

UC San Diego

UC San Diego Electronic Theses and Dissertations

Title

Modern Problems in Mathematical Signal Processing: Quantized Compressed Sensing and Randomized Neural Networks

Permalink

<https://escholarship.org/uc/item/1mh5164c>

Author

Nelson, Aaron Andrew

Publication Date

2019

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA SAN DIEGO

**Modern Problems in Mathematical Signal Processing: Quantized Compressed Sensing
and Randomized Neural Networks**

A dissertation submitted in partial satisfaction of the
requirements for the degree
Doctor of Philosophy

in

Mathematics with a specialization in Computational Science

by

Aaron A. Nelson

Committee in charge:

Professor Rayan Saab, Chair
Professor Yoav Freund
Professor Michael Holst
Professor Melvin Leok
Professor Piya Pal

2019

The dissertation of Aaron A. Nelson is approved, and it is acceptable in quality and form for publication on microfilm and electronically:

Chair

University of California San Diego

2019

TABLE OF CONTENTS

Signature Page	iii
Table of Contents	iv
List of Figures	vi
List of Tables	vii
Acknowledgements	viii
Vita	ix
Abstract of the Dissertation	x
Chapter 1	Introduction and Background	1
	1.1 Mathematical Signal Processing in the Modern World	1
	1.2 Compressed sensing and quantization	2
	1.3 High-dimensional data analysis and machine learning	4
	1.4 Organization of the manuscript	5
Chapter 2	One-Bit Compressed Sensing on Manifolds	7
	2.1 Introduction	7
	2.2 Background and Notation	8
	2.2.1 Introduction to compressed sensing	9
	2.2.2 Introduction to quantization	12
	2.2.3 Notation	13
	2.3 Preliminaries	13
	2.3.1 Noise-shaping quantization methods	13
	2.3.2 Binary embeddings	16
	2.3.3 Geometric Multi-Resolution Analysis	17
	2.4 Problem Formulation and Main Result	19
	2.5 Numerical Simulations	26
	2.6 Acknowledgements	27
Chapter 3	Random Vector Functional Link Networks	28
	3.1 Introduction	28
	3.2 Background and notation	31
	3.2.1 Randomized single layer neural networks	31
	3.2.2 Notation	33
	3.3 Theoretical results in Euclidean space	33
	3.4 Preliminaries	38
	3.4.1 A concentration bound for classic Monte Carlo integration	38

3.4.2	Smooth, compact manifolds in Euclidean space	41
3.5	Proof of Theorem 3.3.1	43
3.5.1	A convolution identity	43
3.5.2	The limit-integral representation	49
3.5.3	Monte-Carlo integral approximation	54
3.5.4	Bounding the asymptotic mean square error	57
3.6	Proofs of Corollary 3.3.3 and Theorem 3.3.4	59
3.6.1	Proof of Corollary 3.3.3	59
3.6.2	Proof of Theorem 3.3.4	61
3.7	Theoretical results on submanifolds of Euclidean space	66
3.7.1	Adapting RVFL networks to d -manifolds	67
3.7.2	Main results on d -manifolds	68
3.8	Numerical Simulations	78
3.9	Acknowledgements	82
Bibliography		83

LIST OF FIGURES

Figure 2.5.1: Log-scale plot of reconstruction error for Algorithm 1 as a function of the oversampling rate.	26
Figure 3.8.1: Log-scale plot of average error for Algorithm 2 as a function of the number of neural network nodes.	81

LIST OF TABLES

Table 3.2.1: List of typical activation functions for neural networks. 32

ACKNOWLEDGEMENTS

Special thanks to Felix Kraher, Sara Krause-Solberg, and Johannes Maly for sharing their GMRA code for use in numerical experiments, which they adapted from that originally provided by Mauro Maggioni.

A conference version of the material presented in Chapter 2 appeared in the proceedings of the 17th Annual International Conference on Sampling Theory and Applications (SampTA), 2019. The dissertation author was the primary researcher and author of this material. Thanks are extended to coauthors Mark Iwen, Eric Lybrand, and Rayan Saab for their contributions and insights aiding in the development of this material.

A journal version of the material presented in Chapter 3 is currently being prepared for submission for publication. The dissertation author was the primary researcher and author of this material. Thanks are extended to coauthors Deanna Needell, Rayan Saab, and Palina Salanevich for their contributions and insights aiding in the development of this material.

VITA

2012	Bachelor of Science in Mathematics with a Minor in Physics, <i>summa cum laude</i> , University of Connecticut
2014	Master of Science in Applied Mathematics, <i>Distinguished Graduate</i> , Air Force Institute of Technology
2019	Doctor of Philosophy in Mathematics with a specialization in Computational Science, University of California San Diego
2012-Present	Commissioned Officer, United States Air Force

PUBLICATIONS

A. S. Bandeira, J. Cahill, D. G. Mixon, A. A. Nelson, *Fundamental Limits of Phase Retrieval*, In: Proceedings of the 10th International Conference on Sampling Theory and Applications (SampTA), 2013.

A. S. Bandeira, J. Cahill, D. G. Mixon, A. A. Nelson, *Saving phase: Injectivity and stability for phase retrieval*, Applied and Computational Harmonic Analysis, 37(1):106–125, 2014.

M. Begue, D. J. Kelleher, A. Nelson, H. Panzo, R. Pellico, A. Teplyaev, *Random walks on barycentric subdivisions and the Strichartz hexacarpent*, Experimental Mathematics, 21(4):402–417, 2012.

M. Fickus, D. G. Mixon, A. A. Nelson, Y. Wang, *Phase retrieval with very few measurements*, Linear Algebra and its Applications, 449:475–499, 2014.

M. A. Iwen, E. Lybrand, A.A. Nelson, R. Saab, *New Algorithms and Improved Guarantees for One-Bit Compressed Sensing on Manifolds*, in: Proceedings of the 17th International Conference on Sampling Theory and Applications (SampTA), 2019.

D. Molitor, D. Needell, A. A. Nelson, R. Saab, P. Salanevich, *Classification scheme for binary data with extensions*, In: Compressed Sensing and its Applications, Springer, 2018.

A. A. Nelson, *About Phase: Synthetic Aperture Radar and the Phase Retrieval Problem*, Air Force Institute of Technology Graduate School of Engineering and Management, Thesis No. AFIT-ENC-14-M-03, 2014.

ABSTRACT OF THE DISSERTATION

**Modern Problems in Mathematical Signal Processing: Quantized Compressed Sensing
and Randomized Neural Networks**

by

Aaron A. Nelson

Doctor of Philosophy in Mathematics with a specialization in Computational Science

University of California San Diego, 2019

Professor Rayan Saab, Chair

We study two problems from mathematical signal processing. First, we consider problem of approximately recovering signals on a smooth, compact manifold from one-bit linear measurements drawn from either a Gaussian ensemble, partial circulant ensemble, or bounded orthonormal ensemble and quantized using $\Sigma\Delta$ or distributed noise-shaping schemes. We construct a convex optimization algorithm for signal recovery that, given a Geometric Multi-Resolution Analysis approximation of the manifold, guarantees signal recovery with high probability. We prove an upper bound on the recovery error which outperforms prior works that use memoryless scalar quantization, requires a simpler analysis, and extends the class of measurements beyond

Gaussians.

Second, we consider the problem of approximation continuous functions on compact domains using neural networks. The learning speed of feed-forward neural networks is notoriously slow and has presented a bottleneck in deep learning applications for several decades. For instance, gradient-based learning algorithms, which are used extensively to train neural networks, tend to work slowly when all of the network parameters must be iteratively tuned. To counter this, both researchers and practitioners have tried introducing randomness to reduce the learning requirement. Based on the original construction of B. Igel'nik and Y.H. Pao, single layer neural-networks with random input-to-hidden layer weights and biases have seen success in practice, but the necessary theoretical justification is lacking. We begin to fill this theoretical gap by providing a (corrected) rigorous proof that the Igel'nik and Pao construction is a universal approximator for continuous functions on compact domains, with ε -error convergence rate inversely proportional to the number of network nodes; we then extend this result to the non-asymptotic setting using a concentration inequality for Monte-Carlo integral approximations. We further adapt this randomized neural network architecture to approximate functions on smooth, compact submanifolds of Euclidean space, providing theoretical guarantees in both the asymptotic and non-asymptotic cases.

The views expressed in this dissertation are those of the author and do not reflect the official policy or position of the U.S. Air Force, Department of Defense, or U.S. Government.

Chapter 1

Introduction and Background

1.1 Mathematical Signal Processing in the Modern World

Mathematical signal processing is a critical component of modern technology and data science – everyday devices ranging from mobile phones and digital cameras to medical imaging and radar systems, as well as internet and wireless communication, audio systems, and chemical or physical sensors are just a few examples of modern technologies which require advanced processing of signals, images, and/or data. Engineering progress related to these technologies goes hand-in-hand with the corresponding applied mathematical theory. Indeed, engineering disciplines provide new, challenging problems and inspiration for mathematics, while mathematics leads to engineering advances by providing rigorous, sophisticated solutions to these problems. The Shannon-Nyquist sampling theory, which has its origins in the 1940s, is considered to be one of the first such developments, setting the foundations for measuring and transforming analog signals into digital form, as well as the fast, efficient transmission of this information across multiple media, both wired and wireless. Such developments in mathematics and engineering lead to the theory of wavelets, first introduced in the 1980s, which in turn has led to new compression, denoising, and other image processing methods. These new methods have provided important

mathematical insights with implications throughout the mathematical disciplines.

When practitioners first began to realize the possibility of sketching analog signals using significantly less numerical information in the early 2000s, the powerful mathematical theory of compressed sensing was born. This theory predicts that the Shannon-Nyquist rate may be significantly overcome, in the sense that compressible (i.e., sparse) signals can be recovered efficiently from what was previously believed to be highly incomplete linear measurements. This surprising discovery led to a flurry of research on the applications of compressed sensing in other areas, including biomedical imaging, radar, and astronomy, to name a few. It has also led to various mathematical results by establishing links between fields such as harmonic analysis, random matrix theory, and convex optimization.

More recently, mathematical signal processing has come to a turning point. Indeed, not only are traditional analog signals like sound, images, and video collected and processed, but entire daily activities are measured in various ways. The amount of data acquired, stored, and transmitted on a daily basis is increasing rapidly, and the ability to efficiently process these massive data sets is becoming ever more important. Together with new technological advances, there is an increasing demand for novel mathematical methods to perform efficient information processing at large scales. As such, real-world data analysis requires fundamentally new ideas and approaches, with significant potential for cross-disciplinary mathematical developments. In fact, mathematical techniques involved in signal and data processing, as well as proofs of corresponding results, often involve various fields including harmonic analysis, optimization, probability theory, numerical linear algebra, and graph theory, among others.

1.2 Compressed sensing and quantization

A major shift in sampling theory occurred over the past fifteen years with the development of compressed sensing, which predicts the robust recovery of sparse signals from vastly incomplete

linear random measurements using efficient methods such as convex optimization or certain greedy algorithms. Indeed, one major difference between compressed sensing and classical sampling theory is that signal recovery is no longer a linear process. Huge research efforts emerged from this new field, including the development of new recovery algorithms, tailored to different applications, as well as the analysis of random measurements, both structured and unstructured. More recently, the extension of compressed sensing methods to more general situations is a popular activity, and new results and breakthroughs are appearing in both theory and practice. For instance, the recovery of low rank matrices from few linear random measurements, which has important applications in high-dimensional data analysis, forms the basis for the development of a robust theory of compressed sensing via non-linear measurements. Moreover, recent results in the well-known phase retrieval problem, which appears in applications such as diffraction imaging and x-ray crystallography, show that it is possible to recover signals from the absolute values of their scalar products with respect to a small number of elements of certain finite frames. Even the mathematical analysis of (structured) random matrices requires sophisticated tools from compressed sensing.

The current state of research on quantization in mathematical signal processing focuses on developing efficient methods for obtaining robust quantizers of analog signals, including low-bit quantization schemes such as Sigma-Delta (abbreviated by $\Sigma\Delta$). The use of quantization in the context of compressed sensing is also a major research topic, including one-bit compressed sensing, extensions to low rank matrix recovery, and the analysis of structured random measurement schemes. Various results from modern embedding theory, which has its origins in the 1980s with the Johnson-Lindenstrauss lemma, have led to practical algorithms for quantized compressed sensing with theoretical recovery guarantees.

1.3 High-dimensional data analysis and machine learning

The analysis of data in high-dimensions poses several theoretical and practical challenges due to the complexity inherent in the ambient system. Recent research efforts have attempted to overcome this problem by using structural data assumptions. In this way, one hopes to reduce the degrees-of-freedom of the system while keeping the problem tractable. Some of the first works along this line of reasoning considered the recovery of sparse signals, while subsequent works analyzed more general union-of-subspaces models and the recovery of low rank matrices, which have applications in the phaseless reconstruction problem and other bilinear inverse problems. Yet another line of work following this approach studies manifold models. Using this approach, one assumes that the structural constraints are given by (unions of finitely many) manifolds having low intrinsic dimension. Arguably, working with manifold models is better adapted to real world data than sparsity, and so has gained traction in recent analysis methods.

Current developments in the field of machine learning, led primarily by the computer science community, are producing exceptional results on highly complex decision tasks. The great empirical success of deep learning shows that algorithms can beat humans in several tasks, such as image classification or learning how to play complex games. However, the mathematical modeling of such algorithms is still in its infancy. Indeed, the theoretical understanding of the behavior of such machine learning techniques remains largely a mystery, and developing mathematical guarantees is crucial for the advancement of the field – without such comprehension, practitioners are operating solely based on empirical evidence. Recently, the mathematical community has contributed novel methods, such as diffusion maps and kernel methods, to aid in the understanding of machine learning techniques. This effort has led to an increased theoretical understanding of and insights into deep learning methods, combining tools from harmonic analysis, high-dimensional optimization, and probability theory. The introduction of methods from high-dimensional data analysis which take advantage of the intrinsic structure of data, (e.g.,

Principal Component Analysis and more sophisticated variants thereof) has also led to a greater theoretical understanding of deep learning, but many open questions remain.

1.4 Organization of the manuscript

In this dissertation, we consider two topics from mathematical signal processing that draw on results from various theoretical backgrounds. First, we consider the problem of recovering a signal from quantized linear measurements; second, we delve into the problem of approximating continuous functions using artificial neural networks. Although their motivations are fundamentally different, the common theme connecting these problems in this work is the desire to leverage underlying signal structure (e.g., sparsity) to improve recovery or estimation accuracy. In particular, we explore the increasingly popular assumption that the signal class we are interested in lies on a low-dimensional submanifold of Euclidean space, which is important for dimensionality reduction. Indeed, much of modern signal processing deals with massive amounts of data lying in high-dimensional spaces, and the computational complexity of practical algorithms suffers as a result. This complexity can often be reduced by designing algorithms that take advantage of intrinsic signal structure.

The remainder of the manuscript is organized as follows: In Chapter 2, we introduce the quantized compressed sensing problem, with a focus on using one-bit quantization schemes in conjunction with random linear measurements of structured signals. Specifically, we study the problem of approximately recovering signals on a low-dimensional submanifold of Euclidean space from one-bit linear measurements drawn from either a Gaussian ensemble, partial circulant ensemble, or bounded orthogonal ensemble and quantized using stable noise-shaping techniques. We construct a convex optimization algorithm for signal recovery that, given a suitable approximation to the manifold, guarantees signal recovery with high probability. We prove an upper bound on the recovery error which outperforms prior works that use memoryless scalar quantization,

requires a simpler analysis, and extends the class of measurements beyond Gaussians. Finally, we illustrate our results with numerical experiments.

In Chapter 3, we introduce a randomized neural network architecture, known as the Random Vector Functional Link, for approximating continuous functions on compact sets. Based on the original construction of B. Igel'nik and Y.H. Pao, these single layer neural-networks use random input-to-hidden layer weights and biases to reduce the complexity and running time of learning algorithms. This approach has seen success in practice, but the necessary theoretical justification is lacking. To begin to fill this theoretical gap, we provide a (corrected) rigorous proof that the Igel'nik and Pao construction is a universal approximator for continuous functions on compact domains, with ε -error convergence rate inversely proportional to the number of network nodes; we then extend this result to the non-asymptotic setting using a concentration inequality for Monte-Carlo integral approximations. We further adapt this randomized neural network architecture to approximate functions on smooth, compact submanifolds of Euclidean space, providing theoretical guarantees in both the asymptotic and non-asymptotic cases that depend on the manifold dimension rather than the ambient dimension. Finally, we illustrate our results on manifolds with numerical experiments.

A conference version of the material presented in Chapter 2 appeared in the proceedings of the 17th Annual International Conference on Sampling Theory and Applications (SampTA), 2019 [ILNS19]. Moreover, a journal version of the material presented in Chapter 3 is currently being prepared for submission for publication.

Chapter 2

One-Bit Compressed Sensing on Manifolds

2.1 Introduction

Compressed sensing [CRT06, D⁺06] demonstrates that structured high dimensional signals such as sparse vectors or low-rank matrices can be recovered from few random linear measurements. Recovery is typically formulated as a convex optimization problem whose minimizer cannot be expressed analytically and must be solved for using numerical algorithms running on digital devices. Thus, it is necessary to consider the effect of quantization in the design of the recovery algorithms. Indeed, sparse vector recovery and low-rank matrix recovery have been studied in the presence of various quantization schemes [GLP⁺10, HS18, JLBB13, LS17, PV13]. We look to extend these results to account for those structured signals that lie on a compact, low-dimensional submanifold of \mathbb{R}^N for which we have a Geometric Multi-Resolution Analysis (GMRA) [ACM12]. Our work is motivated by the results of Iwen et al. in [IKKSM18] where they assume memoryless scalar quantized Gaussian measurements, and we provide better error bounds that hold for a wider class of measurement ensembles.

As in [IKKSM18], a key component of our technique is the GMRA which approximates the manifold at various levels of refinement. At each level the GMRA is a collection of ap-

proximate tangent spaces about certain known "centers", and the quality of the approximation improves with every level. Unlike in [IKKSM18], the quantization schemes that we use are $\Sigma\Delta$ or distributed noise-shaping methods (see, e.g., [GLP⁺10, HS18]) and the compressed sensing measurements that our results apply to include those drawn from Gaussian ensembles, partial circulant ensembles (PCE) or bounded orthogonal ensembles (BOE). Our proposed reconstruction method is summarized in Algorithm 1. This simple algorithm first finds a GMRA center that quantizes to a bit sequence close to the quantized measurements, where "closeness" is determined using a pseudo-metric that respects the quantization; it then optimizes over all points in the associated approximate tangent space to enforce, as much as possible, the consistency of the quantization. Using the results of [HS18] we prove that the quantization error associated with our proposed reconstruction algorithm decays polynomially or exponentially as a function of the number of measurements, depending on the quantization scheme. This greatly improves on the sub-linear error decay associated with scalar quantization in [IKKSM18].

2.2 Background and Notation

Let \mathcal{M} be a smooth, compact submanifold of \mathbb{R}^N . Given a linear transformation $A \in \mathbb{R}^{m \times N}$, a discrete set \mathcal{A} , and a quantization map $\mathcal{Q}: \mathbb{R}^m \rightarrow \mathcal{A}$, we seek to recover a vector $x \in \mathcal{M}$ from the quantized linear measurements $q = \mathcal{Q}(Ax)$. It is assumed that we do not know the structure of \mathcal{M} a priori, but instead have access to a structured dictionary model, which we describe later. In [IKKSM18], Iwen et al. study the case where the signal x is restricted to a manifold $\mathcal{M} \subset \mathbb{S}^{N-1}$ and the quantization scheme is memoryless scalar quantization (MSQ), i.e., A is a standard Gaussian matrix, $\mathcal{A} = \{\pm 1\}$, and $\mathcal{Q}(\cdot) = \text{sign}(\cdot)$. In their paper, Iwen et al. propose an algorithm for recovering x from such measurements and show that the associated error decays like $O(m^{-1/7})$ [IKKSM18]. Such slow error decay, associated with MSQ (see, e.g., [HS18] and references therein), has also been seen in the context of sparse vector recovery in the compressed

sensing literature. Indeed, it is known in that setting that the error under any reconstruction scheme using MSQ measurements cannot decay faster than $O(m^{-1})$ [GVT95] (see also [BJKS14]).

To achieve better error rates than seen in [IKKSM18], one must use more sophisticated quantization schemes. For example, in the sparse vector setting, noise-shaping techniques such as $\Sigma\Delta$ and distributed noise-shaping leverage redundancy of the measurements to ensure error decay like $O(m^{-r})$ or $O(\beta^{-cm})$ for some parameters $r \in \mathbb{N}$, $\beta > 1$ that depend on the quantization scheme (see, e.g., [CG16, SWY18a]). To take advantage of these better decay rates, we study the quantized manifold recovery problem in this setting. In particular, we study the recovery of $x \in \mathcal{M} \subset B_2^N$ from the measurements $q = \mathcal{Q}(Ax)$ using $\Sigma\Delta$ and distributed noise-shaping methods by leveraging an embedding result from [HS18]. This allows us to use not only Gaussian measurement ensembles, as in [IKKSM18], but also more structured systems drawn from partial circulant ensembles (POEs) and bounded orthogonal ensembles (BOEs).

Since our results involve material from several intersecting fields, it is convenient now to provide a few brief introductions. The material that follows is therefore an amalgamation of background information pertaining to compressed sensing and quantization. Much of this information can be found in [HS18] and references therein.

2.2.1 Introduction to compressed sensing

Let $x \in \mathbb{C}^N$ be an unknown vector that we wish to reconstruct from m linear measurements of the form

$$y(j) = \langle a_j, x \rangle + w(j), \quad j = 1, \dots, m,$$

where each $a_j \in \mathbb{C}^N$ is a known measurement vector and $w(j) \in \mathbb{C}$ is unknown measurement noise. The collection of measurements $y = \{y(j)\}_{j=1}^m$ is often written in the matrix form $y = Ax + w$, where the j th row of $A \in \mathbb{C}^{m \times N}$ is the vector a_j and $w = \{w(j)\}_{j=1}^m$. When x is *sparse* (i.e., most

of its entries are zero) or well-approximated by sparse vectors, it can be recovered from the measurements y even when the matrix A is tall (i.e., $m \ll N$) [HS18]. The study of measurement matrices, recovery algorithms, and reconstruction error guarantees in this setting comprises the field of *compressed sensing*. One of the most popular compressed sensing techniques is ℓ_1 -minimization, where the vector x is approximated by

$$x^\sharp := \arg \min_{z \in \mathbb{C}^N} \|z\|_1 \quad \text{subject to} \quad \|Az - y\|_2 \leq \eta, \quad (2.2.1)$$

where $\eta > 0$ is an assumed upper bound on the norm of the noise vector w .

For measurement matrices A that are selected entrywise at random (independently) from Gaussian or Bernoulli distributions, it is known that, with high probability on the draw of the matrix, the solution x^\sharp to (2.2.1) satisfies the error bound

$$\|x - x^\sharp\|_2 \lesssim \eta + \frac{\sigma_k(x)_1}{\sqrt{k}}$$

whenever $m \gtrsim k \log(N/k)$ (see, e.g., [CRT06, D⁺06, MPTJ08, BDDW08]); here, $\sigma_k(x)_1$ denotes the ℓ_1 error of the best k -sparse approximation to x , and is therefore a measure of the best error one may hope to achieve. Reconstruction error guarantees of this form are often obtained by requiring that the measurement matrix satisfy the Restricted Isometry Property (RIP):

Definition 2.2.1 (Restricted Isometry Property (RIP)). A matrix $A \in \mathbb{C}^{m \times N}$ satisfies the (δ_k, k) -Restricted Isometry Property (RIP) if for every k -sparse $x \in \mathbb{C}^N$ we have

$$(1 - \delta_k)^2 \|x\|_2^2 \leq \|Ax\|_2^2 \leq (1 + \delta_k)^2 \|x\|_2^2.$$

Essentially, matrices A satisfying the RIP provide near-isometric embeddings when acting on sparse vectors.

For practical reasons (e.g., physical properties of the measurements, reconstruction al-

gorithm speed, etc.), one may require more structure in the matrix A than the random matrices considered above. This motivated the study of structured random matrices for compressed sensing (see, e.g., [CT06, BWD⁺06, RV08, Rau08, FR17]). Two of the more popular classes of structured random matrices, which we will use in our later algorithms, are *partial circulant ensembles* (POEs) and *bounded orthogonal ensembles* (BOEs). The following definitions of these objects can be found, for example, in [HS18].

Definition 2.2.2 (Partial Circulant Ensemble (PCE)). For $z \in \mathbb{C}^N$, let $H_z \in \mathbb{C}^{N \times N}$ be the circulant matrix given by its action $H_z x = z * x$ on $x \in \mathbb{C}^N$, where $*$ denotes circular convolution. Fix a subset $\Omega \subset \{1, 2, \dots, N\}$ of cardinality m arbitrarily. A matrix $A \in \mathbb{C}^{m \times N}$ is drawn from the partial circulant ensemble associated with Ω by choosing a vector σ whose entries are selected uniformly at random (independently) from the ± 1 Bernoulli distribution and setting the rows of A to be the rows of H_σ indexed by Ω .

Definition 2.2.3 (Bounded Orthogonal Ensemble (BOE)). Let $N^{-1/2}U \in \mathbb{C}^{N \times N}$ be any unitary matrix such that $|U(j, k)| \leq 1$ for all $j, k = 1, \dots, N$. A matrix $A \in \mathbb{C}^{m \times N}$ is drawn from the bounded orthogonal ensemble associated with U by picking each row of A uniformly at random (independently) from the set of all rows of U .

Structured random matrices drawn from both PCEs and BOEs arise naturally in compressed sensing applications, such as radar, wireless channel estimation, and magnetic resonance imaging (see, e.g., [HBRN10, RRT12, Rom09, FKS17, HHL10, LDP07, MAD⁺12, VAH⁺10]). Moreover, PCEs and BOEs satisfy the RIP with high probability whenever $m \gtrsim k \text{polylog}(N)$ [RV08, CGV13, Bou14, HR17] and are fast to implement, featuring prominently in the construction of fast Johnson-Lindenstrauss embeddings [AL13, AC09, KW11].

