibliography

- [1] Beatrice Acciaio, Anastasis Kratsios, and Gudmund Pammer. Designing universal causal deep learning models: the geometric (hyper)transformer. *Math. Finance*, 34(2):671–735, 2024.
- [2] Samuel Ainsworth, Jonathan Hayase, and Siddhartha Srinivasa. Git re-basin: Merging models modulo permutation symmetries. In *The Eleventh International Conference on Learning Representations*, 2023.
- [3] Alexandre Araujo, Aaron J Havens, Blaise Delattre, Alexandre Allauzen, and Bin Hu. A unified algebraic perspective on lipschitz neural networks. In *The Eleventh International Conference on Learning Representations*, 2023.
- [4] Martin Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein generative adversarial networks. In *International conference on machine learning*, pages 214–223. PMLR, 2017.
- [5] Pierre Baldi and Roman Vershynin. The capacity of feedforward neural networks. *Neural networks*, 116:288–311, 2019.
- [6] Peter L Bartlett, Dylan J Foster, and Matus J Telgarsky. Spectrally-normalized

- margin bounds for neural networks. *Advances in neural information processing systems*, 30, 2017.
- [7] Peter L. Bartlett, Nick Harvey, Christopher Liaw, and Abbas Mehrabian. Nearly-tight VC-dimension and pseudodimension bounds for piecewise linear neural networks. *J. Mach. Learn. Res.*, 20:Paper No. 63, 17, 2019.
- [8] Peter L. Bartlett and Shahar Mendelson. Rademacher and Gaussian complexities: risk bounds and structural results. *J. Mach. Learn. Res.*, 3(3):463–482, 2003. Computational learning theory.
- [9] Denis Belomestny, Alexey Naumov, Nikita Puchkin, and Sergey Samsonov. Simultaneous approximation of a smooth function and its derivatives by deep neural networks with piecewise-polynomial activations. *Neural Networks*, 161:242–253, 2023.
- [10] Cyril Bénézet, Ziteng Cheng, and Sebastian Jaimungal. Learning conditional distributions on continuous spaces. *arXiv preprint arXiv:2406.09375*, 2024.
- [11] Aritra Bhowmick, Meenakshi D’Souza, and G Srinivasa Raghavan. Lipbab: Computing exact lipschitz constant of relu networks. In *Artificial Neural Networks and Machine Learning–ICANN 2021: 30th International Conference on Artificial Neural Networks, Bratislava, Slovakia, September 14–17, 2021, Proceedings, Part IV 30*, pages 151–162. Springer, 2021.
- [12] Jiezhong Cao, Langyuan Mo, Yifan Zhang, Kui Jia, Chunhua Shen, and Mingkui Tan. Multi-marginal wasserstein gan. *Advances in Neural Information Processing Systems*, 32, 2019.

- [13] Christa Cuchiero, Philipp Schmock, and Josef Teichmann. Global universal approximation of functional input maps on weighted spaces. *arXiv preprint arXiv:2306.03303*, 2023.
- [14] Weinan E, Jiequn Han, and Arnulf Jentzen. Deep learning-based numerical methods for high-dimensional parabolic partial differential equations and backward stochastic differential equations. *Commun. Math. Stat.*, 5(4):349–380, 2017.
- [15] Dennis Elbrächter, Philipp Grohs, Arnulf Jentzen, and Christoph Schwab. DNN expression rate analysis of high-dimensional PDEs: application to option pricing. *Constr. Approx.*, 55(1):3–71, 2022.
- [16] Dennis Elbrächter, Dmytro Perekrestenko, Philipp Grohs, and Helmut Bölcskei. Deep neural network approximation theory. *IEEE Transactions on Information Theory*, 67(5):2581–2623, 2021.
- [17] Rahim Entezari, Hanie Sedghi, Olga Saukh, and Behnam Neyshabur. The role of permutation invariance in linear mode connectivity of neural networks. In *Sparcity in Neural Networks-Advancing Understanding and Practice: SNN Workshop 2021*, 2021.
- [18] Mahyar Fazlyab, Taha Entesari, Aniket Roy, and Rama Chellappa. Certified robustness via dynamic margin maximization and improved lipschitz regularization. *Advances in Neural Information Processing Systems*, 36, 2024.
- [19] Herbert Federer. *Geometric measure theory*. Springer-Verlag New York, Inc., New York., 1969.