2.2.2 Introduction to quantization

Consider the m linear measurements $y = Ax$ of a signal $x \in \mathbb{R}^N$, where $A \in \mathbb{R}^{m \times N}$. Vector quantization is the process of mapping y to a vector of elements from some finite alphabet \mathcal{A} via a quantization map $\mathcal{Q}: \mathbb{R}^m \rightarrow \mathcal{A}^m$, thus allowing for digital storage and processing. Such digitization is necessary for many applications, and careful selection of the pair $(\mathcal{Q}, \mathcal{A})$ often leads to faster, more efficient algorithms [HS18]. For instance, in the field of binary embeddings one designs the quantization map \mathcal{Q} in such a way that pairwise distances between signals are approximately preserved (i.e., such that \mathcal{Q} is an approximate isometry) [JLBB13, PV12, PV14, YCP15]. On the other hand, in compressed sensing one requires a quantization map \mathcal{Q} and a reconstruction algorithm $\mathcal{R}: \mathcal{A}^m \rightarrow \mathbb{R}^N$ such that the reconstruction error $\|x - \mathcal{R}(\mathcal{Q}(y))\|_2$ is sufficiently small for all sparse (or approximately sparse) vectors (see, e.g., [CGK⁺15]).

Several constructions of quantization maps (and associated reconstruction algorithms) have emerged in the binary embedding and compressed sensing literature. For instance, some of the more intuitive approaches use memoryless scalar quantization (MSQ) [JLBB13, PV12, DJR17], in which one takes (sub)-Gaussian random measurements $y = Ax$ and quantizes to the alphabet $\mathcal{A} = \{\pm 1\}$ using the quantization map $\mathcal{Q}(\cdot) = \text{sign}(\cdot)$. In this way, MSQ simply quantizes the vector $y \in \mathbb{R}^m$ by mapping it to the nearest corner of the Hamming cube. Since recovery algorithms based on MSQ are notoriously suboptimal [HS18], more sophisticated approaches exist and are often based on noise-shaping techniques. These noise-shaping methods, such as $\Sigma\Delta$ quantization [GLP⁺13, KSW12, KSY14] and distributed noise-shaping quantization [Cho13, Huy16], are known for both their computational simplicity and desirable error bounds (as a function of m) [HS18]. Each of these quantization methods employs a state variable $u \in \mathbb{R}^m$ and quantizes measurements in a recursive fashion:

$$q(j) = \mathcal{Q}\left(f(y(j), \dots, y(1), u(j-1), \dots, u(1))\right) \quad j = 1, \dots, m,$$

where f is some function designed for the quantization scheme. The state variable is then updated via the state relation $Ax - q = Hu$, where $H : \mathbb{R}^m \rightarrow \mathbb{R}^m$ is a lower-triangular noise-shaping matrix. Important for the analysis (and for practical reasons) is that H and f are chosen so that whenever $\|Ax\|_\infty$ is bounded, we have that $\|u\|_\infty$ is also bounded; such quantization schemes are said to be *stable*. Since we will use both $\Sigma\Delta$ and distributed noise-shaping quantization methods in our results, we save more detailed discussions of both for a later section.

2.2.3 Notation

We use \gtrsim and \lesssim for inequalities that hold up to a constant; subscripts indicate the constant depends on a specified parameter. For any probability distribution $P : \mathbb{R}^N \rightarrow [0, 1]$, a random variable X distributed according to P is denoted by $X \sim P$, and we write its expectation as $\mathbb{E}X := \int_{\mathbb{R}^N} X dP$. Given a set $T \subset \mathbb{R}^N$, we define its radius to be $\text{rad}(T) := \sup_{x \in T} \|x\|_2$; for $g \sim \text{Norm}(0, I_N)$, the Gaussian mean width of T is defined by $\omega(T) := \mathbb{E} \sup_{x \in T} \langle g, x \rangle$. The open ℓ_p ball of radius $r > 0$ centered at $x \in \mathbb{R}^N$ is denoted by $B_p^N(x, r)$ for all $1 \leq p \leq \infty$; the ℓ_p unit-ball centered at the origin is abbreviated B_p^N .

2.3 Preliminaries

In this section, we introduce some important mathematical concepts that will be used later in our theoretical results. These introductions are by no means exhaustive, and are instead intended to provide a basic understanding that will enable us to properly introduce our main results and supporting material.

2.3.1 Noise-shaping quantization methods

Noise-shaping quantizers, first proposed for analog-to-digital conversion of bandlimited functions (see, e.g., [DD03, Gün03, PST17]), have enjoyed success essentially because they

push the quantization error toward the nullspace of the associated reconstruction operator. Such methods have since been extended to the settings of finite frames (see, e.g., [CGK⁺15]) and compressed sensing [GLP⁺13, KSW12, Cho13]. In fact, the approaches based on $\Sigma\Delta$ quantization and beta encoders (i.e., distributed noise-shaping) have been shown to achieve near-optimal bounds for sub-Gaussian measurements [Cho13, SWY18b].

To explain these methods, let $\mathcal{A}_{L,\delta}$ denote a real quantization alphabet, consisting of $2L$ symmetric levels with spacing 2δ :

$$\mathcal{A}_{L,\delta} := \{ -(2L-1)\delta, -(2L-2)\delta, \dots, -\delta, \delta, \dots, (2L-2)\delta, (2L-1)\delta \}.$$

A noise-shaping quantizer $\mathcal{Q}: \mathbb{R}^m \rightarrow \mathcal{A}_{L,\delta}^m$ is defined using a state variable $u \in \mathbb{R}^m$ such that, for each $y \in \mathbb{R}^m$, the resulting quantization $q := \mathcal{Q}(y)$ satisfies the noise-shaping relation

$$y - q = Hu \tag{2.3.1}$$

for some lower-triangular noise-shaping matrix $H: \mathbb{R}^m \rightarrow \mathbb{R}^m$. Careful attention must be taken when designing such methods to ensure that the process is stable, which we make precise below:

Definition 2.3.1 (Stable noise-shaping quantization). Given a finite alphabet \mathcal{A} and a noise-shaping, lower-triangular matrix $H \in \mathbb{R}^{m \times m}$, the noise-shaping quantization scheme

$$\mathcal{Q}: \mathbb{R}^m \rightarrow \mathcal{A}^m, \quad y \mapsto q = y + Hu$$

is said to be stable if $\|u\|_\infty \leq C$ holds whenever $\|y\|_\infty \leq h < 1$. The upper bound C is referred to as the stability constant of the quantization scheme.

Examples of stable noise-shaping quantizers are the popular $\Sigma\Delta$ and distributed noise-shaping schemes, which we now describe.

For $r \in \mathbb{N}$, the standard r th-order $\Sigma\Delta$ quantization scheme computes a uniformly bounded

solution u to (2.3.1) where $H = D^r$ is the order r difference matrix, which is computed by taking powers of the $m \times m$ first-order difference matrix D :

$$D(j,k) := \begin{cases} 1 & \text{if } j = k, \\ -1 & \text{if } j = k + 1, \\ 0 & \text{otherwise} \end{cases}$$

for all $j, k = 1, \dots, m$. Constructing quantization functions that yield stable methods in this context is far from simple, but several constructions do exist (see, eg., [DD03, Gün03, DKG11]). In particular, one-bit $\Sigma\Delta$ schemes, where $\mathcal{A} = \{\pm 1\}$ and $\delta = 1$, are stable as long as $\|y\|_\infty \leq h < 1$ [KSW12].

Distributed noise-shaping quantization schemes are constructed in a way similar to $\Sigma\Delta$ schemes, albeit with a few key differences. To see this, let $\lambda := m/p$ be the oversampling rate (we assume it to be an integer). Then for any $\beta > 1$, the distributed noise-shaping quantization method computes a uniformly bounded solution u to (2.3.1) where $H = I_p \otimes H_\beta$ is the block-diagonal operator determined by the $\lambda \times \lambda$ matrix

$$H_\beta(j,k) := \begin{cases} 1 & \text{if } j = k, \\ -\beta & \text{if } j = k + 1, \\ 0 & \text{otherwise} \end{cases}$$

for all $j, k = 1, \dots, \lambda$. As in the $\Sigma\Delta$ case, ensuring that these quantization methods are stable is nontrivial. However, stable constructions do exist for Gaussian measurement models [CG16] as well as those drawn from BOEs and PCEs [HS18].

2.3.2 Binary embeddings

General embedding theory is often centered around the idea of obtaining low-dimensional representations of high-dimensional sets that preserve the geometric structure of the original set. The benefits of such embeddings stem from the dimensionality reduction (e.g., reduced storage space and computational time), and therefore play an important role in the signal processing and machine learning communities [HS18]. The first, and perhaps most important, example of low-distortion embeddings is the Johnson-Lindenstrauss lemma:

Lemma 2.3.2 (Johnson-Lindenstrauss Lemma [JL84]). *For every $\varepsilon \in (0, 1)$ and finite collection $T \subset \mathbb{R}^N$ of size $|T| = n$, if $m > 8\varepsilon^{-2} \log(n)$, then there exists a linear map $f: \mathbb{R}^N \rightarrow \mathbb{R}^m$ such that*

$$(1 - \varepsilon) \|u - v\|_2^2 \leq \|f(u) - f(v)\|_2^2 \leq (1 + \varepsilon) \|u - v\|_2^2$$

for all $u, v \in T$.

In words, Lemma 2.3.2 states that one can linearly embed any set of n points in \mathbb{R}^N into \mathbb{R}^m while preserving pairwise Euclidean distance up to ε -Lipschitz distortion, provided $m = O(\varepsilon^{-2} \log(n))$. The Johnson-Lindenstrauss Lemma extends to infinite subsets of \mathbb{R}^N by replacing cardinality with another measure of complexity, like Gaussian mean width [OR15]; moreover, it is known (see, e.g., [KW11]) that matrices with the RIP and randomized column signs provide such embeddings.

Binary embeddings, where the embedding space is instead taken to be $\{\pm 1\}^m$, have gained recent attention in the signal processing and machine learning communities (see, e.g., [WTF09, RL09, SH09, GLGP13]) due to the potential storage and computational benefits. Much effort has been focused on developing fast and efficient binary embeddings in the context of quantization (see [HS18] and references therein).

2.3.3 Geometric Multi-Resolution Analysis

Let $\mathcal{M} \subset \mathbb{R}^N$ be a smooth, compact d -dimensional manifold. In other words, \mathcal{M} is a compact subset of \mathbb{R}^N for which the following holds: There exists a collection $\{(U_\alpha, \phi_\alpha)\}_{\alpha \in \mathcal{S}}$ of subsets $U_\alpha \subset \mathcal{M}$ and homeomorphisms $\phi_\alpha: U_\alpha \rightarrow \mathbb{R}^d$ such that $\cup_{\alpha \in \mathcal{S}} U_\alpha = \mathcal{M}$ and the transition maps $\phi_\beta \circ \phi_\alpha^{-1}: \phi_\alpha(U_\alpha \cap U_\beta) \rightarrow \phi_\beta(U_\alpha \cap U_\beta)$ are C^∞ . The pairs (U_α, ϕ_α) are called *charts* and the collection $\{(U_\alpha, \phi_\alpha)\}_{\alpha \in \mathcal{S}}$ is called an *atlas*.

As previously mentioned, we will assume, as in [IKKSM18], that we do not know the structure of the manifold \mathcal{M} a priori, but instead have access to a structured dictionary model. Clearly, any solution to our signal recovery problem depends on the type of representation we have. In our results, we consider the case where the dictionary model for the manifold is provided by a Geometric Multi-Resolution Analysis (GMRA) approximation of \mathcal{M} [ACM12], which we make precise below, but first we must define a new geometric object: For any set $T \subset \mathbb{R}^N$ and constant $\rho > 0$, let

$$\text{tube}_\rho(T) := \{x \in \mathbb{R}^N : \inf_{y \in T} \|x - y\|_2 \leq \rho\}$$

denote the tube of radius ρ around T . With this definition in hand, we are ready to formally introduce the GMRA approximation:

Definition 2.3.3 (Geometric Multi-Resolution Analysis (GMRA) [IM13]). Let $J \in \mathbb{N}$ be nonzero and $\{K_j\}_{j=1}^J \subset \mathbb{N}$. A *GMRA approximation* of a smooth, compact d -dimensional manifold $\mathcal{M} \subset \mathbb{R}^N$ is a collection $\{(\mathcal{C}_j, \mathcal{P}_j)\}_{j=1}^J$ of centers $\mathcal{C}_j = \{c_{j,k}\}_{k=1}^{K_j}$ and affine projections

$$\mathcal{P}_j = \{P_{j,k}: \mathbb{R}^N \rightarrow \mathbb{R}^N : k \in \{1, \dots, K_j\}\}$$

with the following properties:

1. Affine Projections. Every $P_{j,k}$ is an orthogonal projection onto some d -dimensional affine

space which contains the center $c_{j,k}$.

2. Dyadic Structure. The number of centers at each level is bounded by $|\mathcal{C}_j| = K_j \leq C_C 2^{dj}$ for an absolute constant $C_C \geq 1$. Moreover, there exist $C_1 > 0, C_2 \in (0, 1]$ such that

- (a) $K_j \leq K_{j+1}$ for all $j \in \{1, \dots, J-1\}$,
- (b) $\|c_{j,k_1} - c_{j,k_2}\|_2 > C_1 2^{-j}$ for all $j \in \{1, \dots, J\}$ and $k_1 \neq k_2 \in \{1, \dots, K_j\}$,
- (c) For each $j \in \{1, \dots, J\}$ there exists a parent function $p_j: \{1, \dots, K_j\} \rightarrow \{1, \dots, K_{j-1}\}$ with

$$\|c_{j,k} - c_{j-1,p_j(k)}\|_2 \leq C_2 \min_{k' \in \{1, \dots, K_{j-1}\} \setminus \{p_j(k)\}} \|c_{j,k} - c_{j-1,k'}\|_2.$$

3. Multiscale Approximation. The projectors in \mathcal{P}_j approximate \mathcal{M} in the following sense:

- (a) There exists $j_0 \in \{1, \dots, J-1\}$ such that $c_{j,k} \in \text{tube}_{C_1 2^{-j-2}}(\mathcal{M})$ for all $j \geq j_0$ and $k \in \{1, \dots, K_j\}$.
- (b) For each $j \in \{1, \dots, J\}$ and $z \in \mathbb{R}^N$, let

$$c_{j,k_j(z)} \in \arg \min_{c_{j,k} \in \mathcal{C}_j} \|z - c_{j,k}\|_2. \quad (2.3.2)$$

Then for each $z \in \mathcal{M}$ there exists $C_z > 0$ so that $\|z - P_{j,k_j(z)}z\|_2 \leq C_z 2^{-2j}$ for all $j \in \{1, \dots, J\}$; moreover, for each $z \in K$ there exists $\tilde{C}_z > 0$ so that $\|z - P_{j,k'z}\|_2 \leq \tilde{C}_z 2^{-j}$ whenever $j \in \{1, \dots, J\}$ and $k' \in \{1, \dots, K_j\}$ satisfy

$$\|z - c_{j,k'}\|_2 \leq 16 \max \left\{ \|z - c_{j,k_j(z)}\|_2, C_1 2^{-j-1} \right\}.$$

Remark 2.3.4. By part (1) of Definition 2.3.3, a GMRA approximation of \mathcal{M} represents the manifold as a collection of points (the centers $c_{j,k}$) and corresponding low-dimensional affine spaces (defined by $P_{j,k}$). The refinement levels j control the accuracy of this approximation. Part

(2) of the definition states that the centers are organized in a tree-like structure, while part (3) characterizes the approximation accuracy at each refinement level. Note in particular, by part (3.a) the centers $c_{j,k}$ do not necessarily lie on \mathcal{M} , but cannot be too far away.

2.4 Problem Formulation and Main Result

Let $\mathcal{M} \subset (1 - \mu)B_2^N$ be a smooth, compact d -dimensional submanifold of the unit ℓ_2 -ball in \mathbb{R}^N for some $\mu \in (0, 1)$. We assume that we have access to a GMRA approximation $\{(\mathcal{C}_j, \mathcal{P}_{j,k})\}_j$ of \mathcal{M} , as in Definition 2.3.3. Define the scale- j GMRA approximation

$$\widehat{\mathcal{M}}_j := \{P_{j,k_j(z)}z : \|z\|_2 \leq 1\} \cap B_2^N. \quad (2.4.1)$$

We suppose that j_0 is large enough so that $\sup_{x \in \mathcal{M}} \tilde{C}_x 2^{-j_0} \leq \mu$ to ensure that

$$\{P_{j,k'}x : x \in \mathcal{M}, (j, k') \text{ as in part (3.b) of Definition 2.3.3}\} \subset B_2^N,$$

and further assume that $\text{tube}_{C_1 2^{-j_0-2}}(\mathcal{M}) \subset B_2^N$ which ensures $\mathcal{C}_j \subset \widehat{\mathcal{M}}_j$ for $j \geq j_0$. The number of measurements required for our theoretical guarantees to hold will depend on two notions of complexity of \mathcal{M} and the GMRA approximation, namely, its Gaussian mean width $w(\mathcal{M})$ and radius $\text{rad}(\mathcal{M})$. Finally, for $j \geq j_0$, we define $S := \mathcal{M} \cup \widehat{\mathcal{M}}_j$; essentially, this set defines an enlargement of \mathcal{M} within the ℓ_2 unit ball that will enable us to approximate $x \in \mathcal{M}$ using suitable points from the GMRA approximation of \mathcal{M} . Now, let \mathcal{Q} be either a stable r th order $\Sigma\Delta$ quantizer or stable distributed noise-shaping quantizer for $\beta > 1$ and associated alphabet \mathcal{A} . Let $x \in \mathcal{M}$, $A \in \mathbb{R}^{m \times N}$ be a standard Gaussian matrix (or a matrix drawn from a PCE or BOE), $D_\varepsilon \in \mathbb{R}^{N \times N}$ a diagonal matrix with random signs (independent of A) along the diagonal, $\Phi := AD_\varepsilon$ and $q := \mathcal{Q}(\Phi x)$. Our goal is the following:

Given q and Φ , approximate $x \in \mathcal{M}$ and show that the

associated error bounds decay fast as a function of m .

A useful fact (which we state below in its entirety) is that the binary embeddings provided by $\Sigma\Delta$ and distributed noise-shaping quantization approximately preserve Euclidean distance via a related pseudometric on the quantized vectors. This pseudometric, first defined in [HS18], is constructed by means of a *condensation operator* $V: \mathbb{R}^m \rightarrow \mathbb{R}^p$ in such a way that the state variable is controlled, i.e., $\|VHu\|_2$ is small (as a function of m). Given such a condensation operator, we define the pseudometric by

$$d_V(u, v) := \|V(u - v)\|_2 \quad (2.4.2)$$

for all $u, v \in \mathcal{A}^m$.

For the case of stable $\Sigma\Delta$ quantization, the condensation operator we will use is defined as follows: Let $\lambda := m/p =: r\tilde{\lambda} - r + 1$ for some integer $\tilde{\lambda}$ and consider the row vector $v_{\Sigma\Delta} \in \mathbb{R}^\lambda$ whose j th entry is the corresponding coefficient of the polynomial $(1 + z + \dots + z^{\tilde{\lambda}-1})^r$. Letting \otimes denote the Kronecker product and I_p the $p \times p$ identity matrix, we then define the $\Sigma\Delta$ condensation operator to be

$$\tilde{V}_{\Sigma\Delta} := \frac{1}{\|v_{\Sigma\Delta}\|_2 \sqrt{p}} I_p \otimes v_{\Sigma\Delta}, \quad (2.4.3)$$

which is a $p \times m$ matrix. The distributed noise-shaping condensation operator is defined in a similar way. Indeed, consider the row vector $v_\beta \in \mathbb{R}^\lambda$ whose j th entry is β^{-j} . Then the distributed noise-shaping condensation operator is defined by

$$\tilde{V}_\beta := \frac{1}{\|v_\beta\|_2 \sqrt{p}} I_p \otimes v_\beta. \quad (2.4.4)$$

The condensation operators $\tilde{V}_{\Sigma\Delta}$ and \tilde{V}_β interact with stable quantization schemes in a particularly nice way. To be explicit, we have the following results from [HS18]:

Lemma 2.4.1 (Lemma 4.5 in [HS18]). *For a stable r th order $\Sigma\Delta$ quantization scheme the condensation operator $\tilde{V}_{\Sigma\Delta}: \mathbb{R}^m \rightarrow \mathbb{R}^p$ satisfies $\|\tilde{V}_{\Sigma\Delta}D^r\|_{\infty \rightarrow 2} \leq (8r)^{r+1}\lambda^{-r+1/2}$.*

Lemma 2.4.2 ([HS18]). *For $\beta > 1$, the stable distributed noise-shaping condensation operator $\tilde{V}_\beta: \mathbb{R}^m \rightarrow \mathbb{R}^p$ satisfies $\|\tilde{V}_\beta H\|_{\infty \rightarrow 2} \leq \frac{9}{8}\beta^{-\lambda+1}$.*

Using Lemmas 2.4.1 and 2.4.2, T. Huynh and R. Saab were able to prove that the binary embeddings associated with $\Sigma\Delta$ and distributed noise-shaping quantization approximately preserve Euclidean distance via the pseudometric (2.4.2):

Theorem 2.4.3 (c.f. Theorem 5.2 in [HS18]). *Consider any set $T \subset B_2^N$. Fix $\alpha \in (0, 1)$ and let $\Phi := AD_\varepsilon$ for $A \in \mathbb{R}^{m \times N}$ a standard (sub)-Gaussian matrix, PCE (Definition 2.2.2), or BOE (Definition 2.2.3) and $D_\varepsilon \in \mathbb{R}^{N \times N}$ a diagonal matrix of random signs. Let $\mathcal{Q}: \mathbb{R}^m \rightarrow \sqrt{(1+\alpha)m}\{\pm 1\}$ be the stable quantization scheme corresponding to either r th order $\Sigma\Delta$ quantization or distributed noise-shaping quantization with $\beta \in (1, 10/9]$ and suppose $\lambda = m/p$. Finally, define \tilde{V} as in either (2.4.3) or (2.4.4) for each quantization method (resp.) and let $\gamma = \|\mathbf{v}_{\Sigma\Delta}\|_1 / \|\mathbf{v}_{\Sigma\Delta}\|_2$ or $\gamma = \|\mathbf{v}_\beta\|_1 / \|\mathbf{v}_\beta\|_2$, respectively. Now, fix $\nu > 0$ and suppose that*

$$p \geq c_1 \gamma^2 (1 + \nu)^2 \log^4(N) \frac{\max \left\{ 1, \frac{\omega^2(T-T)}{\text{rad}^2(T-T)} \right\}}{\alpha^2}$$

for some constant $c_1 > 0$. If $m \geq \chi(p)$, then with probability at least $1 - e^{-\nu}$, there exists a constant $c_2 > 0$ such that

$$\left| d_{\tilde{V}}(\mathcal{Q}(\Phi x), \mathcal{Q}(\Phi y)) - \|x - y\|_2 \right| \leq \max \{ \sqrt{\alpha}, \alpha \} \text{rad}(T - T) + c_2 \eta(\lambda)$$

holds for all $x, y \in T$. Here, $\chi(p) = p^{\frac{r-1/2}{r-1}}$ and $\eta(\lambda) = \lambda^{-r+1/2}$ for a stable $\Sigma\Delta$ quantization scheme, while $\chi(p) = p \log(m)$ and $\eta(\lambda) = \beta^{-\lambda+1}$ for distributed noise-shaping quantization.

We are now ready to state our recovery algorithm and its associated error guarantees.

Algorithm 1 Reconstruction Algorithm

Given: noise-shaping quantizer \mathcal{Q} ; measurements $q = \mathcal{Q}(\Phi x)$ for $x \in \mathcal{M}$; GMRA approximation $\{(c_{j,k}, P_{j,k})\}_{k=1}^{K_j}$ of \mathcal{M} at scale $j \geq j_0$; \tilde{V} as in (2.4.3) or (2.4.4).

Ensure: $x^\sharp \approx x$.

Step 1: Find $c_{j,k'} \in \arg \min_{c_{j,k} \in \mathcal{C}_j} \|\tilde{V}(\mathcal{Q}(\Phi c_{j,k}) - q)\|_2$.

Step 2: If $\|\tilde{V}(\mathcal{Q}(\Phi c_{j,k'}) - q)\|_2 = 0$, set $x^\sharp = c_{j,k'}$; else

$$x^\sharp = \arg \min_{z \in \mathbb{R}^N} \left\| \tilde{V}(\Phi z - q) \right\|_2 \quad \text{subject to } z = P_{j,k'}(z), \quad \|z\|_2 \leq 1.$$

Theorem 2.4.4. For $\mathcal{M} \subset (1 - \mu)B_2^N$, $\mu \in (0, 1)$ a smooth, compact d -dimensional manifold with GMRA approximation at scale $j \geq j_0$, let $\Phi := AD_\varepsilon$ for $A \in \mathbb{R}^{m \times N}$ a standard (sub)-Gaussian matrix, PCE (Definition 2.2.2), or BOE (Definition 2.2.3) and $D_\varepsilon \in \mathbb{R}^{N \times N}$ a diagonal matrix of random signs. Then there exist $\alpha \in (0, 1)$, $r \in \mathbb{N}$, and $\beta \in (1, 10/9]$ such that the following holds: Let $\mathcal{Q}: \mathbb{R}^m \rightarrow \sqrt{(1 + \alpha)m}\{\pm 1\}$ be the stable quantization scheme corresponding to either r th order $\Sigma\Delta$ quantization or distributed noise-shaping quantization, where $S := \mathcal{M} \cup \widehat{\mathcal{M}}_j$ is determined by (2.4.1), and suppose $\lambda = m/p$. Now, fix $\nu > 0$ and suppose that

$$p \gtrsim (1 + \nu)^2 \log^4(N) \text{rad}^4(S - S) \frac{\max\{1, \omega^2(S - S) \text{rad}^{-2}(S - S)\}}{\alpha^2}. \quad (2.4.5)$$

If $m \geq \chi(p)$, then with probability at least $1 - e^{-\nu}$, for all $x \in \mathcal{M}$, x^\sharp from Algorithm 1 satisfies

$$\|x^\sharp - x\|_2 \lesssim \tilde{C}_x 2^{-j} + \max\{\sqrt{\alpha}, \alpha\} \text{rad}(S - S) + \sqrt{(1 + \alpha)m} \eta(m/p).$$

Here, $\chi(p) = p^{\frac{r-1/2}{r-1}}$ and $\eta(\lambda) = \lambda^{-r+1/2}$ for a stable $\Sigma\Delta$ quantization scheme, while $\chi(p) = p \log(m)$ and $\eta(\lambda) = \beta^{-\lambda+1}$ for distributed noise-shaping quantization.

Remark 2.4.5. The error bound in Theorem 2.4.4 consists of three terms, each originating from a

different source:

$$\|x^\sharp - x\|_2 \lesssim \underbrace{\tilde{C}_x 2^{-j}}_{\text{GMRA error}} + \underbrace{\max\{\sqrt{\alpha}, \alpha\} \text{rad}(S - S)}_{\text{Manifold complexity}} + \underbrace{\sqrt{(1 + \alpha)m\eta(m/p)}}_{\text{Quantization error}}.$$

By appropriate choice of α and $\lambda = m/p$, the manifold complexity and quantization error terms may be made sufficiently small so that $\|x^\sharp - x\|_2 \lesssim \tilde{C}_x 2^{-j}$ (see [ILNS19] for details); for instance, in the case of $\Sigma\Delta$ quantization, it suffices to choose $\alpha \lesssim \text{rad}^{-2}(S - S)2^{-2(j+1)}$ and $\lambda \gtrsim j^2$ (hence, $r = O(j)$). Essentially, this implies that the accuracy of Algorithm 1 is limited almost entirely by the level of refinement in the GMRA approximation of \mathcal{M} .

Proof of Theorem 2.4.4. Fix $x \in \mathcal{M}$ and let k' be the index of the optimal center $c_{j,k'}$ found in Step 1 of Algorithm 1 and $P_{j,k'}$ the corresponding GMRA projection. By the triangle inequality, we then have

$$\|x^\sharp - x\|_2 \leq \|x^\sharp - P_{j,k'}x\|_2 + \|P_{j,k'}x - x\|_2. \quad (2.4.6)$$

To bound the first term on the right-hand side of (2.4.6), note that optimality of x^\sharp and feasibility of $P_{j,k'}x$ give us

$$\begin{aligned} 0 &\leq \left\| \tilde{V}(\Phi P_{j,k'}x - q) \right\|_2 - \left\| \tilde{V}(\Phi x^\sharp - q) \right\|_2 \\ &= \left\| \tilde{V}(\Phi P_{j,k'}x - q) \right\|_2 - \left\| \tilde{V}(\Phi x^\sharp - \Phi P_{j,k'}x + \Phi P_{j,k'}x - q) \right\|_2 \\ &\leq 2 \left\| \tilde{V}(\Phi P_{j,k'}x - q) \right\|_2 - \left\| \tilde{V}\Phi(x^\sharp - P_{j,k'}x) \right\|_2, \end{aligned}$$

where the final step is an application of the triangle inequality. Now, by the definition of S we have $x^\sharp, P_{j,k'}x \in S$ and, hence, by the proof of Theorem 2.4.3 (specifically, the fact that $\tilde{V}\Phi$ satisfies the

RIP; see equation (5.8) in [HS18]) we have

$$\left\| \tilde{V} \Phi(x^\sharp - P_{j,k'}x) \right\|_2 \geq \|x^\sharp - P_{j,k'}x\|_2 - \max \{ \sqrt{\alpha}, \alpha \} \text{rad}(S - S)$$

with probability exceeding $1 - e^{-\nu}$. Thus, it follows that

$$\|x^\sharp - P_{j,k'}x\|_2 \leq 2 \left\| \tilde{V}(\Phi P_{j,k'}x - q) \right\|_2 + \max \{ \sqrt{\alpha}, \alpha \} \text{rad}(S - S)$$

holds with probability at least $1 - e^{-\nu}$. Now, if we define $q^* := \mathcal{Q}(\Phi P_{j,k'}x)$, then

$$\left\| \tilde{V}(\Phi P_{j,k'}x - q) \right\|_2 \leq \left\| \tilde{V}(\Phi P_{j,k'}x - q^*) \right\|_2 + d_{\tilde{V}}(q, q^*).$$

Since the quantization scheme is stable, Lemmas 2.4.1 and 2.4.2 then enable us to bound $\left\| \tilde{V}(\Phi P_{j,k'}x - q^*) \right\|_2 \lesssim \eta(\lambda)$, while Theorem 2.4.3 implies

$$d_{\tilde{V}}(q, q^*) \leq \|x - P_{j,k'}x\| + \max \{ \sqrt{\alpha}, \alpha \} \text{rad}(S - S) + c_2 \sqrt{(1 + \alpha)m} \eta(\lambda)$$

with probability exceeding $1 - e^{-\nu}$. Hence, up to constant factors (depending on r or β), we have obtained

$$\|x^\sharp - P_{j,k'}x\|_2 \lesssim 2\|x - P_{j,k'}x\| + 3 \max \{ \sqrt{\alpha}, \alpha \} \text{rad}(S - S) + 2(1 + c_2 \sqrt{(1 + \alpha)m}) \eta(\lambda),$$

which combines with (2.4.6) to yield

$$\|x^\sharp - x\|_2 \lesssim \max \{ \sqrt{\alpha}, \alpha \} \text{rad}(S - S) + \sqrt{(1 + \alpha)m} \eta(\lambda) + \|P_{j,k'}x - x\|_2.$$

Therefore, to complete the proof we must show that $\|P_{j,k'}x - x\|_2 \lesssim \tilde{C}_x 2^{-j}$, which follows from

part (3.b) of Definition 2.3.3, provided we can show that

$$\|x - c_{j,k'}\|_2 \leq 16 \max \{ \|x - c_{j,k_j(x)}\|_2, C_1 \cdot 2^{-j-1} \} \quad (2.4.7)$$

holds with high probability. To this end observe that, by Theorem 2.4.3, with probability at least $1 - e^{-\nu}$ we have

$$|d_{\tilde{\nu}}(\mathcal{Q}(\Phi c_{j,k}), q) - \|c_{j,k} - x\|_2| \lesssim \max \{ \sqrt{\alpha}, \alpha \} \text{rad}(S - S) + \sqrt{(1 + \alpha)m} \eta(\lambda)$$

for all $c_{j,k} \in \mathcal{C}_j$. Since optimality of $c_{j,k'}$ in Algorithm 1 gives us

$$d_{\tilde{\nu}}(\mathcal{Q}(\Phi c_{j,k'}), q) \leq d_{\tilde{\nu}}(\mathcal{Q}(\Phi c_{j,k_j(x)}), q),$$

where $k_j(x)$ is as in (2.3.2), it therefore follows that

$$\|x - c_{j,k'}\|_2 \lesssim \|x - c_{j,k_j(x)}\|_2 + \max \{ \sqrt{\alpha}, \alpha \} \text{rad}(S - S) + \sqrt{(1 + \alpha)m} \eta(\lambda).$$

Hence, choosing α sufficiently small and r or β sufficiently large (so that $\eta(\lambda)$ is small) ensures (2.4.7) holds with probability at least $1 - e^{-\nu}$, as desired. \square

Remark 2.4.6. As Lemma 4.3 of [IKKSM18] shows, $\omega(S - S) \lesssim \omega(\mathcal{M}) + \sqrt{dj}$. This is a suitable bound for coarse GMRA scales, i.e. $j \lesssim \log(N)$. However, for $j \gtrsim \log(N)$ one can slightly modify the definition of S and use the bound $\omega(S - S) \lesssim (\omega(\mathcal{M}) + 1) \log(N)$ as proven in Lemma 4.5 of [IKKSM18], albeit this requires some modifications to the proof of Theorem 2.4.4. Please see Remark 4.15 of [IKKSM18] for more details.

2.5 Numerical Simulations

To simulate Algorithm 1, we take $\mathcal{M} = \mathbb{S}^2$ embedded in \mathbb{R}^{20} and construct a GMRA up to level $j_{\max} = 15$ using 20,000 data points sampled uniformly from \mathcal{M} . We randomly select a test set of 100 points $x \in \mathcal{M}$ for use throughout all experiments. In each experiment (i.e., point in Figure 2.5.1), compressed sensing measurements $y = \Phi x = m^{-1/2} A D_\varepsilon x$ are taken for each test point, with $A \sim \text{Norm}(0, I_{m \times N})$ and D_ε a diagonal $N \times N$ matrix of random ± 1 s. We recover x^\sharp from the r th order $\Sigma\Delta$ measurements $\mathcal{Q}(y)$ via Algorithm 1 where, for practical reasons, the alphabet from Theorem 2.4.4 is modified to be $\mathcal{A} = \{\pm 1\}$. We vary $\lambda = m/p$ for fixed r , p , and refinement scale j . The reconstruction error decays as a function of λ until reaching a floor due to the refinement level of the GMRA.




Figure 2.5.1: Log-scale plot of average relative reconstruction error from Algorithm 1 as a function of the oversampling rate $\lambda = m/p$ for $p = 10$. Solid lines correspond to GMRA refinement level $j = 12$; dashed lines to $j = 6$. Blue and red plots represent $r = 2, 4$ (resp.). For each j , reconstruction error decays as a function of λ until reaching a floor due to error in the GMRA approximation of \mathcal{M} .

2.6 Acknowledgements

A conference version of the material presented in this chapter appeared in the proceedings of the 17th Annual International Conference on Sampling Theory and Applications (SampTA), 2019. The dissertation author was the primary researcher and author of this material. Thanks are extended to coauthors Mark Iwen, Eric Lybrand, and Rayan Saab for their contributions and insights aiding in the development of this material.

Chapter 3

Random Vector Functional Link Networks

3.1 Introduction

In recent years, deep learning has triggered an increased interest in neural networks amongst researchers in the machine learning community. So-called deep neural networks model functions using a composition of multiple hidden layers, each transforming (possibly non-linearly) the previous layer(s) before building a final output representation. In machine learning parlance, these layers are determined by sets of *weights* and *biases* which, when adjusted appropriately, allow the network to mimic the action of more general functions. In this way, deep neural networks are fundamentally parametric functions whose parameters may be chosen using optimization techniques to minimize the difference between the network and the function it is intended to model. This difference is typically characterized using a finite set of input signals and their function evaluations (called *training data*); indeed, these function evaluations may be compared to the corresponding network outputs when evaluated on the same set of input signals, and the weights and biases then *learned* by minimizing a given loss function (e.g., sum-of-squares error, cross-entropy, etc.). Unfortunately, estimating a general unknown function using a deep neural network in this manner often requires learning thousands of weights and biases using gradient

descent-based algorithms such as back-propagation, which can be very time consuming, often get stuck in local minima, and can be very sensitive to the distribution of training data used [Sug18]. Moreover, deep neural networks can require massive amounts of training data, and so are typically unreliable for applications with very limited data availability, such as agriculture, healthcare, and ecology [OWB18].

To address some of the difficulties associated with training deep neural networks, both researchers and practitioners have attempted to incorporate randomness in some way. Indeed, randomization-based neural networks that yield closed form solutions typically require less time to train and avoid some of the pitfalls of traditional neural networks trained using back-propagation [Sug18, SKD92, BS95]. One of the popular randomization-based neural network architectures is the Random Vector Functional Link (RVFL) network [PT92, IP95], which is a single layer feed-forward neural network (SLFN) in which the input-to-hidden layer weights and biases are selected randomly and independently from a suitable domain and the remaining hidden-to-output layer weights are learned using training data. Although originally considered in the early- to mid-1990s [PT92, PPS94, IP95, PP95], RVFL networks have had much more recent success in several modern applications, including time-series data prediction [CW99], handwritten word recognition [PP00], visual tracking [ZS17b], signal classification [ZS17a, KSZ18], regression [VPM18], and forecasting [TWY18, DMSP18]. Deep neural network architectures based on RVFL networks have also made their way into more recent literature [HR18, KST19], although traditional, single layer RVFL networks tend to perform just as well as, and with lower training costs than, their multi-layer counterparts [KST19].

Although RVFL networks are proving their usefulness in practice, the supporting theoretical framework is currently lacking [ZWC⁺19]. Most theoretical research into the approximation capabilities of deep neural networks centers around two main concepts: universal approximation of functions on compact domains and point-wise approximation on finite training sets [HZS06]. For instance, in the early 1990s it was shown that multi-layer feed-forward neural networks having

activation functions which are continuous, bounded, and non-constant are universal approximators (in the L^p sense for $1 \leq p < \infty$) of continuous functions on compact domains [Hor91]; this result was later improved to include non-polynomial activation functions [LLPS93]. Likewise, it is known that m distinct observations can be learned with zero training error using SLFNs with at most $n = m$ hidden nodes and (almost) any bounded, non-linear activation function [GB98] or RVFL networks with at most $n = m$ nodes and smooth activation function [HZS06]. The most notable result in the existing literature regarding the universal approximation capability of RVFL networks is due to B. Igelnik and Y.H. Pao in the mid-1990s, who showed that such neural networks can universally approximate continuous functions on compact sets [IP95]; the noticeable lack of results since has left a sizable gap between theory and practice. In this paper, we begin to bridge this gap, bringing the mathematical theory behind RVFL networks into the modern spotlight. Our contributions are as follows: First, we provide a rigorous proof of the original Igelnik and Pao result, that is, that RVFL networks with activation functions which are both absolutely and square integrable are universal approximators (on average) of continuous functions with compact support, provided the number of nodes n is allowed to be infinite [IP95]. Our proof of this result corrects several mistakes from the original paper and formally structures the proof technique in a way that is readily adaptable to other settings. Second, we prove a non-asymptotic version of the original Igelnik and Pao result, showing that RVFL networks having absolutely and square integrable activation functions universally approximate (with high probability) continuous functions with compact support, provided the number of nodes n is large enough (but finite!). Next, we propose a new RVFL network architecture for approximating continuous functions defined on smooth, compact manifolds and show that the original Igelnik and Pao result may be adapted to this setting. Finally, we prove an analogous, probabilistic result for our RVFL network architecture on manifolds.

3.2 Background and notation

In this section, we provide the necessary mathematical background for Random Vector Functional Link (RVFL) networks, as well as some notation that will be used throughout.

3.2.1 Randomized single layer neural networks

Consider a single layer feed-forward neural network (SLFN) with n nodes, which may be regarded as a parametric function $f_n: \mathbb{R}^N \rightarrow \mathbb{R}$ of the form

$$f_n(x) = \sum_{k=1}^n v_k \rho(\langle w_k, x \rangle + b_k), \quad x \in \mathbb{R}^N.$$

Here, the function $\rho: \mathbb{R} \rightarrow \mathbb{R}$ is called an activation function and is potentially non-linear; see Table 3.2.1 for some typical examples used in practical applications. The parameters of the SLFN are the number of nodes $n \in \mathbb{N}$ in the hidden layer, the input-to-hidden layer weights and biases $\{w_k\}_{k=1}^n \subset \mathbb{R}^N$ and $\{b_k\}_{k=1}^n \subset \mathbb{R}$ (resp.), and the hidden-to-output layer weights $\{v_k\}_{k=1}^n \subset \mathbb{R}$. Such neural networks are often used in supervised learning, where the network parameters are learned from training data in order to approximate an unknown function on a given domain. Specifically, given an unknown function $f: \mathbb{R}^N \rightarrow \mathbb{R}$ and a training set $\{(x_k, f(x_k))\}_{k=1}^m$ for some $m \in \mathbb{N}$, one seeks to optimize the parameters of f_n in such a way that $f_n \approx f$. Often the means of measuring approximation error in this setting is via a loss function $\mathcal{L}(x_1, \dots, x_k)$; indeed, a typical loss function is the sum-of-squares error

$$\mathcal{L}(x_1, \dots, x_k) = \frac{1}{m} \sum_{k=1}^m |f(x_k) - f_n(x_k)|^2.$$

The SLFN which approximates f is then determined using an optimization algorithm, such as back-propagation, to find the network parameters which minimize $\mathcal{L}(x_1, \dots, x_k)$. It is known that there exist weights and biases which make the loss function vanish when the number of nodes n

Table 3.2.1: List of typical activation functions for neural networks.

Sigmoid	$\rho(z) = \frac{1}{1+\exp(-z)}$
ReLU	$\rho(z) = \max\{0, z\}$
Sine	$\rho(z) = \sin(z)$
Hardlim	$\rho(z) = \frac{1}{2}(1 + \sin(z))$
Tribas	$\rho(z) = \max\{0, 1 + z \}$
Radbas	$\rho(z) = \exp(-z^2)$
Sign	$\rho(z) = \text{sign}(z)$

is at most m , provided the activation function is bounded, non-linear, and has at least one finite limit at either $\pm\infty$ [GB98].

Unfortunately, optimizing the parameters in SLFNs can be difficult. For instance, any non-linearity in the activation function can cause back-propagation to be very time consuming or get caught in local minima of the loss function [Sug18]. RVFL networks, which are SLFNs where the input-to-hidden layer weights and biases are chosen at random, are designed to avoid these difficulties. Indeed, by eliminating the need to optimize the input-to-hidden layer weights and biases, RVFL networks turn supervised learning into a purely linear problem. To see this, define $\rho(X) \in \mathbb{R}^{n \times m}$ to be the matrix whose j th column is $\{\rho(\langle w_k, x_j \rangle + b_k)\}_{k=1}^n$ and $f(X) \in \mathbb{R}^m$ the vector whose j th entry is $f(x_j)$. Then the vector $v \in \mathbb{R}^n$ of hidden-to-output layer weights is the solution to the matrix-vector equation $f(X) = \rho(X)^T v$, which can be solved by computing the Moore-Penrose pseudoinverse of $\rho(X)^T$. In fact, there exist weights and biases which make the loss function vanish when the number of nodes n is at most m , provided the activation function is smooth [HZS06].

In this paper, we study the uniform approximation capabilities of RVFL networks, specifically, the problem of estimating a continuous, compactly supported function on N -dimensional Euclidean space. Few theoretical results exist in this area; indeed, the most notable result is due to Igel'nik and Pao [IP95], who proved that any $f \in C_c(\mathbb{R}^N)$ may be approximated uniformly (in the sense of mean-square error) by RVFL networks whose activation function satisfies $\rho \in L^1(\mathbb{R}) \cap L^2(\mathbb{R})$, provided the number of nodes n is allowed to be infinite. Although this

result is indeed true, the original proof presented in [IP95] is unclear and not mathematically rigorous. Moreover, the results of [IP95] are both asymptotic in the number of nodes n and highly dependent on the ambient dimension N . In practice, one would prefer *non-asymptotic* results that quantify the relationship between the approximation error and the size of the neural network, as well as take advantage of any lower-dimensional structure exhibited by the function f . We address both of these issues, providing non-asymptotic error bounds for RVFL networks that depend on the complexity of f through the intrinsic dimension of its domain.

3.2.2 Notation

For a function $f: \mathbb{R}^N \rightarrow \mathbb{R}$, the set $\text{supp}(f) \subset \mathbb{R}^N$ denotes the support of f . We denote by $C_c(\mathbb{R}^N)$ and $C_0(\mathbb{R}^N)$ the classes of continuous functions mapping \mathbb{R}^N to \mathbb{R} whose support sets are compact and vanish at infinity, respectively. Given a set $S \subset \mathbb{R}^N$, we define its radius to be $\text{rad}(S) := \sup_{x \in S} \|x\|_2$; moreover, if $d\mu$ denotes the uniform volume measure on S , then we write $\text{vol}(S) := \int_S d\mu$ to represent the volume of S . For any probability distribution $P: \mathbb{R}^N \rightarrow [0, 1]$, a random variable X distributed according to P is denoted by $X \sim P$, and we write its expectation as $\mathbb{E}X := \int_{\mathbb{R}^N} X dP$. The open ℓ_p ball of radius $r > 0$ centered at $x \in \mathbb{R}^N$ is denoted by $B_p^N(x, r)$ for all $1 \leq p \leq \infty$; the ℓ_p unit-ball centered at the origin is abbreviated B_p^N . Given a fixed $\delta > 0$ and a set $S \subset \mathbb{R}^N$, a minimal δ -net for S , which we denote $\mathcal{C}(\delta, S)$ is the smallest subset of S satisfying $S \subset \cup_{x \in \mathcal{C}(\delta, S)} B_2^N(x, \delta)$; the δ -covering number of S is the cardinality of a minimal δ -net for S and is denoted $\mathcal{N}(\delta, S) := |\mathcal{C}(\delta, S)|$.

3.3 Theoretical results in Euclidean space

The first theoretical result for RVFL networks, due to Igelnik and Pao, guarantees that continuous functions can be universally approximated on compact sets using RVFL networks, provided the number of nodes $n \in \mathbb{N}$ in the network is allowed to go to infinity [IP95]. Moreover,

it shows that the mean square error of the approximation vanishes at a rate proportional to $1/n$. At the time, this result was state-of-the-art and justified how RVFL networks were used in practice. However, the original theorem, although correct in spirit, is not technically correct. In fact, several aspects of the proof technique are flawed. Some of the minor flaws are mentioned in [JCIY97], but the subsequent revisions do not address the more major issues that we tackle here. Thus, our first contribution to the theory of RVFL networks is a corrected version of the original Igelnik and Pao theorem:

Theorem 3.3.1 (Igelnik and Pao, 1995). *Let $f \in C_c(\mathbb{R}^N)$ with $K := \text{supp}(f)$ and fix any activation function $\rho \in L^1(\mathbb{R}) \cap L^2(\mathbb{R})$. For any $\varepsilon > 0$, there exist constants $\alpha, \Omega > 0$ and hidden-to-output layer weights $\{v_k\}_{k=1}^n \subset \mathbb{R}$ such that the following holds: If*

$$w_0 \sim \text{Unif}([- \alpha \Omega, \alpha \Omega])^N;$$

$$y_0 \sim \text{Unif}(K);$$

$$u_0 \sim \text{Unif}([- \frac{\pi}{2}(2L+1), \frac{\pi}{2}(2L+1)]), \quad \text{where } L := \lceil \frac{2N}{\pi} \text{rad}(K)\Omega - \frac{1}{2} \rceil;$$

$$b_0 := -\langle w_0, y_0 \rangle - \alpha u_0,$$

and one chooses $\{w_k\}_{k=1}^n, \{b_k\}_{k=1}^n$ as independent draws from the distributions of w_0 and b_0 , respectively, then the sequence of RVFL networks $\{f_n\}_{n=1}^\infty$ defined by

$$f_n(x) := \sum_{k=1}^n v_k \rho(\langle w_k, x \rangle + b_k) \quad \text{for } x \in K$$

satisfies

$$\lim_{n \rightarrow \infty} \mathbb{E} \int_K |f(x) - f_n(x)|^2 dx < \varepsilon,$$

with convergence rate $O(1/n)$.

Remark 3.3.2. We note that, unlike the original theorem statement in [IP95], Theorem 3.3.1 does not show exact convergence of the sequence of constructed RVFL networks f_n to the original function f . Indeed, it only ensures that the limit f_n is ε -close to f . This should still be sufficient for practical applications since, given a desired accuracy level $\varepsilon > 0$, one can find values of α, Ω, n such that this accuracy level is achieved on average. Exact convergence can be proved if one replaces α and Ω by sequences $\{\alpha_n\}_{n=1}^\infty$ and $\{\Omega_n\}_{n=1}^\infty$ of positive numbers, both tending to infinity with n . In this setting, however, there is no guaranteed rate of convergence; moreover, as n increases, the ranges of the random variables $\{w_k\}_{k=1}^n$ and $\{u_k\}_{k=1}^n$ become increasingly larger, which may cause problems in practical applications.

The following simple corollary, which increases the class of activation functions one may use when constructing RVFL networks, also first appeared in [IP95]. We state the corollary here for completeness:

Corollary 3.3.3 (Igelnik and Pao, 1995). *Let $f \in C_c(\mathbb{R}^N)$ with $K := \text{supp}(f)$ and fix any differentiable activation function $\rho: \mathbb{R} \rightarrow \mathbb{R}$ such that $\rho' \in L^1(\mathbb{R}) \cap L^2(\mathbb{R})$. For any $\varepsilon > 0$, there exist constants $\alpha, \Omega > 0$ and hidden-to-output layer weights $\{v_k\}_{k=1}^n \subset \mathbb{R}$ such that the following holds: If*

$$w_0 \sim \text{Unif}([- \alpha \Omega, \alpha \Omega])^N;$$

$$y_0 \sim \text{Unif}(K);$$

$$u_0 \sim \text{Unif}([- \frac{\pi}{2}(2L+1), \frac{\pi}{2}(2L+1)]), \quad \text{where } L := \lceil \frac{2N}{\pi} \text{rad}(K)\Omega - \frac{1}{2} \rceil;$$

$$b_0 := -\langle w_0, y_0 \rangle - \alpha u_0,$$

and one chooses $\{w_k\}_{k=1}^n, \{b_k\}_{k=1}^n$ as independent draws from the distributions of w_0 and b_0 ,

respectively, then the sequence of RVFL networks $\{f_n\}_{n=1}^\infty$ defined by

$$f_n(x) := \sum_{k=1}^n v_k \rho(\langle w_k, x \rangle + b_k) \quad \text{for } x \in K$$

satisfies

$$\lim_{n \rightarrow \infty} \mathbb{E} \int_K |f(x) - f_n(x)|^2 dx < \varepsilon,$$

with convergence rate $O(1/n)$.