- [20] Paul Geuchen, Thomas Heindl, Dominik Stöger, and Felix Voigtländer. Upper and lower bounds for the lipschitz constant of random neural networks. *CoRR*, 2023.
- [21] Robin Giles. A generalization of the strict topology. *Trans. Amer. Math. Soc.*, 161:467–474, 1971.
- [22] Gene H. Golub and Charles F. Van Loan. *Matrix computations*, volume 3 of *Johns Hopkins Series in the Mathematical Sciences*. Johns Hopkins University Press, Baltimore, MD, second edition, 1989.
- [23] Lukas Gonon and Christoph Schwab. Deep ReLU neural networks overcome the curse of dimensionality for partial integrodifferential equations. *Anal. Appl. (Singap.)*, 21(1):1–47, 2023.
- [24] Alexis Goujon, Joaquim Campos, and Michael Unser. Stable parameterization of continuous and piecewise-linear functions. *Appl. Comput. Harmon. Anal.*, 67:Paper No. 101581, 27, 2023.
- [25] Henry Gouk, Eibe Frank, Bernhard Pfahringer, and Michael J Cree. Regularisation of neural networks by enforcing Lipschitz continuity. *Machine Learning*, 110:393–416, 2021.
- [26] Lyudmila Grigoryeva and Juan-Pablo Ortega. Universal discrete-time reservoir computers with stochastic inputs and linear readouts using non-homogeneous state-affine systems. *J. Mach. Learn. Res.*, 19:Paper No. 24, 40, 2018.

- [27] Ingo Gühring, Gitta Kutyniok, and Philipp Petersen. Error bounds for approximations with deep ReLU neural networks in $W^{s,p}$ norms. *Anal. Appl. (Singap.)*, 18(5):803–859, 2020.
- [28] Ingo Gühring and Mones Raslan. Approximation rates for neural networks with encodable weights in smoothness spaces. *Neural Networks*, 134:107–130, 2021.
- [29] Jiequn Han, Arnulf Jentzen, and Weinan E. Solving high-dimensional partial differential equations using deep learning. *Proc. Natl. Acad. Sci. USA*, 115(34):8505–8510, 2018.
- [30] Jakob Heiss, Josef Teichmann, and Hanna Wutte. How (implicit) regularization of relu neural networks characterizes the learned function—part ii: the multi-d case of two layers with random first layer. *arXiv preprint arXiv:2303.11454*, 2023.
- [31] Roger A. Horn and Charles R. Johnson. *Matrix analysis*. Cambridge University Press, Cambridge, second edition, 2013.
- [32] Kurt Hornik, Maxwell Stinchcombe, and Halbert White. Universal approximation of an unknown mapping and its derivatives using multilayer feedforward networks. *Neural networks*, 3(5):551–560, 1990.
- [33] Songyan Hou, Parnian Kassraie, Anastasis Kratsios, Andreas Krause, and Jonas Rothfuss. Instance-dependent generalization bounds via optimal transport. *J. Mach. Learn. Res.*, 24:Paper No. [349], 51, 2023.
- [34] Itay Hubara, Matthieu Courbariaux, Daniel Soudry, Ran El-Yaniv, and Yoshua Bengio. Quantized neural networks: Training neural networks with low precision

- weights and activations. *Journal of Machine Learning Research*, 18(187):1–30, 2018.
- [35] Hui Jin and Guido Montúfar. Implicit bias of gradient descent for mean squared error regression with two-layer wide neural networks. *Journal of Machine Learning Research*, 24(137):1–97, 2023.
- [36] Matt Jordan and Alexandros G Dimakis. Exactly computing the local lipschitz constant of relu networks. *Advances in Neural Information Processing Systems*, 33:7344–7353, 2020.
- [37] Minh Kim. Algebraic characterization of the voronoi cell structure of the a_n lattice. *arXiv preprint arXiv:2304.10186*, 2023.
- [38] A. Kolmogoroff. über die beste Annäherung von Funktionen einer gegebenen Funktionenklasse. *Ann. of Math. (2)*, 37(1):107–110, 1936.
- [39] Alexander Korotin, Daniil Selikhanovych, and Evgeny Burnaev. Neural optimal transport. In *The Eleventh International Conference on Learning Representations*, 2021.
- [40] Anastasis Kratsios, Behnoosh Zamanlooy, Tianlin Liu, and Ivan Dokmanić. Universal approximation under constraints is possible with transformers. In *International Conference on Learning Representations (Spotlight)*, 2022.
- [41] H. W. Kuhn. Some combinatorial lemmas in topology. *IBM J. Res. Develop.*, 4:508–524, 1960.
- [42] Urs Lang and Conrad Plaut. Bilipschitz embeddings of metric spaces into space forms. *Geom. Dedicata*, 87(1-3):285–307, 2001.