One of the drawbacks of Theorem 3.3.1 is that the mean square error guarantee is asymptotic in the number of nodes used in the neural network. This is clearly impractical for applications, and so it is desirable to have a more explicit error bound for each fixed number n of nodes used. To this end, we introduce a new, non-asymptotic version of Theorem 3.3.1, which provides an error guarantee with high probability whenever the number of network nodes is large enough, albeit at the price of an additional Lipschitz requirement on the activation function:

Theorem 3.3.4. *Let $f \in C_c(\mathbb{R}^N)$ with $K := \text{supp}(f)$ and fix any activation function $\rho \in L^1(\mathbb{R}) \cap L^2(\mathbb{R})$. Suppose further that ρ is κ -Lipschitz on \mathbb{R} for some $\kappa > 0$. For any $\varepsilon > 0$ and $\eta \in (0, 1)$, there exist constants $\alpha, \Omega > 0$ and hidden-to-output layer weights $\{v_k\}_{k=1}^n \subset \mathbb{R}$ such that the following holds: Suppose*

$$w_0 \sim \text{Unif}([- \alpha \Omega, \alpha \Omega])^N;$$

$$y_0 \sim \text{Unif}(K);$$

$$u_0 \sim \text{Unif}([- \frac{\pi}{2}(2L+1), \frac{\pi}{2}(2L+1)]), \quad \text{where } L := \lceil \frac{2N}{\pi} \text{rad}(K)\Omega - \frac{1}{2} \rceil;$$

$$b_0 := -\langle w_0, y_0 \rangle - \alpha u_0,$$

and one chooses $\{w_k\}_{k=1}^n, \{b_k\}_{k=1}^n$ as independent draws from the distributions of w_0 and b_0 ,

respectively. For any

$$0 < \delta < \frac{\sqrt{\varepsilon}}{4\sqrt{N}\kappa\alpha^2 M\Omega^{N+2}\text{vol}^{3/2}(K)(1+2N\text{rad}(K))},$$

if one chooses

$$n \geq \frac{2\sqrt{2\text{vol}(K)}Cc \log(3\eta^{-1}\mathcal{N}(\delta, K))}{\sqrt{\varepsilon} \log\left(1 + \frac{C\sqrt{\varepsilon}}{4\sqrt{2N}(2\Omega)^{N+1}\text{rad}(K)\text{vol}^{5/2}(K)\Sigma}\right)},$$

where $M := \sup_{x \in K} |f(x)|$, $c > 0$ is a numerical constant, and C, Σ are constants depending on f and ρ , then the RVFL network defined by

$$f_n(x) := \sum_{k=1}^n v_k \rho(\langle w_k, x \rangle + b_k) \quad \text{for } x \in K$$

satisfies

$$\int_K |f(x) - f_n(x)|^2 dx < \varepsilon$$

with probability at least $1 - \eta$.

The implication of Theorem 3.3.4 is that, given a desired accuracy level $\varepsilon > 0$, one can construct a RVFL network f_n that is ε -close to f with high probability, provided the number of nodes n in the neural network is sufficiently large. In fact, using the coarse estimates

$$\delta \lesssim \sqrt{\frac{\varepsilon}{\text{vol}(K)}} \quad \text{and} \quad n \gtrsim \frac{\log(\mathcal{N}(\delta, K))}{\sqrt{\varepsilon/\text{vol}(K)} \log\left(1 + \sqrt{\varepsilon/\text{vol}(K)}\right)},$$

along with the fact that $\log(1+x) = x + O(x^2)$ for small values of x , the requirement on the

number of nodes behaves like

$$n \gtrsim \frac{\text{vol}(K)}{\varepsilon} \log(\mathcal{N}(\sqrt{\varepsilon/\text{vol}(K)}, K))$$

whenever ε is sufficiently small. Using a simple bound on the covering number, this yields a coarse estimate of $n \gtrsim N\text{vol}(K)\varepsilon^{-1} \log(\text{vol}(K)/\varepsilon)$.

3.4 Preliminaries

In this section, we briefly introduce supporting material and theoretical results which we will need in later sections. This material is far from exhaustive, and is meant to be a survey of definitions, concepts, and key results.

3.4.1 A concentration bound for classic Monte Carlo integration

A crucial piece of the proof technique employed in [IP95], which we will use repeatedly, is the use of the Monte-Carlo method to approximate high-dimensional integrals. As such, we require a basic understanding of Monte-Carlo integration. The following introduction is adapted from the background material in [DKS13].

Let $f: \mathbb{R}^N \rightarrow \mathbb{R}$ and $S \subset \mathbb{R}^N$ a compact set. Suppose we want to estimate the integral $I(f, S) := \int_S f d\mu$, where μ is the uniform measure on S . The classic Monte Carlo method does this by an equal-weight cubature rule,

$$I_n(f, S) := \frac{\text{vol}(S)}{n} \sum_{j=1}^n f(x_j),$$

where $\{x_j\}_{j=1}^n$ are independent identically distributed uniform random samples from S and

$\text{vol}(S) := \int_S d\mu$ is the volume of S . In particular, note that $\mathbb{E}I_n(f, S) = I(f, S)$ and

$$\mathbb{E}I_n(f, S)^2 = \frac{1}{n}(\text{vol}(S)I(f^2, S) + (n-1)I(f, S)^2);$$

if we define the quantity

$$\sigma(f, S)^2 := \frac{I(f^2, S)}{\text{vol}(S)} - \frac{I(f, S)^2}{\text{vol}^2(S)}, \quad (3.4.1)$$

it therefore follows that the random variable $I_n(f)$ has mean $I(f, S)$ and variance $\text{vol}^2(S)\sigma(f)^2/n$.

Hence, by the Central Limit Theorem, provided that $0 < \text{vol}^2(S)\sigma(f, S)^2 < \infty$, we have

$$\lim_{n \rightarrow \infty} \mathbb{P}(|I_n(f, S) - I(f, S)| \leq C\varepsilon(f)) = (2\pi)^{-1/2} \int_{-C}^C e^{-x^2/2} dx$$

for any constant $C > 0$. This yields the following well-known result:

Theorem 3.4.1. *For any $f \in L^2(S, \mu)$, the mean-square error of the Monte Carlo approximation $I_n(f, S)$ satisfies*

$$\mathbb{E}|I_n(f, S) - I(f, S)|^2 = \frac{\text{vol}^2(S)\sigma(f, S)^2}{n},$$

where the expectation is taken with respect to the random variables $\{x_j\}_{j=1}^n$ and $\sigma(f, S)$ is defined in (3.4.1).

In particular, Theorem 3.4.1 implies that

$$\lim_{n \rightarrow \infty} \mathbb{E}|I_n(f, S) - I(f, S)|^2 = 0,$$

with convergence at a rate $O(1/n)$.

In the non-asymptotic setting, an interesting question is how to obtain a useful bound

on the probability $\mathbb{P}(|I_n(f, S) - I(f, S)| \geq t)$ for all $t > 0$. To this end, we now briefly recall a concentration result, which is a generalization of Bennett's inequality:

Theorem 3.4.2 (Theorem 7.6 in [Led01]; see also [Mas98, Tal96]). *Let $F: \mathbb{R}^N \rightarrow \mathbb{R}$ be a measurable function and $\{X_k\}_{k=1}^n \subset \mathbb{R}^N$ be independent identically distributed random variables. Set*

$$Z := \sum_{k=1}^n F(X_k) \quad \text{or} \quad Z := \left| \sum_{k=1}^n F(X_k) \right|$$

and assume there exists a constant $K > 0$ such that $|F| \leq K$ almost surely. Then for every $t > 0$ we have

$$\mathbb{P}(|Z - \mathbb{E}Z| \geq t) \leq 3 \exp\left(-\frac{t}{CK} \log\left(1 + \frac{Kt}{\mathbb{E}\Sigma^2}\right)\right),$$

where $\Sigma^2 := \sum_{k=1}^n F(X_k)^2$ and $C > 0$ is a universal constant.

To apply Theorem 3.4.2 in the classic Monte Carlo setting, we consider the function $F(x) := \text{vol}(S)f(x) - I(f, S)$, so that

$$Z = \sum_{j=1}^n F(x_j) = \sum_{j=1}^n (\text{vol}(S)f(x_j) - I(f, S)) = n(I_n(f, S) - I(f, S)).$$

Observing that $\mathbb{E}F(x_j) = \text{vol}(S)\mathbb{E}f(x_j) - I(f, S) = 0$ and

$$\begin{aligned} \mathbb{E}F(x_j)^2 &= \text{vol}^2(S)\mathbb{E}f(x_j)^2 - 2\text{vol}(S)I(f, S)\mathbb{E}f(x_j) + I(f, S)^2 \\ &= \text{vol}(S)I(f^2, S) - I(f, S)^2 \\ &= \text{vol}^2(S)\sigma(f, S)^2 \end{aligned}$$

for each $j = 1, \dots, n$, as well as $\mathbb{E}Z = n(\mathbb{E}I_n(f, S) - I(f, S)) = 0$, we obtain (via Theorem 3.4.2)

$$\mathbb{P}(n|I_n(f, S) - I(f, S)| \geq t) \leq 3 \exp\left(-\frac{t}{CK} \log\left(1 + \frac{Kt}{n \text{vol}^2(S) \sigma(f, S)^2}\right)\right)$$

for all $t > 0$, where we take

$$K := \sup_{x \in S} |\text{vol}(S)f(x) - I(f, S)|,$$

which is finite by assumptions on S and f . In this way, we obtain the following concentration inequality for the random variable $I_n(f, S)$:

Lemma 3.4.3. *For any $f \in L^2(S)$ and $n \in \mathbb{N}$ we have*

$$\mathbb{P}(|I_n(f, S) - I(f, S)| \geq t) \leq 3 \exp\left(-\frac{nt}{CK} \log\left(1 + \frac{Kt}{\text{vol}^2(S) \sigma(f)^2}\right)\right)$$

for all $t > 0$, provided $|\text{vol}(S)f(x) - I(f, S)| \leq K$ for almost every $x \in S$, where $C > 0$ is a universal constant.

3.4.2 Smooth, compact manifolds in Euclidean space

In this section we review several concepts of smooth manifolds that will be useful to us later. Many of the definitions and results that follow can be found, for instance, in [SCC18]. Let $\mathcal{M} \subset \mathbb{R}^N$ be a smooth, compact d -dimensional manifold. A *chart* for \mathcal{M} is a pair (U, ϕ) such that $U \subset \mathcal{M}$ is an open set and $\phi: U \rightarrow \mathbb{R}^d$ is a homeomorphism. One way to interpret a chart is as a tangent plane at some point $x \in U$; in this way, a chart defines a Euclidean coordinate system on U via the map ϕ . A collection $\{(U_j, \phi_j)\}_{j \in J}$ of charts defines an *atlas* for \mathcal{M} if $\cup_{j \in J} U_j = \mathcal{M}$.

We now define a special collection of functions on \mathcal{M} called a *partition of unity*:

Definition 3.4.4. Let $\mathcal{M} \subset \mathbb{R}^N$ be a smooth manifold. A *partition of unity* of \mathcal{M} with respect to

an open cover $\{U_j\}_{j \in J}$ of \mathcal{M} is a family of nonnegative smooth functions $\{\eta_j\}_{j \in J}$ such that for every $x \in \mathcal{M}$ we have $1 = \sum_{j \in J} \eta_j(x)$ and, for every $j \in J$, $\text{supp}(\eta_j) \subset U_j$.

It is known (see, e.g., [Tu10]) that if \mathcal{M} is compact there exists a partition of unity of \mathcal{M} such that $\text{supp}(\eta_j)$ is compact for all $j \in J$. In particular, such a partition of unity exists for any open cover of \mathcal{M} corresponding to an atlas.

Fix an atlas $\{(U_j, \phi_j)\}_{j \in J}$ for \mathcal{M} , as well as the corresponding, compactly supported partition of unity $\{\eta_j\}_{j \in J}$. Then we have the following, useful result (see, e.g., Lemma 4.8 in [SCC18]):

Lemma 3.4.5. *Let $\mathcal{M} \subset \mathbb{R}^N$ be a smooth, compact manifold with atlas $\{(U_j, \phi_j)\}_{j \in J}$ and compactly supported partition of unity $\{\eta_j\}_{j \in J}$. For any $f \in C(\mathcal{M})$ we have*

$$f(x) = \sum_{\{j \in J: x \in U_j\}} (\hat{f}_j \circ \phi_j)(x)$$

for all $x \in \mathcal{M}$, where

$$\hat{f}_j(z) := \begin{cases} f(\phi_j^{-1}(z)) \eta_j(\phi_j^{-1}(z)) & z \in \phi_j(U_j) \\ 0 & \text{otherwise.} \end{cases}$$

In later sections, it will be useful to know how to use the representation of Lemma 3.4.5 to integrate functions $f \in C(\mathcal{M})$ over \mathcal{M} . To this end, for each $j \in J$, let $D\phi_j(y)$ denote the differential of ϕ_j at $y \in U_j$, which is a map from the tangent space $T_y\mathcal{M}$ into \mathbb{R}^d . In particular, one may interpret $D\phi_j(y)$ as the matrix representation of a basis for the cotangent space at $y \in U_j$. As a result, $D\phi_j(y)$ is necessarily invertible for each $y \in U_j$, and so we know that $|\det(D\phi_j(y))| > 0$ for each $y \in U_j$. Hence, it follows by the change of variables theorem that

$$\int_{\mathcal{M}} f(x) dx = \int_{\mathcal{M}} \sum_{\{j \in J: x \in U_j\}} (\hat{f}_j \circ \phi_j)(x) dx = \sum_{j \in J} \int_{\phi_j(U_j)} \frac{\hat{f}_j(z)}{|\det(D\phi_j(\phi_j^{-1}(z)))|} dz. \quad (3.4.2)$$

3.5 Proof of Theorem 3.3.1

Let $f \in C_c(\mathbb{R}^N)$ with $K := \text{supp}(f)$ and suppose $\varepsilon > 0$ is fixed. Take $\rho \in L^1(\mathbb{R}) \cap L^2(\mathbb{R})$ arbitrarily. We wish to show that there exists a sequence of RVFL networks $\{f_n\}_{n=1}^\infty$ defined on K which satisfy the asymptotic error bound

$$\lim_{n \rightarrow \infty} \mathbb{E} \int_K |f(x) - f_n(x)|^2 dx < \varepsilon.$$

The proof technique we use is based on that introduced by Igel'nik and Pao, and consists of two approximation steps. First, the function f is approximated by an integral over the parameter space using a convolution identity from functional analysis. Then, this integral approximation is again approximated using a linear combination of random realizations of the activation function ρ via the Monte Carlo method. For clarity of presentation, we further break down the proof into four main parts (Sections 3.5.1- 3.5.4), each building upon the previous steps until the proof is complete.

3.5.1 A convolution identity

The first step in the proof of Theorem 3.3.1 is to represent f using a special convolution identity. To this end, we assume without loss of generality that $\rho \in L^1(\mathbb{R})$ such that $\int_{\mathbb{R}} g(x) dx = 1$ and consider the function $h_w: \mathbb{R}^N \rightarrow \mathbb{R}$ defined by

$$h_w(y) := \prod_{j=1}^N w(j) \rho(w(j)y(j)) \tag{3.5.1}$$

for all $y, w \in \mathbb{R}^N$. Observe that h_w may be viewed as a multidimensional bump function formed by taking Cartesian products of ρ ; indeed, the parameter $w \in \mathbb{R}^N$ controls the width of the bump in each of the N coordinate directions. In particular, if each coordinate of w is allowed to grow very large, then h_w becomes very localized near the origin. Objects that behave in this way are

known in the functional analysis literature as approximate δ -functions:

Definition 3.5.1. A sequence of functions $\{\varphi_t\}_{t>0} \subset L^1(\mathbb{R}^N)$ are called *approximate* (or *nascent*) δ -functions if

$$\lim_{t \rightarrow \infty} \int_{\mathbb{R}^N} \varphi_t(x) f(x) dx = f(0) \quad (3.5.2)$$

for all $f \in C_c(\mathbb{R}^N)$. For such functions, we write $\delta_0(x) = \lim_{t \rightarrow \infty} \varphi_t(x)$ for all $x \in \mathbb{R}^N$, where δ_0 denotes the N -dimensional Dirac δ -function centered at the origin.

Given $\varphi \in L^1(\mathbb{R}^N)$ with $\int_{\mathbb{R}^N} \varphi(x) dx = 1$, one may construct approximate δ -functions for $t > 0$ by defining $\varphi_t(x) := t^N \varphi(tx)$ for all $x \in \mathbb{R}^N$ [SW71]. Such sequences of approximate δ -functions are also called *approximate identity sequences* [Rud91] since they satisfy a particularly nice identity with respect to convolution, namely, $\lim_{t \rightarrow \infty} \|f * \varphi_t - f\|_1 = 0$ for all $f \in C_c(\mathbb{R}^N)$ (see, e.g., [Rud91, Theorem 6.32]). In fact, such an identity holds much more generally:

Lemma 3.5.2. [SW71, Theorem 1.18] Let $\varphi \in L^1(\mathbb{R}^N)$ with $\int_{\mathbb{R}^N} \varphi(x) dx = 1$ and for $t > 0$ define $\varphi_t(x) := t^N \varphi(tx)$ for all $x \in \mathbb{R}^N$. If $f \in L^p(\mathbb{R}^N)$ for $1 \leq p < \infty$ (or $f \in C_0(\mathbb{R}^N) \subset L^\infty(\mathbb{R}^N)$ for $p = \infty$), then $\lim_{t \rightarrow \infty} \|f * \varphi_t - f\|_p = 0$.

Motivated by this result, it is reasonable to suppose that the function h_w satisfies a similar identity. In particular, for any $f \in C_0(\mathbb{R}^N)$ one might suspect that

$$f(x) = \lim_{|w| \rightarrow \infty} (f * h_w)(x) \quad (3.5.3)$$

holds uniformly for all $x \in \mathbb{R}^N$; here, we write $\lim_{|w| \rightarrow \infty}$ to mean the limit as each coordinate $\{w(j)\}_{j=1}^N$ grows to infinity simultaneously. To prove (3.5.3), it would suffice to have $\lim_{|w| \rightarrow \infty} \|f * h_w - f\|_\infty = 0$ for all $f \in C_0(\mathbb{R}^N)$; indeed, since convolutions of $L^1(\mathbb{R}^N)$ and $L^\infty(\mathbb{R}^N)$ functions are uniformly continuous and bounded, this identity implies (3.5.3) by simply observing

that $h_w \in L^1(\mathbb{R}^N)$ and $f \in C_0(\mathbb{R}^N) \subset L^\infty(\mathbb{R}^N)$. Unfortunately, such an identity does not immediately follow from Lemma 3.5.2 as h_w is not constructed in the same way as the approximate identity φ_t . We can, however, prove the identity using the same proof technique from [SW71]:

Lemma 3.5.3. *Let $\rho \in L^1(\mathbb{R})$ with $\int_{\mathbb{R}} \rho(x) dx = 1$ and define $h_w \in L^1(\mathbb{R}^N)$ as in (3.5.1) for all $w \in \mathbb{R}^N$. Then we have*

$$\lim_{|w| \rightarrow \infty} \sup_{x \in \mathbb{R}^N} |(f * h_w)(x) - f(x)| = 0$$

for all $f \in C_0(\mathbb{R}^N)$.

Proof. By symmetry of the convolution operator in its arguments, we have

$$\begin{aligned} \sup_{x \in \mathbb{R}^N} |(f * h_w)(x) - f(x)| &= \sup_{x \in \mathbb{R}^N} \left| \int_{\mathbb{R}^N} f(y) h_w(x-y) dy - f(x) \right| \\ &= \sup_{x \in \mathbb{R}^N} \left| \int_{\mathbb{R}^N} f(x-y) h_w(y) dy - f(x) \right|. \end{aligned}$$

Since a simple substitution yields $1 = \int_{\mathbb{R}^N} \rho(x) dx = \int_{\mathbb{R}^N} h_w(x) dx$, an application of Minkowski's integral inequality (see, e.g., [SP70, Section A.1] or [HLC⁺52, Theorem 202]) for $L^\infty(\mathbb{R}^N)$ gives us

$$\begin{aligned} \sup_{x \in \mathbb{R}^N} |(f * h_w)(x) - f(x)| &= \sup_{x \in \mathbb{R}^N} \left| \int_{\mathbb{R}^N} (f(x-y) - f(x)) h_w(y) dy \right| \\ &\leq \int_{\mathbb{R}^N} |h_w(y)| \sup_{x \in \mathbb{R}^N} |f(x) - f(x-y)| dy. \end{aligned}$$

Finally, expanding the function h_w , we obtain

$$\begin{aligned} \sup_{x \in \mathbb{R}^N} |(f * h_w)(x) - f(x)| &\leq \int_{\mathbb{R}^N} \left(\prod_{j=1}^N w(j) |\rho(w(j)y(j))| \right) \sup_{x \in \mathbb{R}^N} |f(x) - f(x-y)| dy \\ &= \int_{\mathbb{R}^N} \left(\prod_{j=1}^N |\rho(z(j))| \right) \sup_{x \in \mathbb{R}^N} |f(x) - f(x - z \circ w^{-1})| dz, \end{aligned}$$

where we have used the substitution $z = y \circ w$; here, \circ denotes the Hadamard (entrywise) product, and we denote by $w^{-1} \in \mathbb{R}^N$ the vector whose j th entry is $1/w(j)$. Taking limits on both sides of this expression and observing that

$$\int_{\mathbb{R}^N} \left(\prod_{j=1}^N |\rho(z(j))| \right) \sup_{x \in \mathbb{R}^N} |f(x) - f(x - z \circ w^{-1})| dz \leq 2 \|\rho\|_1^N \sup_{x \in \mathbb{R}^N} |f(x)| < \infty,$$

it follows by the Dominated Convergence Theorem that

$$\lim_{|w| \rightarrow \infty} \sup_{x \in \mathbb{R}^N} |(f * h_w)(x) - f(x)| \leq \int_{\mathbb{R}^N} \left(\prod_{j=1}^N |\rho(z(j))| \right) \lim_{|w| \rightarrow \infty} \sup_{x \in \mathbb{R}^N} |f(x) - f(x - z \circ w^{-1})| dz,$$

and so it suffices to show that

$$\lim_{|w| \rightarrow \infty} \sup_{x \in \mathbb{R}^N} |f(x) - f(x - z \circ w^{-1})| = 0$$

for all $z \in \mathbb{R}^N$. To this end, let $\varepsilon > 0$ and $z \in \mathbb{R}^N$ be arbitrary. Since $f \in C_0(\mathbb{R}^N)$, there exists $r > 0$ sufficiently large such that $|f(x)| < \varepsilon/2$ for all $x \in \mathbb{R}^N \setminus \overline{B(0, r)}$, where $\overline{B(0, r)} \subset \mathbb{R}^N$ is the closed ball of radius r centered at the origin. Let $\mathcal{B} := \overline{B(0, r + \|z \circ w^{-1}\|_2)}$, so that for each $x \in \mathbb{R}^N \setminus \mathcal{B}$ we have both x and $x - z \circ w^{-1}$ in $\mathbb{R}^N \setminus \overline{B(0, r)}$. Thus, both $|f(x)| < \varepsilon/2$ and $|f(x - z \circ w^{-1})| < \varepsilon/2$, implying that

$$\sup_{x \in \mathbb{R}^N \setminus \mathcal{B}} |f(x) - f(x - z \circ w^{-1})| < \varepsilon.$$

Hence, we obtain

$$\begin{aligned} & \lim_{|w| \rightarrow \infty} \sup_{x \in \mathbb{R}^N} |f(x) - f(x - z \circ w^{-1})| \\ & \leq \lim_{|w| \rightarrow \infty} \max \left\{ \sup_{x \in \mathcal{B}} |f(x) - f(x - z \circ w^{-1})|, \sup_{x \in \mathbb{R}^N \setminus \mathcal{B}} |f(x) - f(x - z \circ w^{-1})| \right\} \end{aligned}$$

$$< \max \left\{ \varepsilon, \lim_{|w| \rightarrow \infty} \sup_{x \in \mathcal{B}} |f(x) - f(x - z \circ w^{-1})| \right\}.$$

Now, as \mathcal{B} is a compact subset of \mathbb{R}^N , the continuous function f is uniformly continuous on \mathcal{B} , and so the remaining limit and supremum may be freely interchanged, whereby continuity of f yields

$$\lim_{|w| \rightarrow \infty} \sup_{x \in \mathcal{B}} |f(x) - f(x - z \circ w^{-1})| = \sup_{x \in \mathcal{B}} \lim_{|w| \rightarrow \infty} |f(x) - f(x - z \circ w^{-1})| = 0.$$

Since $\varepsilon > 0$ may be taken arbitrarily small, we have proved the result. \square

As alluded to earlier, given $f \in C_0(\mathbb{R}^N)$, Lemma 3.5.3 implies that (3.5.3) holds uniformly for all $x \in \mathbb{R}^N$, that is,

$$\lim_{|w| \rightarrow \infty} \sup_{x \in \mathbb{R}^N} |(f * h_w)(x) - f(x)| = 0.$$

In particular, since both f and $f * h_w$ are uniformly continuous and bounded, we may swap the order of the limit and supremum operators to obtain

$$\sup_{x \in \mathbb{R}^N} \left| \lim_{|w| \rightarrow \infty} (f * h_w)(x) - f(x) \right| = 0. \quad (3.5.4)$$

Hence, we have $f(x) = \lim_{|w| \rightarrow \infty} (f * h_w)(x)$ uniformly for all $x \in \mathbb{R}^N$.

With (3.5.4) in hand, we may now use l'Hôpital's rule to show that

$$f(x) = \lim_{|w| \rightarrow \infty} (f * h_w)(x) = \lim_{\Omega \rightarrow \infty} \frac{1}{\Omega^N} \int_{[0, \Omega]^N} (f * h_w)(x) dy$$

holds uniformly for all $x \in \mathbb{R}^N$. Indeed, consider functions F and G which act on Borel subsets of

\mathbb{R}^N as follows:

$$F(A) := \int_A (f * h_w)(x) dy \quad \text{and} \quad G(A) := \int_A dy.$$

Choosing $A = [0, \Omega]^N$, the Lebesgue Differentiation Theorem states that

$$\frac{d}{d\Omega} F([0, \Omega]^N) = (f * h_w)(x) \Big|_{w=[\Omega, \dots, \Omega]} \quad \text{and} \quad \frac{d}{d\Omega} G([0, \Omega]^N) = 1$$

(in one-dimension, this is simply the Fundamental Theorem of Calculus). Now, as both $F([0, \Omega]^N)$ and $G([0, \Omega]^N)$ are unbounded as Ω tends to infinity, we may apply l'Hôpital's rule to obtain

$$\lim_{\Omega \rightarrow \infty} \frac{F([0, \Omega]^N)}{G([0, \Omega]^N)} = \lim_{\Omega \rightarrow \infty} (f * h_w)(x) \Big|_{w=[\Omega, \dots, \Omega]}.$$

Simplifying the left-hand side of this equation and making a substitution on the right-hand side, we have obtained

$$\lim_{\Omega \rightarrow \infty} \frac{1}{\Omega^N} \int_{[0, \Omega]^N} (f * h_w)(x) dy = \lim_{|w| \rightarrow \infty} (f * h_w)(x),$$

which is the desired equality. In summary, we have proved the following:

Lemma 3.5.4. *Let $f \in C_0(\mathbb{R}^N)$ and $\rho \in L^1(\mathbb{R})$ with $\int_{\mathbb{R}} \rho(z) dz = 1$. Define $h_w \in L^1(\mathbb{R}^N)$ as in (3.5.1) for all $w \in \mathbb{R}^N$. Then we have*

$$f(x) = \lim_{\Omega \rightarrow \infty} \frac{1}{\Omega^N} \int_{[0, \Omega]^N} (f * h_w)(x) dy \tag{3.5.5}$$

uniformly for all $x \in \mathbb{R}^N$.

3.5.2 The limit-integral representation

The next step in the proof of Theorem 3.3.1 is to represent f as the limiting value of a multidimensional integral over the parameter space. In particular, we seek to replace $(f * h_w)(x)$ in the convolution identity (3.5.5) with a function of the form $\int_K F(y)\rho(\langle w, x \rangle + b(y))dy$, as this will introduce the RVFL structure we require. To achieve this, we first use a truncated cosine function in place of the activation function ρ and then use identity (3.5.4) to switch back to a general activation function.

For each fixed $\Omega > 0$, let $L = L(\Omega) := \lceil \frac{2N}{\pi} \text{rad}(K)\Omega - \frac{1}{2} \rceil$ and define $\cos_\Omega: \mathbb{R} \rightarrow [-1, 1]$ by

$$\cos_\Omega(x) := \begin{cases} \cos(x) & x \in [-\frac{1}{2}(2L+1)\pi, \frac{1}{2}(2L+1)\pi], \\ 0 & \text{otherwise.} \end{cases} \quad (3.5.6)$$

Since $\cos_\Omega \in L^1(\mathbb{R}) \cap L^2(\mathbb{R})$, consider the function h_w defined in (3.5.1) with ρ replaced by \cos_Ω . Then we have

$$\begin{aligned} (f * h_w)(x) &= \int_{\mathbb{R}^N} f(y) \left(\prod_{j=1}^N w(j) \cos_\Omega \left(w(j)(x(j) - y(j)) \right) \right) dy \\ &= \int_{\mathbb{R}^N} f(y) \Delta(w, x - y) \left(\prod_{j=1}^N w(j) \right) dy \end{aligned}$$

for all $x \in \mathbb{R}^N$, where we define

$$\Delta(w, z) := \prod_{j=1}^N \cos_\Omega(w(j)z(j))$$

for all $w, z \in \mathbb{R}^N$. When substituted into (3.5.5), this yields the representation

$$f(x) = \lim_{\Omega \rightarrow \infty} \frac{1}{\Omega^N} \int_{\mathbb{R}^N \times [0, \Omega]^N} f(y) \Delta(w, x - y) \left(\prod_{j=1}^N w(j) \right) dy dw \quad (3.5.7)$$

uniformly for all $x \in \mathbb{R}^N$. In order to introduce the inner-product structure present in RVFL networks, we would like to convert the product in Δ to a summation. Now, if we consider the more general product $\prod_{j=1}^N \cos(z(j))$, using the sum and difference identity $2 \cos(a) \cos(b) = \cos(a-b) + \cos(a+b)$ iteratively yields

$$\prod_{j=1}^N \cos(z(j)) = \frac{1}{2^N} \sum_{\pm} \cos(\pm z(1) \pm \dots \pm z(N)),$$

where the summation is taken over all combinations of \pm appearing inside the cosine; for instance, in the case $N = 3$ we have

$$\begin{aligned} \prod_{j=1}^3 \cos(z(j)) &= \frac{1}{2} \cos(z(1)) \left(\cos(z(2) - z(3)) + \cos(z(2) + z(3)) \right) \\ &= \frac{1}{4} \left(\cos(z(1) - z(2) + z(3)) + \cos(z(1) + z(2) - z(3)) \right. \\ &\quad \left. + \cos(z(1) - z(2) - z(3)) + \cos(z(1) + z(2) + z(3)) \right), \end{aligned}$$

at which point multiplying by $1 = \cos(0)$ and applying the sum and difference identity four more times yields

$$\prod_{j=1}^3 \cos(z(j)) = \frac{1}{8} \sum_{\pm} \cos(\pm z(1) \pm z(2) \pm z(3)).$$

To apply the same procedure for the product in Δ , first observe that we have chosen the value of L in a particularly nice way, so that

$$-\frac{\pi}{2}(2L+1) \leq \sum_{j=1}^N \left(\pm w(j)(x(j) - y(j)) \right) \leq \frac{\pi}{2}(2L+1)$$

for any $w \in [0, \Omega]$, $x, y \in K$, and all combinations of sign choices. Hence, we may apply the sum and difference identity $2 \cos_{\Omega}(a) \cos_{\Omega}(b) = \cos_{\Omega}(a-b) + \cos_{\Omega}(a+b)$ inside $\Delta(w, x - y)$ in the

same iterative way to obtain

$$\begin{aligned}\Delta(w, x - y) &= \prod_{j=1}^N \cos_{\Omega} \left(w(j)(x(j) - y(j)) \right) \\ &= \frac{1}{2^N} \sum_{\pm} \cos_{\Omega} \left(\pm w(1)(x(1) - y(1)) \pm \cdots \pm w(N)(x(N) - y(N)) \right)\end{aligned}$$

for all $w \in [0, \Omega]$ and $x, y \in K$. Now, noting that for each $j = 1, \dots, N$ and any constant C the symmetry of \cos_{Ω} gives us

$$\begin{aligned}\int_0^{\Omega} w(j) \left(\cos_{\Omega} \left(w(j)(x(j) - y(j)) + C \right) + \cos_{\Omega} \left(-w(j)(x(j) - y(j)) + C \right) \right) dy(j) \\ = \int_0^{\Omega} w(j) \cos_{\Omega} \left(w(j)(x(j) - y(j)) + C \right) dy(j) \\ \quad - \int_{-\Omega}^0 w(j) \cos_{\Omega} \left(w(j)(x(j) - y(j)) + C \right) dy(j) \\ = \int_{-\Omega}^{\Omega} |w(j)| \cos_{\Omega} \left(w(j)(x(j) - y(j)) + C \right) dy(j),\end{aligned}$$

by replacing each variable $-w(j)$ in $\Delta(w, x - y)$ with $w(j)$ we may write

$$\int_{[0, \Omega]^N} \Delta(w, x - y) \left(\prod_{j=1}^N w(j) \right) dy = \frac{1}{2^N} \int_{[-\Omega, \Omega]^N} \cos_{\Omega} (\langle w, x - y \rangle) \left| \prod_{j=1}^N w(j) \right| dy$$

for all $x, y \in K$. Plugging this expression into (3.5.7), it follows that

$$f(x) = \lim_{\Omega \rightarrow \infty} \frac{1}{(2\Omega)^N} \int_{K \times [-\Omega, \Omega]^N} f(y) \cos_{\Omega} (\langle w, x - y \rangle) \left| \prod_{j=1}^N w(j) \right| dy dw \quad (3.5.8)$$

holds uniformly for all $x \in K$.

With the representation (3.5.8) in hand, we now seek to reintroduce the general activation function ρ . To this end, since $\cos_{\Omega} \in C_c(\mathbb{R}) \subset C_0(\mathbb{R})$ we may apply the convolution identity (3.5.4) with f replaced by \cos_{Ω} to obtain $\cos_{\Omega}(z) = \lim_{\alpha \rightarrow \infty} (\cos_{\Omega} * h_{\alpha})(z)$ uniformly for all $z \in \mathbb{R}$, where h_{α} is the one-dimensional version of h_w as defined in (3.5.1). Using this

representation of \cos_Ω in (3.5.8), it follows that

$$f(x) = \lim_{\Omega \rightarrow \infty} \frac{1}{(2\Omega)^N} \int_{K \times [-\Omega, \Omega]^N} f(y) \left(\lim_{\alpha \rightarrow \infty} (\cos_\Omega * h_\alpha)(\langle w, x - y \rangle) \right) \left| \prod_{j=1}^N w(j) \right| dy dw$$

holds uniformly for all $x \in K$. Since f is continuous and the convolution $\cos_\Omega * h_\alpha$ is uniformly continuous and bounded, the fact that the domain $K \times [-\Omega, \Omega]^N$ is compact then allows us to bring the limit as α tends to infinity outside the integral in this expression via the Dominated Convergence Theorem, which gives us

$$f(x) = \lim_{\Omega \rightarrow \infty} \lim_{\alpha \rightarrow \infty} \frac{1}{(2\Omega)^N} \int_{K \times [-\Omega, \Omega]^N} f(y) (\cos_\Omega * h_\alpha)(\langle w, x - y \rangle) \left| \prod_{j=1}^N w(j) \right| dy dw \quad (3.5.9)$$

uniformly for every $x \in K$.

Remark 3.5.5. It should be noted that we are unable to swap the order of the limits in (3.5.9); indeed, our use of (3.5.4) is no longer valid in this case, as \cos_Ω is not in $C_0(\mathbb{R})$ when Ω is allowed to be infinite.

To complete this step of the proof, observe that the definition of \cos_Ω allows us to write

$$(\cos_\Omega * h_\alpha)(z) = \alpha \int_{\mathbb{R}} \cos_\Omega(u) \rho(\alpha(z - u)) du = \alpha \int_{-\frac{\pi}{2}(2L+1)}^{\frac{\pi}{2}(2L+1)} \cos_\Omega(u) \rho(\alpha(z - u)) du \quad (3.5.10)$$

uniformly for all $z \in \mathbb{R}$. By substituting (3.5.10) into (3.5.9), we then obtain

$$f(x) = \lim_{\Omega \rightarrow \infty} \lim_{\alpha \rightarrow \infty} \frac{\alpha}{(2\Omega)^N} \int_{K(\Omega)} f(y) \cos_\Omega(u) \rho(\alpha(\langle w, x - y \rangle - u)) \left| \prod_{j=1}^N w(j) \right| dy dw du$$

uniformly for all $x \in K$, where $K(\Omega) := K \times [-\Omega, \Omega]^N \times [-\frac{\pi}{2}(2L+1), \frac{\pi}{2}(2L+1)]$. In this way, if

we define

$$F_{\alpha,\Omega}(y, w, u) := \frac{\alpha}{(2\Omega)^N} \left| \prod_{j=1}^N w(j) \right| f(y) \cos_{\Omega}(u), \quad (3.5.11)$$

$$b_{\alpha}(y, w, u) := -\alpha(\langle w, y \rangle + u)$$

for $y, w \in \mathbb{R}^N$ and $u \in \mathbb{R}$, we have proved the following result:

Lemma 3.5.6. *Let $f \in C_c(\mathbb{R}^N)$ and $\rho \in L^1(\mathbb{R})$ with $K := \text{supp}(f)$ and $\int_{\mathbb{R}} \rho(z) dz = 1$. Define $F_{\alpha,\Omega}$ and b_{α} as in (3.5.11) for all $\Omega \in \mathbb{R}^N$ and $\alpha \in \mathbb{R}$. Then we have*

$$f(x) = \lim_{\Omega \rightarrow \infty} \lim_{\alpha \rightarrow \infty} \int_{K(\Omega)} F_{\alpha,\Omega}(y, w, u) \rho(\alpha \langle w, x \rangle + b_{\alpha}(y, w, u)) dy dw du \quad (3.5.12)$$

uniformly for every $x \in K$, where $K(\Omega) := K \times [-\Omega, \Omega]^N \times [-\frac{\pi}{2}(2L+1), \frac{\pi}{2}(2L+1)]$ and $L := \lceil \frac{2N}{\pi} \text{rad}(K) \Omega - \frac{1}{2} \rceil$.

Remark 3.5.7. We will see in Section 3.5.3 that the RVFL networks f_n will be built using random samples drawn independently and uniformly from the domain $K(\Omega)$. Since the range $[-\frac{\pi}{2}(2L+1), \frac{\pi}{2}(2L+1)]$ is potentially quite large (compared to Ω), for practical purposes we may instead use the domain $K \times [-\Omega, \Omega]^N \times [\Omega, \Omega]$. Indeed, by defining the truncation errors

$$\tilde{v}(x) := \frac{1}{(2\Omega)^N} \int_{K \times [-\Omega, \Omega]^N} v(\langle w, x - y \rangle) f(y) \left| \prod_{j=1}^N w(j) \right| dy dw,$$

$$v(z) := \alpha \int_{-\infty}^{-\Omega} \cos_{\Omega}(u) \rho(\alpha(z - u)) du + \alpha \int_{\Omega}^{\infty} \cos_{\Omega}(u) \rho(\alpha(z - u)) du$$

for all $x \in \mathbb{R}^N$ and $z \in \mathbb{R}$, the representation (3.5.12) then becomes

$$f(x) = \lim_{\Omega \rightarrow \infty} \lim_{\alpha \rightarrow \infty} \left(\tilde{v}(x) + \int_{K \times [-\Omega, \Omega]^N \times [\Omega, \Omega]} F_{\alpha,\Omega}(y, w, u) \rho(\alpha \langle w, x \rangle + b_{\alpha}(y, w, u)) dy dw du \right)$$

uniformly for all $x \in K$; in particular,

$$|\tilde{\mathbf{v}}(x)| \lesssim M \text{vol}(K) \left(\|\rho\|_1 - \inf_{\substack{w \in [-\Omega, \Omega]^N \\ x, y \in K}} \int_{-\alpha(\Omega + \langle w, x-y \rangle)}^{\alpha(\Omega + \langle w, x-y \rangle)} |\rho(u)| \, \mathrm{d}u \right),$$

where $M := \sup_{x \in K} |f(x)| < \infty$, which decays to zero as α tends to infinity at least as fast as the tails of $\rho \in L^1(\mathbb{R})$.

3.5.3 Monte-Carlo integral approximation

The next step in the proof of Theorem 3.3.1 is to approximate the integral in (3.5.12) using the Monte-Carlo method. To this end, let $\{y_k\}_{k=1}^n$, $\{w_k\}_{k=1}^n$, and $\{u_k\}_{k=1}^n$ be independent samples drawn uniformly from K , $[-\Omega, \Omega]^N$, and $[-\frac{\pi}{2}(2L+1), \frac{\pi}{2}(2L+1)]$, respectively, and consider the sequence of random variables $\{I_n(x)\}_{n=1}^\infty$ defined by

$$I_n(x) := \frac{\text{vol}(K(\Omega))}{n} \sum_{k=1}^n F_{\alpha, \Omega}(y_k, w_k, u_k) \rho(\alpha \langle w_k, x \rangle + b_\alpha(y_k, w_k, u_k)) \quad (3.5.13)$$

for each $x \in K$, where we note that $\text{vol}(K(\Omega)) = (2\Omega)^N \pi(2L+1) \text{vol}(K)$. If we also define

$$I(x; p) := \int_{K(\Omega)} \left(F_{\alpha, \Omega}(y, w, u) \rho(\alpha \langle w, x \rangle + b_\alpha(y, w, u)) \right)^p \, \mathrm{d}y \mathrm{d}w \mathrm{d}u \quad (3.5.14)$$

for $x \in K$ and $p \in \mathbb{N}$, then we want to show that

$$\lim_{n \rightarrow \infty} \mathbb{E} \int_K |I(x; 1) - I_n(x)|^2 \, \mathrm{d}x = 0 \quad (3.5.15)$$

with convergence rate $O(1/n)$, where the expectation is taken with respect to the joint distribution of the random samples $\{y_k\}_{k=1}^n$, $\{w_k\}_{k=1}^n$, and $\{u_k\}_{k=1}^n$. For this, it suffices to find a constant

$C_{f,\rho,\alpha,\Omega,N} < \infty$ independent of n satisfying

$$\int_K \mathbb{E} |I(x;1) - I_n(x)|^2 dx \leq \frac{C_{f,\rho,\alpha,\Omega,N}}{n};$$

indeed, an application of Fubini's theorem would then yield

$$\mathbb{E} \int_K |I(x;1) - I_n(x)|^2 dx \leq \frac{C_{f,\rho,\alpha,\Omega,N}}{n},$$

which implies (3.5.15). To determine such a constant, we first observe by an application of Theorem 3.4.1 that

$$\mathbb{E} |I(x;1) - I_n(x)|^2 = \frac{\text{vol}^2(K(\Omega)) \sigma(x)^2}{n},$$

where we define the variance term

$$\sigma(x)^2 := \frac{I(x;2)}{\text{vol}(K(\Omega))} - \frac{I(x;1)^2}{\text{vol}^2(K(\Omega))}$$

for $x \in K$. Noting that

$$|F_{\alpha,\Omega}(y,w,u)| = \frac{\alpha}{(2\Omega)^N} \left| \prod_{j=1}^N w(j) \right| |f(y)| |\cos_{\Omega}(u)| \leq \frac{\alpha M}{2^N}$$

for all $y, w \in \mathbb{R}^N$ and $u \in \mathbb{R}$, where $M := \sup_{x \in K} |f(x)| < \infty$, observe that a simple bound on this variance term is

$$\sigma(x)^2 \leq \frac{I(x;2)}{\text{vol}(K(\Omega))} \leq \frac{\alpha^2 M^2}{2^{2N} \text{vol}(K(\Omega))} \int_{K(\Omega)} \left| \rho(\alpha \langle w, x \rangle + b_{\alpha}(y, w, u)) \right|^2 dy dw du. \quad (3.5.16)$$

Since we assume $\rho \in L^2(\mathbb{R})$, we then have

$$\begin{aligned}
\int_K \mathbb{E}|I(x; 1) - I_n(x)|^2 dx &= \frac{\text{vol}^2(K(\Omega))}{n} \int_K \sigma(x)^2 dx \\
&\leq \frac{\alpha^2 M^2 \text{vol}(K(\Omega))}{2^{2N} n} \int_{K \times K(\Omega)} \left| \rho(\alpha \langle w, x \rangle + b_\alpha(y, w, u)) \right|^2 dx dy dw du \\
&\leq \frac{\alpha^2 M^2 \text{vol}(K(\Omega))}{2^{2N} n} \int_{K(\Omega)} \|\rho\|_2^2 dy dw du \\
&= \frac{\alpha^2 M^2 \text{vol}^2(K(\Omega)) \|\rho\|_2^2}{2^{2N} n}.
\end{aligned}$$

Substituting the value of $\text{vol}(K(\Omega))$, this means that

$$C_{f, \rho, \alpha, \Omega, N} := \alpha^2 M^2 \Omega^{2N} \pi^2 (2L + 1)^2 \text{vol}^2(K) \|\rho\|_2^2$$

is a suitable choice for the desired constant.

Now that we have established (3.5.15), we may rewrite the random variables $I_n(x)$ in a more convenient form. To this end, we change the domain of the random samples $\{w_k\}_{k=1}^n$ to $[-\alpha\Omega, \alpha\Omega]^N$ and define the new random variables $\{b_k\}_{k=1}^n \subset \mathbb{R}$ by $b_k := -(\langle w_k, y_k \rangle + \alpha u_k)$ for each $k = 1, \dots, n$. In this way, if we denote

$$v_k := \frac{\text{vol}(K(\Omega))}{n} F_{\alpha, \Omega} \left(y_k, \frac{w_k}{\alpha^N}, u_k \right)$$

for each $k = 1, \dots, n$, the random variables $\{f_n\}_{n=1}^\infty$ defined by

$$f_n(x) := \sum_{k=1}^n v_k \rho(\langle w_k, x \rangle + b_k) \tag{3.5.17}$$

satisfy $f_n(x) = I_n(x)$ for every $x \in K$. Combining this with (3.5.15), we have proved the following:

Lemma 3.5.8. *Let $f \in C_c(\mathbb{R}^N)$ and $\rho \in L^1(\mathbb{R}) \cap L^2(\mathbb{R})$ with $K := \text{supp}(f)$ and $\int_{\mathbb{R}} \rho(z) dz = 1$. Define f_n as in (3.5.17) for each $n \in \mathbb{N}$ and $F_{\alpha, \Omega}, b_\alpha$ as in (3.5.11) for all $\Omega \in \mathbb{R}^N$ and $\alpha \in \mathbb{R}$.*

Then we have

$$\lim_{n \rightarrow \infty} \mathbb{E} \int_K \left| \int_{K(\Omega)} F_{\alpha, \Omega}(y, w, u) \rho(\alpha \langle w, x \rangle + b_\alpha(y, w, u)) dy dw du - f_n(x) \right|^2 dx = 0, \quad (3.5.18)$$

where $K(\Omega) := K \times [-\Omega, \Omega]^N \times [-\frac{\pi}{2}(2L+1), \frac{\pi}{2}(2L+1)]$ and $L := \lceil \frac{2N}{\pi} \text{rad}(K)\Omega - \frac{1}{2} \rceil$, with convergence rate $O(1/n)$.

3.5.4 Bounding the asymptotic mean square error

The fourth and final step in the proof of Theorem 3.3.1 is to combine the limit representation (3.5.12) with the Monte-Carlo error guarantee (3.5.18) and show that, given any $\varepsilon > 0$, there exist $\alpha, \Omega > 0$ such that

$$\lim_{n \rightarrow \infty} \mathbb{E} \int_K |f(x) - f_n(x)|^2 dx < \varepsilon.$$

To this end, let $\varepsilon' > 0$ be arbitrary and consider the integral $I(x; p)$ in (3.5.14) for $x \in K$ and $p \in \mathbb{N}$. By (3.5.12), there exists $\alpha, \Omega > 0$ such that $|f(x) - I(x; 1)| < \varepsilon'$ holds uniformly for every $x \in K$, and so it follows that

$$|f(x) - f_n(x)| < \varepsilon' + |I(x; 1) - f_n(x)|$$

for every $x \in K$. Hence, we have

$$\begin{aligned} & \lim_{n \rightarrow \infty} \mathbb{E} \int_K |f(x) - f_n(x)|^2 dx \\ & < (\varepsilon')^2 \text{vol}(K) + \lim_{n \rightarrow \infty} \mathbb{E} \int_K |I(x; 1) - f_n(x)|^2 dx + 2\varepsilon' \lim_{n \rightarrow \infty} \mathbb{E} \int_K (I(x; 1) - f_n(x)) dx. \end{aligned} \quad (3.5.19)$$

By (3.5.18), we know that the second term on the right-hand side of (3.5.19) vanishes at a rate proportional to $1/n$. On the other hand, the third term on the right-hand side of (3.5.19) vanishes by applying Fubini's Theorem and observing that $\mathbb{E}f_n(x) = I(x; 1)$ for all $n \in \mathbb{N}$ and $x \in K$. Therefore, we have

$$\lim_{n \rightarrow \infty} \mathbb{E} \int_K |f(x) - f_n(x)|^2 dx < (\varepsilon')^2 \text{vol}(K)$$

with convergence rate $O(1/n)$, and so the proof is completed by taking $\varepsilon' = \sqrt{\varepsilon/\text{vol}(K)}$ and choosing $\alpha, \Omega > 0$ accordingly.

It remains only to verify our use of Fubini's Theorem in evaluating the final term on the right-hand side of (3.5.19). To this end, recall that the Monte Carlo integral approximation f_n satisfies $\lim_{n \rightarrow \infty} (I(x; 1) - f_n(x)) \sim \text{Norm}(0, \sigma(x)^2)$ via the Central Limit Theorem. Hence, we have

$$\mathbb{E} \lim_{n \rightarrow \infty} |I(x; 1) - f_n(x)| \leq \sigma(x) \sqrt{\frac{2}{\pi}}. \quad (3.5.20)$$

Since have already seen in (3.5.16) that

$$\sigma(x) \leq \frac{\alpha M}{2^N \sqrt{\text{vol}(K(\Omega))}} \left(\int_{K(\Omega)} |\rho(\alpha \langle w, x \rangle + b_\alpha(y, w, u))|^2 dy dw du \right)^{1/2}$$

for all $x \in K$, observing that

$$\begin{aligned} \int_{K(\Omega)} |\rho(\alpha \langle w, x \rangle + b_\alpha(y, w, u))|^2 dy dw du &= \int_{K(\Omega)} |\rho(\alpha \langle w, x - y \rangle - \alpha u)|^2 dy dw du \\ &\leq \int_{[-\Omega, \Omega]^N \times [-\frac{\pi}{2}(2L+1), \frac{\pi}{2}(2L+1)]} \|\rho\|_2^2 dw du \\ &= \frac{\text{vol}(K(\Omega))}{\text{vol}(K)} \|\rho\|_2^2, \end{aligned}$$

we obtain the bound

$$\int_K \mathbb{E} \lim_{n \rightarrow \infty} |I(x; 1) - f_n(x)| dx \leq \sqrt{\frac{2}{\pi}} \int_K \sigma(x) dx \leq \frac{\alpha M \|\rho\|_2 \sqrt{\text{vol}(K)}}{2^{N-1/2} \sqrt{\pi}},$$

which is necessarily finite. Therefore, we may apply both Fubini's Theorem and the Dominated Convergence Theorem to obtain

$$\int_K \mathbb{E} \lim_{n \rightarrow \infty} (I(x; 1) - f_n(x)) dx = \lim_{n \rightarrow \infty} \int_K \mathbb{E} (I(x; 1) - f_n(x)) dx = \lim_{n \rightarrow \infty} \mathbb{E} \int_K (I(x; 1) - f_n(x)) dx,$$

as desired.