- [43] Zhen Lei, Lei Shi, and Chenyu Zeng. Solving parametric partial differential equations with deep rectified quadratic unit neural networks. *Journal of Scientific Computing*, 93(3):80, 2022.
- [44] Mario Lezcano-Casado and David Martinez-Rubio. Cheap orthogonal constraints in neural networks: A simple parametrization of the orthogonal and unitary group. In *International Conference on Machine Learning*, pages 3794–3803. PMLR, 2019.
- [45] Shuai Li, Kui Jia, Yuxin Wen, Tongliang Liu, and Dacheng Tao. Orthogonal deep neural networks. *IEEE transactions on pattern analysis and machine intelligence*, 43(4):1352–1368, 2019.
- [46] Jianfeng Lu, Zuowei Shen, Haizhao Yang, and Shijun Zhang. Deep network approximation for smooth functions. *SIAM Journal on Mathematical Analysis*, 53(5):5465–5506, 2021.
- [47] Raffaele Marino and Nicolas Macris. Solving non-linear Kolmogorov equations in large dimensions by using deep learning: a numerical comparison of discretization schemes. *J. Sci. Comput.*, 94(1):Paper No. 8, 31, 2023.
- [48] Manor Mendel and Assaf Naor. Nonlinear spectral calculus and super-expanders. *Publ. Math. Inst. Hautes Études Sci.*, 119:1–95, 2014.
- [49] Takeru Miyato, Toshiki Kataoka, Masanori Koyama, and Yuichi Yoshida. Spectral normalization for generative adversarial networks. *arXiv preprint arXiv:1802.05957*, 2018.

- [50] Rotem Mulayoff, Tomer Michaeli, and Daniel Soudry. The implicit bias of minima stability: A view from function space. *Advances in Neural Information Processing Systems*, 34:17749–17761, 2021.
- [51] James R. Munkres. *Topology*. Prentice Hall, Inc., Upper Saddle River, NJ, second edition, 2000.
- [52] Ariel Neufeld and Philipp Schmocker. Universal approximation property of random neural networks. *arXiv preprint arXiv:2312.08410*, 2023.
- [53] Behnam Neyshabur, Srinadh Bhojanapalli, David Mcallester, and Nati Srebro. Exploring generalization in deep learning. In I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017.
- [54] Behnam Neyshabur, Russ R Salakhutdinov, and Nati Srebro. Path-sgd: Path-normalized optimization in deep neural networks. *Advances in neural information processing systems*, 28, 2015.
- [55] Piotr W. Nowak. Coarse embeddings of metric spaces into Banach spaces. *Proc. Amer. Math. Soc.*, 133(9):2589–2596, 2005.
- [56] J. A. A. Opschoor, Ch. Schwab, and J. Zech. Exponential ReLU DNN expression of holomorphic maps in high dimension. *Constr. Approx.*, 55(1):537–582, 2022.
- [57] Sejun Park, Chulhee Yun, Jaeho Lee, and Jinwoo Shin. Minimum width for universal approximation. In *International Conference on Learning Representations*, 2021.

- [58] Patricia Pauli, Anne Koch, Julian Berberich, Paul Kohler, and Frank Allgöwer. Training robust neural networks using lipschitz bounds. *IEEE Control Systems Letters*, 6:121–126, 2021.
- [59] Isaac Z. Pesenson. Estimates of Kolmogorov, Gelfand and linear n -widths on compact Riemannian manifolds. *Proc. Amer. Math. Soc.*, 144(7):2985–2998, 2016.
- [60] Philipp Petersen, Mones Raslan, and Felix Voigtlaender. Topological properties of the set of functions generated by neural networks of fixed size. *Found. Comput. Math.*, 21(2):375–444, 2021.
- [61] Philipp Petersen and Felix Voigtlaender. Optimal approximation of piecewise smooth functions using deep relu neural networks. *Neural Networks*, 108:296–330, 2018.
- [62] Philipp Petersen and Jakob Zech. Mathematical theory of deep learning. *arXiv preprint arXiv:2407.18384*, 2024.
- [63] Guergana Petrova and Przemyslaw Wojtaszczyk. Limitations on approximation by deep and shallow neural networks. *Journal of Machine Learning Research*, 24(353):1–38, 2023.
- [64] Allan Pinkus. n -widths in approximation theory, volume 7 of *Ergebnisse der Mathematik und ihrer Grenzgebiete (3) [Results in Mathematics and Related Areas (3)]*. Springer-Verlag, Berlin, 1985.
- [65] Allan Pinkus. Approximation theory of the MLP model in neural networks. In