3.6 Proofs of Corollary 3.3.3 and Theorem 3.3.4

In this section we prove the remaining results for RVFL networks in \mathbb{R}^N . The proof techniques rely heavily on that used to prove Theorem 3.3.1, and so we will refer back to the results in that proof as needed.

3.6.1 Proof of Corollary 3.3.3

Let $f \in C_c(\mathbb{R}^N)$ with $K := \text{supp}(f)$ and suppose $\varepsilon > 0$ is fixed. Take the activation function $\rho: \mathbb{R} \rightarrow \mathbb{R}$ to be differentiable with $\rho' \in L^1(\mathbb{R}) \cap L^2(\mathbb{R})$. We wish to show that there exists a sequence of RVFL networks $\{f_n\}_{n=1}^\infty$ defined on K which satisfy the asymptotic error bound

$$\lim_{n \rightarrow \infty} \mathbb{E} \int_K |f(x) - f_n(x)|^2 dx < \varepsilon.$$

The proof of this result is a minor modification of the first two steps in the proof of Theorem 3.3.1.

To begin, note that ρ' satisfies the assumptions on ρ in Theorem 3.3.1. Hence, we may

use Lemma 3.5.4 with h_w defined by

$$h_w(y) := \prod_{j=1}^N w(j) \rho'(w(j)y(j))$$

for all $y, w \in \mathbb{R}^N$ to obtain the representation (3.5.5) for all $x \in \mathbb{R}^N$, which leads to the representation (3.5.9). Now, since (3.5.10) gives us

$$(\cos_\Omega * h_\alpha)(z) = \alpha \int_{\mathbb{R}} \cos_\Omega(u) \rho'(\alpha(z-u)) du$$

uniformly for all $z \in \mathbb{R}$, recalling the definition of \cos_Ω in (3.5.6) and integrating by parts, we obtain

$$\begin{aligned} (\cos_\Omega * h_\alpha)(z) &= \alpha \int_{\mathbb{R}} \cos_\Omega(u) \rho'(\alpha(z-u)) du \\ &= - \int_{-\frac{\pi}{2}(2L+1)}^{\frac{\pi}{2}(2L+1)} \cos_\Omega(u) d\rho(\alpha(z-u)) \\ &= -\cos_\Omega(u) \rho(\alpha(z-u)) \Big|_{-\frac{\pi}{2}(2L+1)}^{\frac{\pi}{2}(2L+1)} + \int_{-\frac{\pi}{2}(2L+1)}^{\frac{\pi}{2}(2L+1)} \rho(\alpha(z-u)) d\cos_\Omega(u) \\ &= - \int_{\mathbb{R}} \sin_\Omega(u) \rho(\alpha(z-u)) du \end{aligned}$$

for all $z \in \mathbb{R}$, where $L := \lceil \frac{2N}{\pi} \text{rad}(K)\Omega - \frac{1}{2} \rceil$ and $\sin_\Omega: \mathbb{R} \rightarrow [-1, 1]$ is defined analogously to (3.5.6). Substituting this representation of $(\cos_\Omega * h_\alpha)(z)$ into (3.5.9) then yields

$$f(x) = \lim_{\Omega \rightarrow \infty} \lim_{\alpha \rightarrow \infty} \frac{-\alpha}{(2\Omega)^N} \int_{K(\Omega)} f(y) \sin_\Omega(\langle w, x-y \rangle) \rho(\alpha(z-u)) \Big| \prod_{j=1}^N w(j) \Big| dy dw du$$

uniformly for every $x \in K$. Thus, if we replace the definition of $F_{\alpha, \Omega}$ in (3.5.11) by

$$F_{\alpha, \Omega}(y, w, u) := \frac{-\alpha}{(2\Omega)^N} \Big| \prod_{j=1}^N w(j) \Big| f(y) \sin_\Omega(u)$$

for $y, w \in \mathbb{R}^N$ and $u \in \mathbb{R}$, we again obtain the uniform representation (3.5.12) for all $x \in K$. The remainder of the proof proceeds from this point exactly as in the proof of Theorem 3.3.1.

3.6.2 Proof of Theorem 3.3.4

Let $f \in C_c(\mathbb{R}^N)$ with $K := \text{supp}(f)$ and suppose $\varepsilon > 0$, $\eta \in (0, 1)$ are fixed. Take any κ -Lipschitz activation function $\rho \in L^1(\mathbb{R}) \cap L^2(\mathbb{R})$ arbitrarily. We wish to show that there exists a sequence of RVFL networks $\{f_n\}_{n=1}^\infty$ defined on K which satisfy the error bound

$$\int_K |f(x) - f_n(x)|^2 dx < \varepsilon$$

with probability at least $1 - \eta$ when n is chosen sufficiently large. The proof is obtained by modifying the proof of Theorem 3.3.1 for the asymptotic case.

We begin by repeating the first two steps in the proof of Theorem 3.3.1 from Sections 3.5.1 and 3.5.2. In particular, by Lemma 3.5.6 we have the representation (3.5.12), namely,

$$f(x) = \lim_{\Omega \rightarrow \infty} \lim_{\alpha \rightarrow \infty} \int_{K(\Omega)} F_{\alpha, \Omega}(y, w, u) \rho(\alpha \langle w, x \rangle + b_\alpha(y, w, u)) dy dw du$$

holds uniformly for all $x \in K$. Hence, if we define the random variables f_n and I_n from Section 3.5.3 as in (3.5.17) and (3.5.13), respectively, we seek a uniform bound on the quantity

$$|f(x) - f_n(x)| \leq |f(x) - I(x; 1)| + |I_n(x) - I(x; 1)|$$

over the compact set K , where $I(x; 1)$ is given by (3.5.14) for all $x \in K$. Since equation (3.5.12) allows us to fix $\alpha, \Omega > 0$ such that

$$|f(x) - I(x; 1)| = \left| f(x) - \int_{K(\Omega)} F_{\alpha, \Omega}(y, w, u) \rho(\alpha \langle w, x \rangle + b) dy dw du \right| < \sqrt{\frac{\varepsilon}{2 \text{vol}(K)}}$$

holds for every $x \in K$ simultaneously, the result follows if we show $|I_n(x) - I(x; 1)| < \sqrt{\varepsilon/2\text{vol}(K)}$ uniformly for all $x \in K$ with high probability, since this would yield

$$\int_K |f(x) - f_n(x)|^2 dx \leq \int_K |f(x) - I(x; 1)|^2 dx + \int_K |I_n(x) - I(x; 1)|^2 dx < \varepsilon$$

with high probability. To this end, for $\delta > 0$ let $\mathcal{C}(\delta, K) \subset K$ denote a minimal δ -net for K , with cardinality $\mathcal{N}(\delta, K)$. Now, fix $x \in K$ and consider the inequality

$$|I_n(x) - I(x; 1)| \leq \underbrace{|I_n(x) - I_n(z)|}_{(*)} + \underbrace{|I_n(z) - I(z; 1)|}_{(**)} + \underbrace{|I(x; 1) - I(z; 1)|}_{(***)}, \quad (3.6.1)$$

where $z \in \mathcal{C}(\delta, K)$ is such that $\|x - z\|_2 < \delta$. We will obtain the desired bound on (3.6.1) by bounding each of the terms $(*)$, $(**)$, and $(***)$ separately.

First, we consider the term $(*)$. Recalling the definition of I_n , observe that we have

$$\begin{aligned} (*) &= \frac{(2\Omega)^{N+1}\text{vol}(K)}{n} \left| \sum_{k=1}^n F_{\alpha, \Omega}(y_k, w_k, u_k) \left(\rho(\alpha \langle w_k, x \rangle + b_\alpha(y_k, w_k, u_k)) \right. \right. \\ &\quad \left. \left. - \rho(\alpha \langle w_k, z \rangle + b_\alpha(y_k, w_k, u_k)) \right) \right| \\ &\leq \frac{2\alpha M \Omega^{N+1}\text{vol}(K)}{n} \sum_{k=1}^n \left| \rho(\alpha \langle w_k, x \rangle + b_\alpha(y_k, w_k, u_k)) - \rho(\alpha \langle w_k, z \rangle + b_\alpha(y_k, w_k, u_k)) \right| \\ &\leq 2\alpha M \Omega^{N+1}\text{vol}(K) R_{\alpha, \Omega}(x, z), \end{aligned}$$

where $M := \sup_{x \in K} |f(x)|$ and we define

$$R_{\alpha, \Omega}(x, z) := \sup_{\substack{y \in K \\ w \in [-\Omega, \Omega]^N \\ u \in [-\Omega, \Omega]}} \left| \rho(\alpha \langle w, x \rangle + b_\alpha(y, w, u)) - \rho(\alpha \langle w, z \rangle + b_\alpha(y, w, u)) \right|.$$

Now, since ρ is assumed to be κ -Lipschitz, we have

$$\begin{aligned} & \left| \rho(\alpha\langle w, x \rangle + b_\alpha(y, w, u)) - \rho(\alpha\langle w, z \rangle + b_\alpha(y, w, u)) \right| \\ &= \left| \rho(\alpha(\langle w, x - y \rangle - u)) - \rho(\alpha(\langle w, z - y \rangle - u)) \right| \leq \kappa\alpha|\langle w, x - z \rangle| \end{aligned}$$

for any $y \in K$, $w \in [-\Omega, \Omega]^N$, and $u \in [-\Omega, \Omega]$. Hence, an application of the Cauchy-Schwarz inequality yields $R_{\alpha, \Omega}(x, z) \leq \kappa\alpha\Omega\delta\sqrt{N}$ for all $x \in K$, from which it follows that

$$(*) \leq 2M\sqrt{N}\kappa\delta\alpha^2\Omega^{N+2}\text{vol}(K) \quad (3.6.2)$$

holds for all $x \in K$.

Next, we bound $(***)$ using a similar approach to that used to bound $(*)$. Indeed, by the definition of $I(\cdot; 1)$ we have

$$\begin{aligned} (***) &= \left| \int_{K(\Omega)} F_{\alpha, \Omega}(y, w, u) \left(\rho(\alpha\langle w, x \rangle + b_\alpha(y, w, u)) - \rho(\alpha\langle w, z \rangle + b_\alpha(y, w, u)) \right) dydwdu \right| \\ &\leq \frac{\alpha M}{2^N} \int_{K(\Omega)} \left| \rho(\alpha\langle w, x \rangle + b_\alpha(y, w, u)) - \rho(\alpha\langle w, z \rangle + b_\alpha(y, w, u)) \right| dydwdu \\ &\leq \frac{\alpha M \text{vol}(K(\Omega))}{2^N} R_{\alpha, \Omega}(x, z). \end{aligned}$$

Using the fact that $R_{\alpha, \Omega}(x, z) \leq \kappa\alpha\Omega\delta\sqrt{N}$ for all $x \in K$, it follows that

$$(***) \leq \frac{M\sqrt{N}\kappa\delta\alpha^2\Omega\text{vol}(K(\Omega))}{2^N} \quad (3.6.3)$$

holds for all $x \in K$.

Notice that the inequalities (3.6.2) and (3.6.3) are deterministic. In fact, both can be controlled by choosing an appropriate value for δ in the net $\mathcal{C}(\delta, K)$. To see this, fix $\varepsilon' > 0$ arbitrarily and recall that $\text{vol}(K(\Omega)) = (2\Omega)^N \pi(2L+1)\text{vol}(K)$. A simple computation then

shows that $(*) + (***) < \varepsilon'$ whenever

$$\delta < \frac{\varepsilon'}{2\sqrt{N}\kappa\alpha^2M\Omega^{N+2}\text{vol}(K)(1+2N\text{rad}(K))}. \quad (3.6.4)$$

We now continue to bound $(**)$ uniformly for $x \in K$. Unlike $(*)$ and $(***)$, we cannot bound this term deterministically. However, since $f_n \in L^2(K(\Omega))$, we may apply Lemma 3.4.3 to obtain the tail bound

$$\mathbb{P}((**) \geq t) \leq 3 \exp\left(-\frac{nt}{C_z c} \log\left(1 + \frac{C_z t}{\text{vol}^2(K(\Omega))\sigma(z)^2}\right)\right)$$

for all $t > 0$, where $c > 0$ is a numerical constant and

$$C_z := \text{ess sup}_{k \in \{1, \dots, n\}} \left| \text{vol}(K(\Omega)) F_{\alpha, \Omega}\left(y_k, \frac{w_k}{\alpha^N}, u_k\right) \rho(\langle w_k, z \rangle + b_k) - I(z; 1) \right|,$$

$$\sigma(z)^2 := \frac{I(z; 2)}{\text{vol}(K(\Omega))} - \frac{I(z; 1)^2}{\text{vol}^2(K(\Omega))}$$

for all $z \in \mathcal{C}(\delta, K)$. Taking

$$C := \sup_{z \in \mathcal{C}(\delta, K)} C_z \quad \text{and} \quad \Sigma := \sup_{z \in \mathcal{C}(\delta, K)} \sigma(z)^2, \quad (3.6.5)$$

which are now fixed constants describing the complexity of the function $F_{\alpha, \Omega}\rho$, if we choose the number of nodes such that

$$n \geq \frac{Cc \log(3\eta^{-1} \mathcal{N}(\delta, K))}{t \log\left(1 + \frac{Ct}{\text{vol}^2(K(\Omega))\Sigma}\right)}, \quad (3.6.6)$$

then a union bound yields $(**) < t$ simultaneously for all $z \in \mathcal{C}(\delta, K)$ with probability at least

$1 - \eta$. Combined with the bounds (3.6.2) and (3.6.3), it follows from (3.6.1) that

$$|I_n(x) - I(x; 1)| < \varepsilon' + t$$

simultaneously for all $x \in K$ with probability at least $1 - \eta$, provided δ and n satisfy (3.6.4) and (3.6.6), respectively. Since we require $|I_n(x) - I(x; 1)| < \sqrt{\varepsilon/2\text{vol}(K)}$, the proof is then completed by setting $\varepsilon' + t = \sqrt{\varepsilon/2\text{vol}(K)}$ and choosing δ and n accordingly. In particular, it suffices to choose $\varepsilon' = t = \frac{1}{2}\sqrt{\varepsilon/2\text{vol}(K)}$, so that (3.6.4) and (3.6.6) become

$$\begin{aligned} \delta &< \frac{\sqrt{\varepsilon}}{4\sqrt{N}\kappa\alpha^2 M\Omega^{N+2}\text{vol}^{3/2}(K)(1+2N\text{rad}(K))}, \\ n &\geq \frac{2\sqrt{2\text{vol}(K)}Cc \log(3\eta^{-1}\mathcal{N}(\delta, K))}{\sqrt{\varepsilon} \log\left(1 + \frac{C\sqrt{\varepsilon}}{2\sqrt{2\text{vol}^{5/2}(K(\Omega))\Sigma}}\right)}, \end{aligned}$$

as desired.

Remark 3.6.1. The κ -Lipschitz assumption on the activation function ρ may likely be removed for most practical applications. Indeed, since $(***)$ can be bounded instead by leveraging continuity of the L^1 norm with respect to translation, the only term whose bound depends on the Lipschitz property of ρ is $(*)$. However, there is randomness in I_n that we did not use to obtain the bound (3.6.2), and this randomness may be enough to control $(*)$ in most cases. To see this, recall that in bounding $(*)$ we require control over quantities of the form

$$\left| \rho\left(\alpha(\langle w_k, x - y_k \rangle - u_k)\right) - \rho\left(\alpha(\langle w_k, z - y_k \rangle - u_k)\right) \right|.$$

For most practical realizations of ρ , this difference will be small with high probability (on the draws of y_k, w_k, u_k) whenever $\|x - z\|_2$ is sufficiently small.

3.7 Theoretical results on submanifolds of Euclidean space

The constructions of RVFL networks presented in Theorems 3.3.1 and 3.3.4 depend heavily on the dimension of the ambient space \mathbb{R}^N . Indeed, the random variables used to construct the input-to-hidden layer weights and biases for these neural networks are N -dimensional objects; moreover, we saw that the lower bound on the number of nodes in Theorem 3.3.4 behaves like $n \gtrsim N \text{vol}(K) \varepsilon^{-1} \log(\text{vol}(K)/\varepsilon)$. If the ambient dimension is small, these dependencies do not present much of a problem. However, many modern applications require the ambient dimension to be large. Fortunately, a common assumption in practice is that signals of interest have structure (e.g., sparsity) that effectively reduces their complexity. Good theoretical results and algorithms typically depend on this intrinsic dimension rather than the ambient dimension. For this reason, it is desirable to obtain approximation results for RVFL networks that leverage the underlying structure of the signal class of interest, namely, the domain of $f \in C_c(\mathbb{R}^N)$.

One way to introduce lower-dimensional structure in the context of RVFL networks is to assume that $\text{supp}(f)$ lies on a subspace of \mathbb{R}^N . More generally, and motivated by applications, we may consider the case where $\text{supp}(f)$ is actually a submanifold \mathbb{R}^N . To this end, for the remainder of this section we take $\mathcal{M} \subset \mathbb{R}^N$ to be a smooth, compact d -dimensional manifold and consider the problem of approximating functions $f \in C(\mathcal{M})$ using RVFL networks. As we will see, RVFL networks in this setting yield theoretical guarantees that replace the dependencies of Theorems 3.3.1 and 3.3.4 on the ambient dimension N with dependencies on the manifold dimension d . Indeed, one might expect to see the random variables $\{w_k\}_{k=1}^n, \{b_k\}_{k=1}^n$ be d -dimensional objects (rather than N -dimensional) and that the lower bound on the number of network nodes in Theorem 3.3.4 scales like $n \gtrsim d \text{vol}(M) \varepsilon^{-1} \log(\text{vol}(M)/\varepsilon)$.

3.7.1 Adapting RVFL networks to d -manifolds

As in Section 3.4.2, let $\{(U_j, \phi_j)\}_{j \in J}$ be an atlas for the smooth, compact d -dimensional manifold $\mathcal{M} \subset \mathbb{R}^N$ with corresponding, compactly supported partition of unity $\{\eta_j\}_{j \in J}$. Since \mathcal{M} is compact, we assume without loss of generality that $|J| < \infty$; indeed, if the coordinate maps ϕ_j are sufficiently nice¹, then there exists an atlas for \mathcal{M} with $|J| \lesssim 2^d d \log(d) \text{vol}(\mathcal{M}) r^{-d}$, where $r = \sup_{j \in J} \text{rad}(U_j)$ is the largest chart radius [SCC18]. Now, for $f \in C(\mathcal{M})$, Lemma 3.4.5 implies that

$$f(x) = \sum_{\{j \in J: x \in U_j\}} (\hat{f}_j \circ \phi_j)(x) \quad (3.7.1)$$

for all $x \in \mathcal{M}$, where

$$\hat{f}_j(z) := \begin{cases} f(\phi_j^{-1}(z)) \eta_j(\phi_j^{-1}(z)) & z \in \phi_j(U_j) \\ 0 & \text{otherwise.} \end{cases}$$

As we will see, the fact that \mathcal{M} is smooth and compact implies $\hat{f}_j \in C_c(\mathbb{R}^d)$ for each $j \in J$, and so we may approximate each \hat{f}_j using RVFL networks on \mathbb{R}^d as in Theorems 3.3.1 and 3.3.4. In this way, it is reasonable to expect that f can be approximated on \mathcal{M} using a linear combination of these low-dimensional RVFL networks. To be clear, we propose approximating f on \mathcal{M} via the following process:

1. For each $j \in J$, approximate \hat{f}_j uniformly on $\phi_j(U_j) \subset \mathbb{R}^d$ using a RVFL network \tilde{f}_{n_j} as in Theorems 3.3.1 and 3.3.4;

¹For instance, one may construct the atlas $\{(U_j, \phi_j)\}_{j \in J}$ by intersecting \mathcal{M} with ℓ_2 balls in \mathbb{R}^N of sufficiently small radii so that each set U_j is diffeomorphic to a ℓ_2 ball in \mathbb{R}^d with coordinate map ϕ_j close to the identity.

2. Approximate f uniformly on \mathcal{M} by summing these RVFL networks over J , i.e.,

$$f(x) \approx \sum_{\{j \in J: x \in U_j\}} (\tilde{f}_{n_j} \circ \phi_j)(x)$$

for all $x \in \mathcal{M}$.

3.7.2 Main results on d -manifolds

Using the construction presented in Section 3.7.1, we now prove the manifold-equivalents of Theorems 3.3.1 and 3.3.4. For notational clarity, from here forward we use $\lim_{\{n_j\}_{j \in J} \rightarrow \infty}$ to mean the limit as each n_j tends to infinity simultaneously. The first theorem that we prove is an asymptotic approximation result for continuous functions on manifolds using the RVFL network construction presented in Section 3.7.1. This theorem is the manifold-equivalent of Theorem 3.3.1:

Theorem 3.7.1. *Let $\mathcal{M} \subset \mathbb{R}^N$ be a smooth, compact d -dimensional manifold with finite atlas $\{(U_j, \phi_j)\}_{j \in J}$ and $f \in C(\mathcal{M})$. Fix any activation function $\rho \in L^1(\mathbb{R}) \cap L^2(\mathbb{R})$. For any $\varepsilon > 0$, there exist constants $\alpha_j, \Omega_j > 0$ and hidden-to-output layer weights $\{v_k^{(j)}\}_{k=1}^{n_j} \subset \mathbb{R}$ for each $j \in J$ such that the following holds: If*

$$\begin{aligned} w_0^{(j)} &\sim \text{Unif}([- \alpha_j \Omega_j, \alpha_j \Omega_j])^d; \\ y_0^{(j)} &\sim \text{Unif}(\phi_j(U_j)); \\ u_0^{(j)} &\sim \text{Unif}([-\frac{\pi}{2}(2L_j + 1), \frac{\pi}{2}(2L_j + 1)]), \quad \text{where } L_j := \lceil \frac{2d}{\pi} \text{rad}(\phi_j(U_j)) \Omega_j - \frac{1}{2} \rceil; \\ b_0^{(j)} &:= -\langle w_0^{(j)}, y_0^{(j)} \rangle - \alpha_j u_0^{(j)}, \end{aligned}$$

and one chooses $\{w_k^{(j)}\}_{k=1}^{n_j}, \{b_k^{(j)}\}_{k=1}^{n_j}$ as independent draws from the distributions of $w_0^{(j)}$ and

$b_0^{(j)}$ for each $j \in J$, respectively, then the sequences of RVFL networks $\{\tilde{f}_{n_j}\}_{n_j=1}^\infty$ defined by

$$\tilde{f}_{n_j}(z) := \sum_{k=1}^{n_j} v_k^{(j)} \rho(\langle w_k^{(j)}, z \rangle + b_k^{(j)}), \quad \text{for } z \in \phi_j(U_j)$$

satisfy

$$\lim_{\{n_j\}_{j \in J} \rightarrow \infty} \mathbb{E} \int_{\mathcal{M}} \left| f(x) - \sum_{\{j \in J: x \in U_j\}} (\tilde{f}_{n_j} \circ \phi_j)(x) \right|^2 dx < \varepsilon$$

with convergence rate $O(1/\min_{j \in J} n_j)$.

Proof. We wish to show that there exist sequences of RVFL networks $\{\tilde{f}_{n_j}\}_{n_j=1}^\infty$ defined on $\phi_j(U_j)$ for each $j \in J$ which together satisfy the asymptotic error bound

$$\lim_{\{n_j\}_{j \in J} \rightarrow \infty} \mathbb{E} \int_{\mathcal{M}} \left| f(x) - \sum_{\{j \in J: x \in U_j\}} (\tilde{f}_{n_j} \circ \phi_j)(x) \right|^2 dx < \varepsilon.$$

We will do so by leveraging the result of Theorem 3.3.1 on each $\phi_j(U_j) \subset \mathbb{R}^d$.

To begin, recall that we may apply the representation (3.7.1) for f on each chart (U_j, ϕ_j) ; the RVFL networks \tilde{f}_{n_j} we seek are approximations of the functions \hat{f}_j in this expansion. Now, as $\text{supp}(\eta_j) \subset U_j$ is compact for each $j \in J$, it follows that each set $\phi_j(\text{supp}(\eta_j))$ is a compact subset of \mathbb{R}^d . Moreover, because $\hat{f}_j(z) \neq 0$ if and only if $z \in \phi_j(U_j)$ and $\phi_j^{-1}(z) \in \text{supp}(\eta_j) \subset U_j$, we have that $\hat{f}_j = \hat{f}_j|_{\phi_j(\text{supp}(\eta_j))}$ is supported on a compact set. Hence, $\hat{f}_j \in C_c(\mathbb{R}^d)$ for each $j \in J$, and so we may apply Lemma 3.5.6 to obtain the uniform limit representation (3.5.12) on $\phi_j(U_j)$, that is,

$$\hat{f}_j(z) = \lim_{\Omega_j \rightarrow \infty} \lim_{\alpha_j \rightarrow \infty} \int_{K(\Omega_j)} F_{\alpha_j, \Omega_j}(y, w, u) \rho(\alpha_j \langle w, z \rangle + b_{\alpha_j}(y, w, u)) dy dw du, \quad (3.7.2)$$

where we define $K(\Omega_j) := \phi_j(U_j) \times [-\Omega_j, \Omega_j]^d \times [-\frac{\pi}{2}(2L_j + 1), \frac{\pi}{2}(2L_j + 1)]$. In this way, as in

Section 3.5.4, by (3.5.12) we know that for any $\varepsilon_j > 0$ there exist $\alpha_j, \Omega_j > 0$ such that

$$|\hat{f}_j(z) - I^{(j)}(z; 1)| < \sqrt{\frac{\varepsilon_j}{\text{vol}(\phi_j(U_j))}} \quad (3.7.3)$$

holds for each $z \in \phi_j(U_j)$ simultaneously, where $I^{(j)}(\cdot; p)$ is as in (3.5.14), as well as the asymptotic error bound that is the final result of Theorem 3.3.1, namely

$$\lim_{n_j \rightarrow \infty} \mathbb{E} \int_{\phi_j(U_j)} |\hat{f}_j(z) - \tilde{f}_{n_j}(z)|^2 dz < \varepsilon_j. \quad (3.7.4)$$

With these results in hand, we may now continue with the main body of the proof.

Since the representation (3.7.1) for f on each chart (U_j, ϕ_j) yields

$$\left| f(x) - \sum_{\{j \in J: x \in U_j\}} (\tilde{f}_{n_j} \circ \phi_j)(x) \right| \leq \sum_{\{j \in J: x \in U_j\}} \left| (\hat{f}_j \circ \phi_j)(x) - (\tilde{f}_{n_j} \circ \phi_j)(x) \right|$$

for all $x \in \mathcal{M}$, the mean square error of our RVFL approximation may be bounded by

$$\begin{aligned} & \mathbb{E} \int_{\mathcal{M}} \left| f(x) - \sum_{\{j \in J: x \in U_j\}} (\tilde{f}_{n_j} \circ \phi_j)(x) \right|^2 dx \\ & \leq \underbrace{\mathbb{E} \int_{\mathcal{M}} \sum_{\{j \in J: x \in U_j\}} \left| (\hat{f}_j \circ \phi_j)(x) - (\tilde{f}_{n_j} \circ \phi_j)(x) \right|^2 dx}_{(*)} \\ & \quad + 2 \underbrace{\mathbb{E} \int_{\mathcal{M}} \sum_{\{j \neq k \in J: x \in U_j \cap U_k\}} \left((\hat{f}_j \circ \phi_j)(x) - (\tilde{f}_{n_j} \circ \phi_j)(x) \right) \left((\hat{f}_k \circ \phi_k)(x) - (\tilde{f}_{n_k} \circ \phi_k)(x) \right) dx}_{(**)}. \end{aligned} \quad (3.7.5)$$

To bound (*), note that the change of variables (3.4.2) implies

$$\int_{\mathcal{M}} \sum_{\{j \in J: x \in U_j\}} \left| (\hat{f}_j \circ \phi_j)(x) - (\tilde{f}_{n_j} \circ \phi_j)(x) \right|^2 dx = \sum_{j \in J} \int_{\phi_j(U_j)} \frac{|\hat{f}_j(z) - \tilde{f}_{n_j}(z)|^2}{|\det(D\phi_j(\phi_j^{-1}(z)))|} dz$$

for each $j \in J$. Defining $\beta_j := \inf_{y \in U_j} |\det(D\phi_j(y))|$, which is necessarily bounded away from zero for each $j \in J$ by compactness of \mathcal{M} , we therefore have

$$(*) \leq \sum_{j \in J} \beta_j^{-1} \mathbb{E} \int_{\phi_j(U_j)} |\hat{f}_j(z) - \tilde{f}_{n_j}(z)|^2 dz.$$

Hence, applying (3.7.4) for each $j \in J$ yields

$$\lim_{n_j \rightarrow \infty} (*) \leq \sum_{j \in J} \beta_j^{-1} \lim_{n_j \rightarrow \infty} \mathbb{E} \int_{\phi_j(U_j)} |\hat{f}_j(z) - \tilde{f}_{n_j}(z)|^2 dz < \sum_{j \in J} \frac{\varepsilon_j}{\beta_j} \quad (3.7.6)$$

with convergence rate $O(1/n_j)$. For the term (**), we first use Fubini's Theorem to swap the order of integrals and then appeal to independence of the random variables \tilde{f}_{n_j} and \tilde{f}_{n_k} for $j \neq k \in J$, giving us

$$\begin{aligned} (**) &= \sum_{j \neq k \in J} \int_{U_j \cap U_k} \mathbb{E} \left((\hat{f}_j \circ \phi_j)(x) - (\tilde{f}_{n_j} \circ \phi_j)(x) \right) \mathbb{E} \left((\hat{f}_k \circ \phi_k)(x) - (\tilde{f}_{n_k} \circ \phi_k)(x) \right) dx \\ &= \sum_{j \neq k \in J} \int_{U_j \cap U_k} \left((\hat{f}_j \circ \phi_j)(x) - I^{(j)}(\phi_j(x); 1) \right) \left((\hat{f}_k \circ \phi_k)(x) - I^{(k)}(\phi_k(x); 1) \right) dx. \end{aligned}$$

Since the Cauchy-Schwarz inequality yields

$$\begin{aligned} &\int_{U_j \cap U_k} \left((\hat{f}_j \circ \phi_j)(x) - I^{(j)}(\phi_j(x); 1) \right) \left((\hat{f}_k \circ \phi_k)(x) - I^{(k)}(\phi_k(x); 1) \right) dx \\ &\leq \left(\int_{U_j} \left| (\hat{f}_j \circ \phi_j)(x) - I^{(j)}(\phi_j(x); 1) \right|^2 dx \right)^{1/2} \left(\int_{U_k} \left| (\hat{f}_k \circ \phi_k)(x) - I^{(k)}(\phi_k(x); 1) \right|^2 dx \right)^{1/2} \end{aligned}$$

for $j \neq k \in J$, another application of the change of variables (3.4.2) allows us to write

$$(**) \leq \sum_{j \neq k \in J} \left(\int_{\phi_j(U_j)} \frac{|\hat{f}_j(z) - I^{(j)}(z; 1)|^2}{|D\phi_j(\phi_j^{-1}(z))|} dz \right)^{1/2} \left(\int_{\phi_k(U_k)} \frac{|\hat{f}_k(z) - I^{(k)}(z; 1)|^2}{|D\phi_k(\phi_k^{-1}(z))|} dz \right)^{1/2}.$$

Combining (3.7.3) with the notation $\beta_j := \inf_{y \in U_j} |\det(D\phi_j(y))|$, it follows that

$$(**) < \sum_{j \neq k \in J} \sqrt{\frac{\varepsilon_j \varepsilon_k}{\beta_j \beta_k}}, \quad (3.7.7)$$

which is independent of n_j and n_k .

With the bounds (3.7.6) and (3.7.7) in hand, taking limits in (3.7.5) yields

$$\lim_{\{n_j\}_{j \in J} \rightarrow \infty} \mathbb{E} \int_{\mathcal{M}} \left| f(x) - \sum_{\{j \in J: x \in U_j\}} (\tilde{f}_{n_j} \circ \phi_j)(x) \right|^2 dx < \sum_{j \in J} \frac{\varepsilon_j}{\beta_j} + 2 \sum_{j \neq k \in J} \sqrt{\frac{\varepsilon_j \varepsilon_k}{\beta_j \beta_k}} = \left(\sum_{j \in J} \sqrt{\frac{\varepsilon_j}{\beta_j}} \right)^2$$

with convergence rate $O(1/\min_{j \in J} n_j)$, and so the proof is completed by taking each $\varepsilon_j > 0$ in such a way that

$$\varepsilon = \left(\sum_{j \in J} \sqrt{\frac{\varepsilon_j}{\beta_j}} \right)^2,$$

and choosing $\alpha_j, \Omega_j > 0$ accordingly for each $j \in J$.

It remains only to verify our use of Fubini's Theorem in bounding (3.7.7). To this end, recall from (3.5.20) that

$$\mathbb{E} \lim_{n_j \rightarrow \infty} \left| (\hat{f}_j \circ \phi_j)(x) - (\tilde{f}_{n_j} \circ \phi_j)(x) \right| \leq \sigma_j(\phi_j(x)) \sqrt{\frac{2}{\pi}}$$

for each $x \in U_j$. Hence, an application of the Cauchy-Schwarz inequality implies

$$\begin{aligned} & \int_{U_j \cap U_k} \mathbb{E} \lim_{n_j, n_k \rightarrow \infty} \left| (\hat{f}_j \circ \phi_j)(x) - (\tilde{f}_{n_j} \circ \phi_j)(x) \right| \left| (\hat{f}_k \circ \phi_k)(x) - (\tilde{f}_{n_k} \circ \phi_k)(x) \right| dx \\ & \leq \frac{2}{\pi} \int_{U_j \cap U_k} \sigma_j(\phi_j(x)) \sigma_k(\phi_k(x)) dx \\ & \leq \frac{2}{\pi} \left(\int_{U_j} \sigma_j(\phi_j(x))^2 dx \right)^{1/2} \left(\int_{U_k} \sigma_k(\phi_k(x))^2 dx \right)^{1/2}. \end{aligned}$$

Combining this with (3.5.16), we obtain the bound

$$\int_{U_j} \sigma_j(\phi_j(x))^2 dx \leq \frac{\alpha_j^2 M_j^2 \|\rho\|_2^2 \text{vol}(U_j)}{2^{2d} \text{vol}(\phi_j(U_j))},$$

and so it follows that

$$\begin{aligned} & \int_{U_j \cap U_k} \mathbb{E} \lim_{n_j, n_k \rightarrow \infty} \left| (\hat{f}_j \circ \phi_j)(x) - (\tilde{f}_{n_j} \circ \phi_j)(x) \right| \left| (\hat{f}_k \circ \phi_k)(x) - (\tilde{f}_{n_k} \circ \phi_k)(x) \right| dx \\ & \leq \frac{\alpha_j \alpha_k M_j M_k \|\rho\|_2^2}{2^{2d-1} \pi} \sqrt{\frac{\text{vol}(U_j) \text{vol}(U_k)}{\text{vol}(\phi_j(U_j)) \text{vol}(\phi_k(U_k))}} \end{aligned}$$

holds for all $j \neq k \in J$, which is necessarily finite. Hence, we may apply Fubini's Theorem and the Dominated Convergence Theorem to obtain

$$\begin{aligned} & \int_{U_j \cap U_k} \mathbb{E} \lim_{n_j, n_k \rightarrow \infty} \left((\hat{f}_j \circ \phi_j)(x) - (\tilde{f}_{n_j} \circ \phi_j)(x) \right) \left((\hat{f}_k \circ \phi_k)(x) - (\tilde{f}_{n_k} \circ \phi_k)(x) \right) dx \\ & = \lim_{n_j, n_k \rightarrow \infty} \int_{U_j \cap U_k} \mathbb{E} \left((\hat{f}_j \circ \phi_j)(x) - (\tilde{f}_{n_j} \circ \phi_j)(x) \right) \left((\hat{f}_k \circ \phi_k)(x) - (\tilde{f}_{n_k} \circ \phi_k)(x) \right) dx \\ & = \lim_{n_j, n_k \rightarrow \infty} \mathbb{E} \int_{U_j \cap U_k} \left((\hat{f}_j \circ \phi_j)(x) - (\tilde{f}_{n_j} \circ \phi_j)(x) \right) \left((\hat{f}_k \circ \phi_k)(x) - (\tilde{f}_{n_k} \circ \phi_k)(x) \right) dx \end{aligned}$$

for all $j \neq k \in J$, as desired. \square

The biggest takeaway from Theorem 3.7.1 is that the same asymptotic mean-square error behavior we saw in the RVFL network architecture of Theorem 3.3.1 holds for our RVFL network construction on manifolds, with the added benefit that the input-to-hidden layer weights and biases are now d -dimensional random variables rather than N -dimensional. Provided the size of the atlas $|J|$ isn't too large, this significantly reduces the number of random variables that must be generated to produce a uniform approximation of $f \in C(\mathcal{M})$.

One might expect to see a similar reduction in dimension dependence for the non-asymptotic case if the RVFL network construction of Section 3.7.1 is used. Indeed, our next

theorem, which is the manifold-equivalent of Theorem 3.3.4, makes this explicit:

Theorem 3.7.2. *Let $\mathcal{M} \subset \mathbb{R}^N$ be a smooth, compact d -dimensional manifold with finite atlas $\{(U_j, \phi_j)\}_{j \in J}$ and $f \in C(\mathcal{M})$. Fix any activation function $\rho \in L^1(\mathbb{R}) \cap L^2(\mathbb{R})$ such that ρ is κ -Lipschitz on \mathbb{R} for some $\kappa > 0$. For any $\varepsilon > 0$ and $\eta \in (0, 1)$, there exist constants $\alpha_j, \Omega_j > 0$ and hidden-to-output layer weights $\{v_k^{(j)}\}_{k=1}^{n_j} \subset \mathbb{R}$ for each $j \in J$ such that the following holds:
Suppose*

$$\begin{aligned} w_0^{(j)} &\sim \text{Unif}([- \alpha_j \Omega_j, \alpha_j \Omega_j])^d; \\ y_0^{(j)} &\sim \text{Unif}(\phi_j(U_j)); \\ u_0^{(j)} &\sim \text{Unif}([- \frac{\pi}{2}(2L_j + 1), \frac{\pi}{2}(2L_j + 1)]), \quad \text{where } L_j := \lceil \frac{2d}{\pi} \text{rad}(\phi_j(U_j)) \Omega_j - \frac{1}{2} \rceil; \\ b_0^{(j)} &:= -\langle w_0^{(j)}, y_0^{(j)} \rangle - \alpha_j u_0^{(j)}, \end{aligned}$$

and one chooses $\{w_k^{(j)}\}_{k=1}^{n_j}, \{b_k^{(j)}\}_{k=1}^{n_j}$ as independent draws from the distributions of $w_0^{(j)}$ and $b_0^{(j)}$ for each $j \in J$, respectively. For any

$$0 < \delta_j < \frac{\sqrt{\varepsilon}}{4\sqrt{2d|J|\text{vol}(\mathcal{M})\kappa\alpha_j^2 M_j \Omega_j^{d+2} \text{vol}(\phi_j(U_j))(1 + 2d \text{rad}(\phi_j(U_j)))}},$$

if one chooses

$$n_j \geq \frac{4\sqrt{|J|\text{vol}(\mathcal{M})}C^{(j)}c \log(3|J|\eta^{-1}\mathcal{N}(\delta_j, \phi_j(U_j)))}{\sqrt{\varepsilon} \log\left(1 + \frac{C^{(j)}\sqrt{\varepsilon}}{8\sqrt{|J|\text{vol}(\mathcal{M})}d(2\Omega_j)^{d+1}\text{rad}(\phi_j(U_j))\text{vol}^2(\phi_j(U_j))\Sigma^{(j)}}\right)},$$

where $M_j := \sup_{z \in \phi_j(U_j)} |\hat{f}_j(z)|$, $c > 0$ is a numerical constant, and $C^{(j)}, \Sigma^{(j)}$ are constants depending on \hat{f}_j and ρ for each $j \in J$, then the sequences of RVFL networks $\{\tilde{f}_{n_j}\}_{n_j=1}^{\infty}$ defined by

$$\tilde{f}_{n_j}(z) := \sum_{k=1}^{n_j} v_k^{(j)} \rho(\langle w_k^{(j)}, z \rangle + b_k^{(j)}), \quad \text{for } z \in \phi_j(U_j)$$

satisfy

$$\int_{\mathcal{M}} \left| f(x) - \sum_{\{j \in J: x \in U_j\}} (\tilde{f}_{n_j} \circ \phi_j)(x) \right|^2 dx < \varepsilon$$

with probability at least $1 - \eta$.

Proof. We wish to show that there exist sequences of RVFL networks $\{\tilde{f}_{n_j}\}_{n_j=1}^{\infty}$ defined on $\phi_j(U_j)$ for each $j \in J$ which together satisfy the error bound

$$\int_{\mathcal{M}} \left| f(x) - \sum_{\{j \in J: x \in U_j\}} (\tilde{f}_{n_j} \circ \phi_j)(x) \right|^2 dx < \varepsilon$$

with probability at least $1 - \eta$ when $\{n_j\}_{j \in J}$ are chosen sufficiently large. The proof is obtained by showing that

$$\left| f(x) - \sum_{\{j \in J: x \in U_j\}} (\tilde{f}_{n_j} \circ \phi_j)(x) \right| < \sqrt{\frac{\varepsilon}{\text{vol}(\mathcal{M})}} \quad (3.7.8)$$

holds uniformly for $x \in \mathcal{M}$ with high probability.

We begin as in the proof of Theorem 3.7.1 by applying the representation (3.7.1) for f on each chart (U_j, ϕ_j) , which gives us

$$\left| f(x) - \sum_{\{j \in J: x \in U_j\}} (\tilde{f}_{n_j} \circ \phi_j)(x) \right| \leq \sum_{\{j \in J: x \in U_j\}} \left| (\hat{f}_j \circ \phi_j)(x) - (\tilde{f}_{n_j} \circ \phi_j)(x) \right| \quad (3.7.9)$$

for all $x \in \mathcal{M}$. Now, since we have already seen that $\hat{f}_j \in C_c(\mathbb{R}^d)$ for each $j \in J$, Theorem 3.3.4 implies that for any $\varepsilon_j > 0$, there exist constants $\alpha_j, \Omega_j > 0$ and hidden-to-output layer weights $\{v_k^{(j)}\}_{k=1}^{n_j} \subset \mathbb{R}$ for each $j \in J$ such that for any

$$\delta_j < \frac{\sqrt{\varepsilon_j}}{4\sqrt{d}\kappa\alpha_j^2 M_j \Omega_j^{d+2} \text{vol}^{3/2}(\phi_j(U_j))(1 + 2d\text{rad}(\phi_j(U_j)))} \quad (3.7.10)$$

we have

$$\left| \hat{f}_j(z) - \tilde{f}_{n_j}(z) \right| < \sqrt{\frac{2\varepsilon_j}{\text{vol}(\phi_j(U_j))}}$$

uniformly for all $z \in \phi_j(U_j)$ with probability at least $1 - \eta_j$, provided the number of nodes n_j satisfies

$$n_j \geq \frac{2\sqrt{2\text{vol}(\phi_j(U_j))}C^{(j)}c \log(3\eta_j^{-1}\mathcal{N}(\delta_j, \phi_j(U_j)))}{\sqrt{\varepsilon_j} \log\left(1 + \frac{C^{(j)}\sqrt{\varepsilon_j}}{4\sqrt{2d}(2\Omega_j)^{d+1}\text{rad}(\phi_j(U_j))\text{vol}^{5/2}(\phi_j(U_j))^{\Sigma^{(j)}}}\right)}, \quad (3.7.11)$$

where $c > 0$ is a numerical constant and $C^{(j)}, \Sigma^{(j)}$ are as in (3.6.5). Indeed, it suffices to choose the coefficients

$$v_k^{(j)} := \frac{\text{vol}(K(\Omega_j))}{n_j} F_{\alpha_j, \Omega_j} \left(y_k^{(j)}, \frac{w_k^{(j)}}{\alpha_j^d}, u_k^{(j)} \right)$$

for each $k = 1, \dots, n_j$, where

$$K(\Omega_j) := \phi_j(U_j) \times [-\alpha_j\Omega_j, \alpha_j\Omega_j]^d \times \left[-\frac{\pi}{2}(2L_j + 1), \frac{\pi}{2}(2L_j + 1)\right]$$

for each $j \in J$. Combined with (3.7.9), choosing δ_j and n_j satisfying (3.7.10) and (3.7.11), respectively, then yields

$$\left| f(x) - \sum_{\{j \in J: x \in U_j\}} (\tilde{f}_{n_j} \circ \phi_j)(x) \right| < \sum_{\{j \in J: x \in U_j\}} \sqrt{\frac{2\varepsilon_j}{\text{vol}(\phi_j(U_j))}} \leq \sum_{j \in J} \sqrt{\frac{2\varepsilon_j}{\text{vol}(\phi_j(U_j))}}$$

for all $x \in \mathcal{M}$ with probability at least $1 - \sum_{\{j \in J: x \in U_j\}} \eta_j \geq 1 - \sum_{j \in J} \eta_j$. Since we require that (3.7.8) holds for all $x \in \mathcal{M}$ with probability at least $1 - \eta$, the proof is then completed by

choosing $\{\varepsilon_j\}_{j \in J}$ and $\{\eta_j\}_{j \in J}$ such that

$$\varepsilon = 2\text{vol}(\mathcal{M}) \left(\sum_{j \in J} \sqrt{\frac{\varepsilon_j}{\text{vol}(\phi_j(U_j))}} \right)^2 \quad \text{and} \quad \eta = \sum_{j \in J} \eta_j.$$

In particular, it suffices to choose

$$\varepsilon_j = \frac{\text{vol}(\phi_j(U_j)) \varepsilon}{2|J|\text{vol}(\mathcal{M})}$$

and $\eta_j = \eta/|J|$ for each $j \in J$, so that (3.7.10) and (3.7.11) become

$$\begin{aligned} \delta_j &< \frac{\sqrt{\varepsilon}}{4\sqrt{2d|J|\text{vol}(\mathcal{M})\kappa\alpha_j^2 M_j \Omega_j^{d+2} \text{vol}(\phi_j(U_j))(1+2d\text{rad}(\phi_j(U_j)))}}, \\ n_j &\geq \frac{4\sqrt{|J|\text{vol}(\mathcal{M})} C^{(j)} c \log(3|J|\eta^{-1} \mathcal{N}(\delta_j, \phi_j(U_j)))}{\sqrt{\varepsilon} \log\left(1 + \frac{C^{(j)} \sqrt{\varepsilon}}{8\sqrt{|J|\text{vol}(\mathcal{M})} d(2\Omega_j)^{d+1} \text{rad}(\phi_j(U_j)) \text{vol}^2(\phi_j(U_j)) \Sigma^{(j)}}\right)}, \end{aligned}$$

as desired. □

As alluded to earlier, an important implication of Theorems 3.7.1 and 3.7.2 is that the random variables $\{w_k^{(j)}\}_{k=1}^{n_j}$ and $\{b_k^{(j)}\}_{k=1}^{n_j}$ are d -dimensional objects for each $j \in J$. Likewise, for small ε , Theorem 3.7.2 shows that the number of nodes behaves roughly like

$$n_j \gtrsim d|J|\text{vol}(\mathcal{M})\varepsilon^{-1} \log(\text{vol}(\mathcal{M})/\varepsilon)$$

for each $j \in J$. Thus, introducing the manifold structure removes the dependencies on the ambient dimension N , replacing them instead with the intrinsic dimension of \mathcal{M} and the complexity of the atlas $\{(U_j, \phi_j)\}_{j \in J}$.

Remark 3.7.3. The bounds on the covering radii δ_j and hidden layer nodes n_j needed for each chart in Theorem 3.7.2 are not optimal. Indeed, these bounds may be further improved if one uses the local structure of the manifold, through quantities such as its *curvature* and *reach*. In particular,

the appearance of $|J|$ in both bounds may be significantly improved upon if the manifold is locally well-behaved.

3.8 Numerical Simulations

In this section we provide numerical evidence to support the result of Theorem 3.7.2. Let $\mathcal{M} \subset \mathbb{R}^N$ be a smooth, compact d -dimensional manifold. Since having access to an atlas for \mathcal{M} is not necessarily practical, we assume instead that we have a suitable approximation to \mathcal{M} . For our purposes, we will use a Geometric Multi-Resolution Analysis (GMRA) [ACM11] approximation of \mathcal{M} , as previously used in Chapter 2. Recall from Definition 2.3.3 that such a GMRA approximation provides a collection $\{(\mathcal{C}_j, \mathcal{P}_j)\}_{j \in \{1, \dots, J\}}$ of centers $\mathcal{C}_j = \{c_{j,k}\}_{k=1}^{K_j} \subset \mathbb{R}^N$ and affine projections $\mathcal{P}_j = \{P_{j,k}\}_{k=1}^{K_j}$ on \mathbb{R}^N such that, for each $j \in \{1, \dots, J\}$, the pairs $\{(c_{j,k}, P_{j,k})\}_{k=1}^{K_j}$ define d -dimensional affine spaces that approximate \mathcal{M} with increasing accuracy in the following sense: For every $x \in \mathcal{M}$, there exists $\tilde{C}_x > 0$ and $k' \in \{1, \dots, K_j\}$ such that

$$\|x - P_{j,k'}x\|_2 \leq \tilde{C}_x 2^{-j} \quad (3.8.1)$$

holds whenever $\|x - c_{j,k'}\|_2$ is sufficiently small (see part (3.b) of Definition 2.3.3). In this way, a GMRA approximation of \mathcal{M} essentially provides a collection of approximate tangent spaces to \mathcal{M} . Hence, a GMRA approximation having fine enough resolution (i.e., large enough j) is a good substitution for an atlas.

Let $\{(c_{j,k}, P_{j,k})\}_{k=1}^{K_j}$ be a GMRA approximation of \mathcal{M} for refinement level $j \geq j_0$. Since the affine spaces defined by $(c_{j,k}, P_{j,k})$ for each $k \in \{1, \dots, K_j\}$ are d -dimensional, we will approximate f on \mathcal{M} by projecting it (in an appropriate sense) onto these affine spaces and approximating each projection using an RVFL network on \mathbb{R}^d . To make this more precise observe that, since each affine projection acts on $x \in \mathcal{M}$ as $P_{j,k}x = c_{j,k} + \Phi_{j,k}(x - c_{j,k})$ for some orthogonal

projection $\Phi_{j,k}: \mathbb{R}^N \rightarrow \mathbb{R}^N$, for each $k \in \{1, \dots, K_j\}$ we have

$$f(P_{j,k}x) = f(c_{j,k} + \Phi_{j,k}(x - c_{j,k})) = f((I_N - \Phi_{j,k})c_{j,k} + U_{j,k}D_{j,k}V_{j,k}^T x),$$

where $\Phi_{j,k} = U_{j,k}D_{j,k}V_{j,k}^T$ is the compact singular value decomposition (SVD) of $\Phi_{j,k}$ (i.e., only the left and right singular vectors corresponding to nonzero singular values are computed). In particular, the matrix of right-singular vectors $V_{j,k}: \mathbb{R}^d \rightarrow \mathbb{R}^N$ enables us to define a function $\hat{f}_{j,k}: \mathbb{R}^d \rightarrow \mathbb{R}$, given by

$$\hat{f}_{j,k}(z) := f((I_N - \Phi_{j,k})c_{j,k} + U_{j,k}D_{j,k}z), \quad z \in \mathbb{R}^d, \quad (3.8.2)$$

which satisfies $\hat{f}_{j,k}(V_{j,k}^T x) = f(P_{j,k}x)$ for all $x \in \mathcal{M}$. By continuity of f and (3.8.1), this means that for any $\varepsilon > 0$ there exists $j \geq j_0$ such that $|f(P_{j,k}x) - \hat{f}_{j,k}(V_{j,k}^T x)| < \varepsilon$ for some $k \in \{1, \dots, K_j\}$. For such $k \in \{1, \dots, K_j\}$, we may therefore approximate f on the affine space associated with $(c_{j,k}, P_{j,k})$ by approximating $\hat{f}_{j,k}$ using a RVFL network $\tilde{f}_{n_{j,k}}: \mathbb{R}^d \rightarrow \mathbb{R}$ of the form

$$\tilde{f}_{n_{j,k}}(z) := \sum_{\ell=1}^{n_{j,k}} v_{\ell}^{(j,k)} \rho(\langle w_{\ell}^{(j,k)}, z \rangle + b_{\ell}^{(j,k)}), \quad (3.8.3)$$

where $\{w_{\ell}^{(j,k)}\}_{\ell=1}^{n_{j,k}} \subset \mathbb{R}^d$ and $\{b_{\ell}^{(j,k)}\}_{\ell=1}^{n_{j,k}} \subset \mathbb{R}$ are random input-to-hidden layer weights and biases (resp.) and the hidden-to-output layer weights $\{v_{\ell}^{(j,k)}\}_{\ell=1}^{n_{j,k}} \subset \mathbb{R}$ are learned. Choosing the random input-to-hidden layer weights and biases as in Theorem 3.3.4 then guarantees that $|f(P_{j,k}x) - \tilde{f}_{n_{j,k}}(V_{j,k}^T x)|$ is small with high probability whenever $n_{j,k}$ is sufficiently large.