- Acta numerica, 1999*, volume 8 of *Acta Numer.*, pages 143–195. Cambridge Univ. Press, Cambridge, 1999.
- [66] Eric Qu, Xufang Luo, and Dongsheng Li. Data continuity matters: Improving sequence modeling with lipschitz regularizer. In *The Eleventh International Conference on Learning Representations, 2023*.
- [67] 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.
- [68] Christoph Schwab and Jakob Zech. Deep learning in high dimension: neural network expression rates for analytic functions in $L^2(\mathbb{R}^d, \gamma_d)$. *SIAM/ASA J. Uncertain. Quantif.*, 11(1):199–234, 2023.
- [69] Zuowei Shen, Haizhao Yang, and Shijun Zhang. Nonlinear approximation via compositions. *Neural Networks*, 119:74–84, 2019.
- [70] Zuowei Shen, Haizhao Yang, and Shijun Zhang. Optimal approximation rate of ReLU networks in terms of width and depth. *J. Math. Pures Appl. (9)*, 157:101–135, 2022.
- [71] Jonathan W Siegel. Optimal approximation rates for deep relu neural networks on sobolev and besov spaces. *Journal of Machine Learning Research*, 24(357):1–52, 2023.
- [72] Hossein Taheri and Christos Thrampoulidis. Generalization and stability of interpolating neural networks with minimal width. *Journal of Machine Learning Research*, 25(156):1–41, 2024.

- [73] Matus Telgarsky. Benefits of depth in neural networks. In *Conference on learning theory*, pages 1517–1539. PMLR, 2016.
- [74] Vlad Timofte, Aida Timofte, and Liaqat Ali Khan. Stone-Weierstrass and extension theorems in the nonlocally convex case. *J. Math. Anal. Appl.*, 462(2):1536–1554, 2018.
- [75] Christopher Todd. Stone-Weierstrass theorems for the strict topology. *Proc. Amer. Math. Soc.*, 16:654–659, 1965.
- [76] Yusuke Tsuzuku, Issei Sato, and Masashi Sugiyama. Lipschitz-margin training: Scalable certification of perturbation invariance for deep neural networks. *Advances in neural information processing systems*, 31, 2018.
- [77] A. W. van der Vaart and Jon A. Wellner. *Weak convergence and empirical processes—with applications to statistics*. Springer Series in Statistics. Springer, Cham, second edition, [2023] ©2023.
- [78] Gal Vardi, Gilad Yehudai, and Ohad Shamir. On the optimal memorization power of relu neural networks. In *International Conference on Learning Representations*, 2021.
- [79] Roman Vershynin. *High-dimensional probability*, volume 47 of *Cambridge Series in Statistical and Probabilistic Mathematics*. Cambridge University Press, Cambridge, 2018. An introduction with applications in data science, With a foreword by Sara van de Geer.
- [80] Roman Vershynin. Memory capacity of neural networks with threshold and rectified linear unit activations. *SIAM J. Math. Data Sci.*, 2(4):1004–1033, 2020.

- [81] Aladin Virmaux and Kevin Scaman. Lipschitz regularity of deep neural networks: analysis and efficient estimation. *Advances in Neural Information Processing Systems*, 31, 2018.
- [82] Anton Xue, Lars Lindemann, Alexander Robey, Hamed Hassani, George J Pappas, and Rajeev Alur. Chordal sparsity for lipschitz constant estimation of deep neural networks. In *2022 IEEE 61st Conference on Decision and Control (CDC)*, pages 3389–3396. IEEE, 2022.
- [83] Dmitry Yarotsky. Error bounds for approximations with deep relu networks. *Neural networks*, 94:103–114, 2017.
- [84] Dmitry Yarotsky. Optimal approximation of continuous functions by very deep relu networks. In Sébastien Bubeck, Vianney Perchet, and Philippe Rigollet, editors, *Proceedings of the 31st Conference On Learning Theory*, volume 75 of *Proceedings of Machine Learning Research*, pages 639–649. PMLR, 06–09 Jul 2018.
- [85] Dmitry Yarotsky and Anton Zhevnerchuk. The phase diagram of approximation rates for deep neural networks. *Advances in neural information processing systems*, 33:13005–13015, 2020.