In light of the above discussion, we propose the following RVFL network construction for approximating functions $f \in C(\mathcal{M})$: Given a GMRA approximation of \mathcal{M} with resolution $j \geq j_0$, construct and train RVFL networks of the form (3.8.3) for each $k \in \{1, \dots, K_j\}$. Then,

Algorithm 2 Approximation Algorithm

Given: $f \in C(\mathcal{M})$; GMRA approximation $\{(c_{j,k}, P_{j,k})\}_{k=1}^{K_j}$ of \mathcal{M} at scale $j \geq j_0$

Ensure: $y^\sharp \approx f(x)$ for any $x \in \mathcal{M}$

Step 1: For each $k \in \{1, \dots, K_j\}$, construct and train a RVFL network $\tilde{f}_{n_{j,k}}$ of the form (3.8.3)

Step 2: For any $x \in \mathcal{M}$, find $c_{j,k'} \in \arg \min_{c_{j,k} \in \mathcal{C}_j} \|x - c_{j,k}\|_2$

Step 3: Set $y^\sharp = \tilde{f}_{n_{j,k'}}$

given $x \in \mathcal{M}$ and $\varepsilon > 0$, choose $k' \in \{1, \dots, K_j\}$ such that

$$c_{j,k'} \in \arg \min_{c_{j,k} \in \mathcal{C}_j} \|x - c_{j,k}\|_2$$

and evaluate $\tilde{f}_{n_{j,k'}}(x)$ to approximate $f(x)$. We summarize this algorithm in Algorithm 2. Since part (3.b) of Definition 2.3.3 implies $\|x - P_{j,k'}x\|_2 \leq C_x 2^{-2j}$ holds for our choice of $k' \in \{1, \dots, K_j\}$, for large enough j , continuity of f and Lemma 3.5.8 imply that

$$|f(x) - \tilde{f}_{n_{j,k'}}(x)| \leq |f(x) - \hat{f}_{j,k'}(V_{j,k'}^T x)| + |\hat{f}_{j,k'}(V_{j,k'}^T x) - \tilde{f}_{n_{j,k'}}(V_{j,k'}^T x)| < \varepsilon$$

holds with high probability, provided $n_{j,k'}$ satisfies the requirements of Theorem 3.3.4.

As a technical point, in the RVFL network construction proposed above we require that the function f be defined in a sufficiently large region around the manifold. Essentially, we need to ensure that f is continuously defined on the set $S := \mathcal{M} \cup \widehat{\mathcal{M}}_j$, where $\widehat{\mathcal{M}}_j$ is the scale- j GMRA approximation defined in (2.4.1), that is,

$$\widehat{\mathcal{M}}_j := \{P_{j,k_j(z)}z : \|z\|_2 \leq \text{rad}(\mathcal{M})\} \cap B_2^N(0, \text{rad}(\mathcal{M})).$$

This ensures that f can be evaluated on the affine subspaces given by the GMRA.

To simulate Algorithm 2, we take $\mathcal{M} = \mathbb{S}^2$ embedded in \mathbb{R}^{20} and construct a GMRA up to level $j_{\max} = 15$ using 20,000 data points sampled uniformly from \mathcal{M} . Given $j \leq j_{\max}$, we generate

RVFL networks $\hat{f}_{n_{j,k}}: \mathbb{R}^2 \rightarrow \mathbb{R}$ as in (3.8.3) and train them on $V_{j,k}^T(B_2^N(c_{j,k}, r) \cap T_{j,k}) \subset \mathbb{R}^d$ using the training pairs $\{(V_{k,j}^T x_\ell, f(P_{j,k} x_\ell))\}_{\ell=1}^p$, where $T_{k,j}$ is the affine space generated by $(c_{j,k}, P_{j,k})$. For simplicity, we fix $n_{j,k} = n$ to be constant for all $k \in \{1, \dots, K_j\}$ and use a single, fixed pair of parameters $\alpha, \Omega > 0$ when constructing all RVFL networks. We then randomly select a test set of 200 points $x \in \mathcal{M}$ for use throughout all experiments. In each experiment (i.e., point in Figure 3.8.1), we use Algorithm 2 to produce an approximation $y^\sharp = \tilde{f}_{n_{j,k'}}(x)$ of $f(x)$. Figure 3.8.1 displays the mean relative error in these approximations for varying numbers of nodes n ; to construct this plot, f is taken to be the exponential $f(x) = \exp(\sum_{k=1}^N x(k))$ and ρ the hyperbolic secant function. Notice that for small numbers of nodes the RVFL networks are not very good at approximating f , regardless of the choice of $\alpha, \Omega > 0$. However, the error decays as the number of nodes increases until reaching a floor due to error inherent in the GMRA approximation.




Figure 3.8.1: Log-scale plot of average relative error for Algorithm 2 as a function of the number of nodes n in each RVFL network. Black, blue, and red lines correspond to GMRA refinement levels $j = 12$, $j = 9$, and $j = 6$ (resp.). For each j , we fix $\alpha = 2$ and vary $\Omega = 10, 15$ (solid and dashed lines, resp.). Reconstruction error decays as a function of n until reaching a floor due to error in the GMRA approximation of \mathcal{M} .

3.9 Acknowledgements

A journal version of the material presented in this chapter is currently being prepared for submission for publication. The dissertation author was the primary researcher and author of this material. Thanks are extended to coauthors Deanna Needell, Rayan Saab, and Palina Salanevich for their contributions and insights aiding in the development of this material.

Bibliography

- [AC09] Nir Ailon and Bernard Chazelle. The fast johnson–lindenstrauss transform and approximate nearest neighbors. *SIAM Journal on computing*, 39(1):302–322, 2009.
- [ACM11] William K. Allard, Guangliang Chen, and Mauro Maggioni. Multi-scale geometric methods for data sets ii: Geometric multi-resolution analysis. 2011.
- [ACM12] William K. Allard, Guangliang Chen, and Mauro Maggioni. Multi-scale geometric methods for data sets ii: Geometric multi-resolution analysis. *Applied and Computational Harmonic Analysis*, 32(3):435 – 462, 2012.
- [AL13] Nir Ailon and Edo Liberty. An almost optimal unrestricted fast johnson-lindenstrauss transform. *ACM Transactions on Algorithms (TALG)*, 9(3):21, 2013.
- [BDDW08] Richard Baraniuk, Mark Davenport, Ronald DeVore, and Michael Wakin. A simple proof of the restricted isometry property for random matrices. *Constructive Approximation*, 28(3):253–263, 2008.
- [BJKS14] Petros T. Boufounos, Laurent Jacques, Felix Kraher, and Rayan Saab. Quantization and compressive sensing. *preprint arXiv:1405.1194*, 2014.
- [Bou14] Jean Bourgain. An improved estimate in the restricted isometry problem. In *Geometric aspects of functional analysis*, pages 65–70. Springer, 2014.
- [BS95] Hubert A.B. Te Braake and Gerrit Van Straten. Random activation weight neural net (rawn) for fast non-iterative training. *Engineering Applications of Artificial Intelligence*, 8(1):71 – 80, 1995.
- [BWD⁺06] Richard G Baraniuk, Michael Wakin, Marco F Duarte, Joel A Tropp, and Dror Baron. Random filters for compressive sampling and reconstruction. In *IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, volume 3, pages III–872, 2006.
- [CG16] Evan Chou and C Sinan Güntürk. Distributed noise-shaping quantization: I. beta duals of finite frames and near-optimal quantization of random measurements. *Constructive Approximation*, 44(1):1–22, 2016.

- [CGK⁺15] Evan Chou, C Sinan Güntürk, Felix Krahmer, Rayan Saab, and Özgür Yılmaz. Noise-shaping quantization methods for frame-based and compressive sampling systems. In *Sampling theory, a renaissance*, pages 157–184. Springer, 2015.
- [CGV13] Mahdi Cheraghchi, Venkatesan Guruswami, and Ameya Velingker. Restricted isometry of fourier matrices and list decodability of random linear codes. *SIAM Journal on Computing*, 42(5):1888–1914, 2013.
- [Cho13] Evan Chou. *Beta-duals of frames and applications to problems in quantization*. PhD thesis, New York University, 2013.
- [CRT06] Emmanuel J Candes, Justin K Romberg, and Terence Tao. Stable signal recovery from incomplete and inaccurate measurements. *Communications on Pure and Applied Mathematics: A Journal Issued by the Courant Institute of Mathematical Sciences*, 59(8):1207–1223, 2006.
- [CT06] Emmanuel J Candes and Terence Tao. Near-optimal signal recovery from random projections: Universal encoding strategies? *IEEE transactions on information theory*, 52(12):5406–5425, 2006.
- [CW99] C. L. P. Chen and J. Z. Wan. A rapid learning and dynamic stepwise updating algorithm for flat neural networks and the application to time-series prediction. *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)*, 29(1):62–72, Feb 1999.
- [D⁺06] David L Donoho et al. Compressed sensing. *IEEE Transactions on information theory*, 52(4):1289–1306, 2006.
- [DD03] Ingrid Daubechies and Ron DeVore. Approximating a bandlimited function using very coarsely quantized data: A family of stable sigma-delta modulators of arbitrary order. *Annals of mathematics*, pages 679–710, 2003.
- [DJR17] Sjoerd Dirksen, Hans Christian Jung, and Holger Rauhut. One-bit compressed sensing with partial gaussian circulant matrices. *arXiv preprint arXiv:1710.03287*, 2017.
- [DKG11] Percy Deift, Felix Krahmer, and C Sinan Güntürk. An optimal family of exponentially accurate one-bit sigma-delta quantization schemes. *Communications on Pure and Applied Mathematics*, 64(7):883–919, 2011.
- [DKS13] Josef Dick, Frances Y Kuo, and Ian H Sloan. High-dimensional integration: the quasi-monte carlo way. *Acta Numerica*, 22:133–288, 2013.
- [DMSP18] Yajnaseni Dash, Saroj Kanta Mishra, Sandeep Sahany, and Bijaya Ketan Panigrahi. Indian summer monsoon rainfall prediction: A comparison of iterative and non-iterative approaches. *Applied Soft Computing*, 70:1122 – 1134, 2018.

- [FKS17] Joe-Mei Feng, Felix Krahmer, and Rayan Saab. Quantized compressed sensing for partial random circulant matrices. In *2017 International Conference on Sampling Theory and Applications (SampTA)*, pages 236–240. IEEE, 2017.
- [FR17] Simon Foucart and Holger Rauhut. A mathematical introduction to compressive sensing. *Bull. Am. Math.*, 54:151–165, 2017.
- [GB98] Guang-Bin Huang and H. A. Babri. Upper bounds on the number of hidden neurons in feedforward networks with arbitrary bounded nonlinear activation functions. *IEEE Transactions on Neural Networks*, 9(1):224–229, Jan 1998.
- [GLGP13] Y. Gong, S. Lazebnik, A. Gordo, and F. Perronnin. Iterative quantization: A procrustean approach to learning binary codes for large-scale image retrieval. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(12):2916–2929, Dec 2013.
- [GLP⁺10] C. Sinan Güntürk, Mark Lammers, Alex Powell, Rayan Saab, and Özgür Yılmaz. Sigma delta quantization for compressed sensing. In *Information Sciences and Systems (CISS), 2010 44th Annual Conference on*, pages 1–6. IEEE, 2010.
- [GLP⁺13] C Sinan Güntürk, Mark Lammers, Alexander M Powell, Rayan Saab, and Ö Yılmaz. Sobolev duals for random frames and $\sigma\delta$ quantization of compressed sensing measurements. *Foundations of Computational mathematics*, 13(1):1–36, 2013.
- [Gün03] C Sinan Güntürk. One-bit sigma-delta quantization with exponential accuracy. *Communications on Pure and Applied Mathematics: A Journal Issued by the Courant Institute of Mathematical Sciences*, 56(11):1608–1630, 2003.
- [GVT95] Vivek K. Goyal, Martin Vetterli, and Nguyen T. Thao. Quantization of overcomplete expansions. In *Data Compression Conference, 1995. DCC'95. Proceedings*, pages 13–22. IEEE, 1995.
- [HBRN10] Jarvis Haupt, Waheed U Bajwa, Gil Raz, and Robert Nowak. Toeplitz compressed sensing matrices with applications to sparse channel estimation. *IEEE transactions on information theory*, 56(11):5862–5875, 2010.
- [HHL10] Justin P Haldar, Diego Hernando, and Zhi-Pei Liang. Compressed-sensing mri with random encoding. *IEEE transactions on Medical Imaging*, 30(4):893–903, 2010.
- [HLC⁺52] G.H. Hardy, J.E. Littlewood, Karreman Mathematics Research Collection, G. Pólya, D.E. Littlewood, and G. Pólya. *Inequalities*. Cambridge Mathematical Library. Cambridge University Press, 1952.
- [Hor91] Kurt Hornik. Approximation capabilities of multilayer feedforward networks. *Neural Networks*, 4(2):251 – 257, 1991.

- [HR17] Ishay Haviv and Oded Regev. The restricted isometry property of subsampled fourier matrices. In *Geometric Aspects of Functional Analysis*, pages 163–179. Springer, 2017.
- [HR18] Pablo A. Henríquez and Gonzalo Ruz. Twitter sentiment classification based on deep random vector functional link. pages 1–6, 07 2018.
- [HS18] Thang Huynh and Rayan Saab. Fast binary embeddings, and quantized compressed sensing with structured matrices. *preprint arXiv:1801.08639*, 2018.
- [Huy16] Thang Huynh. *Accurate quantization in redundant systems: From frames to compressive sampling and phase retrieval*. PhD thesis, New York University, 2016.
- [HZS06] Guang-Bin Huang, Qin-Yu Zhu, and Chee-Kheong Siew. Extreme learning machine: Theory and applications. *Neurocomputing*, 70(1):489 – 501, 2006. Neural Networks.
- [IKKSM18] Mark A. Iwen, Felix Krahmer, Sara Krause-Solberg, and Johannes Maly. On recovery guarantees for one-bit compressed sensing on manifolds. *preprint arXiv:1807.06490*, 2018.
- [ILNS19] Mark A Iwen, Eric Lybrand, Aaron A Nelson, and Rayan Saab. New algorithms and improved guarantees for one-bit compressed sensing on manifolds. In *Sampling Theory and Applications (SampTA), 2019 17th Annual International Conference on*, 2019.
- [IM13] Mark A. Iwen and Mauro Maggioni. Approximation of points on low-dimensional manifolds via random linear projections. *Information and Inference*, 2(1):1–31, 2013.
- [IP95] Boris Igel'nik and Yoh-Han Pao. Stochastic choice of basis functions in adaptive function approximation and the functional-link net. *IEEE Transactions on Neural Networks*, 6(6):1320–1329, 1995.
- [JCIY97] Jin-Yan Li, W. S. Chow, B. Igel'nik, and Yoh-Han Pao. Comments on "stochastic choice of basis functions in adaptive function approximation and the functional-link net" [with reply]. *IEEE Transactions on Neural Networks*, 8(2):452–454, March 1997.
- [JL84] William B Johnson and Joram Lindenstrauss. Extensions of lipschitz mappings into a hilbert space. *Contemporary mathematics*, 26(189-206):1, 1984.
- [JLBB13] Laurent Jacques, Jason N Laska, Petros T Boufounos, and Richard G Baraniuk. Robust 1-bit compressive sensing via binary stable embeddings of sparse vectors. *IEEE Transactions on Information Theory*, 59(4):2082–2102, 2013.
- [KST19] Rakesh Katuwal, Ponnuthurai Suganthan, and M Tanveer. Random vector functional link neural network based ensemble deep learning. 06 2019.

- [KSW12] Felix Krahmer, Rayan Saab, and Rachel Ward. Root-exponential accuracy for coarse quantization of finite frame expansions. *IEEE Transactions on Information Theory*, 58(2):1069–1079, 2012.
- [KSY14] Felix Krahmer, Rayan Saab, and Özgür Yılmaz. Sigma–delta quantization of sub-gaussian frame expansions and its application to compressed sensing. *Information and Inference: A Journal of the IMA*, 3(1):40–58, 2014.
- [KSZ18] Rakesh Katuwal, P.N. Suganthan, and Le Zhang. An ensemble of decision trees with random vector functional link networks for multi-class classification. *Applied Soft Computing*, 70:1146 – 1153, 2018.
- [KW11] Felix Krahmer and Rachel Ward. New and improved johnson–lindenstrauss embeddings via the restricted isometry property. *SIAM Journal on Mathematical Analysis*, 43(3):1269–1281, 2011.
- [LDP07] Michael Lustig, David Donoho, and John M Pauly. Sparse mri: The application of compressed sensing for rapid mr imaging. *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, 58(6):1182–1195, 2007.
- [Led01] Michel Ledoux. *The concentration of measure phenomenon*. Number 89. American Mathematical Soc., 2001.
- [LLPS93] Moshe Leshno, Vladimir Ya. Lin, Allan Pinkus, and Shimon Schocken. Multilayer feedforward networks with a nonpolynomial activation function can approximate any function. *Neural Networks*, 6(6):861 – 867, 1993.
- [LS17] Eric Lybrand and Rayan Saab. Quantization for low-rank matrix recovery. *Information and Inference*, 2017.
- [MAD⁺12] Mark Murphy, Marcus Alley, James Demmel, Kurt Keutzer, Shreyas Vasanawala, and Michael Lustig. Fast ℓ_1 -spirit compressed sensing parallel imaging mri: Scalable parallel implementation and clinically feasible runtime. *IEEE transactions on medical imaging*, 31(6):1250–1262, 2012.
- [Mas98] P Massart. About the constants in Talagrand’s deviation inequalities for empirical processes. Technical report, tech. rep., Laboratoire de statistiques, Université Paris Sud, 1998.
- [MPTJ08] Shahar Mendelson, Alain Pajor, and Nicole Tomczak-Jaegermann. Uniform uncertainty principle for bernoulli and subgaussian ensembles. *Constructive Approximation*, 28(3):277–289, 2008.
- [OR15] Samet Oymak and Ben Recht. Near-optimal bounds for binary embeddings of arbitrary sets. *arXiv preprint arXiv:1512.04433*, 2015.

- [OWB18] Matthew Olson, Abraham J. Wyner, and Richard Berk. Modern neural networks generalize on small data sets. In *Proceedings of the 32Nd International Conference on Neural Information Processing Systems, NIPS' 18*, pages 3623–3632, USA, 2018. Curran Associates Inc.
- [PP95] Yoh-Han Pao and Stephen M. Phillips. The functional link net and learning optimal control. *Neurocomputing*, 9(2):149 – 164, 1995. Control and Robotics, Part II.
- [PP00] Gwang Hoon Park and Yoh Han Pao. Unconstrained word-based approach for off-line script recognition using density-based random-vector functional-link net. *Neurocomputing*, 31(1):45 – 65, 2000.
- [PPS94] Yoh-Han Pao, Gwang-Hoon Park, and Dejan J. Sobajic. Learning and generalization characteristics of the random vector functional-link net. *Neurocomputing*, 6(2):163 – 180, 1994. Backpropagation, Part IV.
- [PST17] Shanthi Pavan, Richard Schreier, and Gabor C Temes. *Understanding delta-sigma data converters*. John Wiley & Sons, 2017.
- [PT92] Y. H. Pao and Y. Takefuji. Functional-link net computing: theory, system architecture, and functionalities. *Computer*, 25(5):76–79, May 1992.
- [PV12] Yaniv Plan and Roman Vershynin. Robust 1-bit compressed sensing and sparse logistic regression: A convex programming approach. *IEEE Transactions on Information Theory*, 59(1):482–494, 2012.
- [PV13] Yaniv Plan and Roman Vershynin. One-bit compressed sensing by linear programming. *Communications on Pure and Applied Mathematics*, 66(8):1275–1297, 2013.
- [PV14] Yaniv Plan and Roman Vershynin. Dimension reduction by random hyperplane tessellations. *Discrete & Computational Geometry*, 51(2):438–461, 2014.
- [Rau08] Holger Rauhut. Stability results for random sampling of sparse trigonometric polynomials. *IEEE Transactions on Information theory*, 54(12):5661–5670, 2008.
- [RL09] Maxim Raginsky and Svetlana Lazebnik. Locality-sensitive binary codes from shift-invariant kernels. In *Advances in neural information processing systems*, pages 1509–1517, 2009.
- [Rom09] Justin Romberg. Compressive sensing by random convolution. *SIAM Journal on Imaging Sciences*, 2(4):1098–1128, 2009.
- [RRT12] Holger Rauhut, Justin Romberg, and Joel A Tropp. Restricted isometries for partial random circulant matrices. *Applied and Computational Harmonic Analysis*, 32(2):242–254, 2012.

- [Rud91] W. Rudin. *Functional Analysis*. International series in pure and applied mathematics. McGraw-Hill, 1991.
- [RV08] Mark Rudelson and Roman Vershynin. On sparse reconstruction from fourier and gaussian measurements. *Communications on Pure and Applied Mathematics: A Journal Issued by the Courant Institute of Mathematical Sciences*, 61(8):1025–1045, 2008.
- [SCC18] Uri Shaham, Alexander Cloninger, and Ronald R. Coifman. Provable approximation properties for deep neural networks. *Applied and Computational Harmonic Analysis*, 44(3):537 – 557, 2018.
- [SH09] Ruslan Salakhutdinov and Geoffrey Hinton. Semantic hashing. *International Journal of Approximate Reasoning*, 50(7):969–978, 2009.
- [SKD92] W. F. Schmidt, M. A. Kraaijveld, and R. P. W. Duin. Feedforward neural networks with random weights. In *Proceedings., 11th IAPR International Conference on Pattern Recognition. Vol.II. Conference B: Pattern Recognition Methodology and Systems*, pages 1–4, Aug 1992.
- [SP70] E.M. Stein and Princeton University Press. *Singular Integrals and Differentiability Properties of Functions*. Monographs in harmonic analysis. Princeton University Press, 1970.
- [Sug18] Ponnuthurai Nagaratnam Suganthan. Letter: On non-iterative learning algorithms with closed-form solution. *Appl. Soft Comput.*, 70:1078–1082, 2018.
- [SW71] E.M. Stein and G. Weiss. *Introduction to Fourier Analysis on Euclidean Spaces*. Mathematical Series. Princeton University Press, 1971.
- [SWY18a] Rayan Saab, Rongrong Wang, and Özgür Yılmaz. From compressed sensing to compressed bit-streams: practical encoders, tractable decoders. *IEEE Transactions on Information Theory*, 64(9):6098–6114, 2018.
- [SWY18b] Rayan Saab, Rongrong Wang, and Özgür Yılmaz. Quantization of compressive samples with stable and robust recovery. *Applied and Computational Harmonic Analysis*, 44(1):123–143, 2018.
- [Tal96] Michel Talagrand. New concentration inequalities in product spaces. *Inventiones mathematicae*, 126(3):505–563, 1996.
- [Tu10] L.W. Tu. *An Introduction to Manifolds*. Universitext. Springer New York, 2010.
- [TWY18] Ling Tang, Yao Wu, and Lean Yu. A non-iterative decomposition-ensemble learning paradigm using rvfl network for crude oil price forecasting. *Applied Soft Computing*, 70:1097 – 1108, 2018.

- [VAH⁺10] Shreyas S Vasanawala, Marcus T Alley, Brian A Hargreaves, Richard A Barth, John M Pauly, and Michael Lustig. Improved pediatric mr imaging with compressed sensing. *Radiology*, 256(2):607–616, 2010.
- [VPM18] Najdan Vuković, Milica Petrović, and Zoran Miljković. A comprehensive experimental evaluation of orthogonal polynomial expanded random vector functional link neural networks for regression. *Applied Soft Computing*, 70:1083 – 1096, 2018.
- [WTF09] Yair Weiss, Antonio Torralba, and Rob Fergus. Spectral hashing. In *Advances in neural information processing systems*, pages 1753–1760, 2009.
- [YCP15] Xinyang Yi, Constantine Caramanis, and Eric Price. Binary embedding: Fundamental limits and fast algorithm. In *International Conference on Machine Learning*, pages 2162–2170, 2015.
- [ZS17a] L. Zhang and P. N. Suganthan. Benchmarking ensemble classifiers with novel co-trained kernel ridge regression and random vector functional link ensembles [research frontier]. *IEEE Computational Intelligence Magazine*, 12(4):61–72, Nov 2017.
- [ZS17b] L. Zhang and P. N. Suganthan. Visual tracking with convolutional random vector functional link network. *IEEE Transactions on Cybernetics*, 47(10):3243–3253, Oct 2017.
- [ZWC⁺19] Yongshan Zhang, Jia Wu, Zhihua Cai, Bo Du, and Philip S. Yu. An unsupervised parameter learning model for rvfl neural network. *Neural Networks*, 112:85 – 97, 2019.