

Duality for Neural Networks through Reproducing Kernel Banach Spaces

Len Spek^{1,*}, Tjeerd Jan Heeringa¹, Felix Schwenninger², and Christoph Brune¹

¹*Mathematics of Imaging & AI, University of Twente, Enschede, The Netherlands*

²*Mathematics of Systems Theory, University of Twente, Enschede, The Netherlands*

February 2023

Abstract

Reproducing Kernel Hilbert spaces (RKHS) have been a very successful tool in various areas of machine learning. Recently, Barron spaces have been used to prove bounds on the generalisation error for neural networks. Unfortunately, Barron spaces cannot be understood in terms of RKHS due to the strong nonlinear coupling of the weights. This can be solved by using the more general Reproducing Kernel Banach spaces (RKBS). We show that these Barron spaces belong to a class of integral RKBS. This class can also be understood as an infinite union of RKHS spaces. Furthermore, we show that the dual space of such RKBSs, is again an RKBS where the roles of the data and parameters are interchanged, forming an adjoint pair of RKBSs including a reproducing kernel. This allows us to construct the saddle point problem for neural networks, which can be used in the whole field of primal-dual optimisation.

1 Introduction

Neural networks are often considered to be black boxes from a mathematical perspective. However, a lot of recent progress has been made by E et al. [2019b] by proving generalisation bounds and approximation errors for two layer neural networks. The key tool used in these proofs is the Barron space, a space of functions which can be represented by an infinite width neural networks with bounded weights.

Reproducing Kernel Hilbert spaces (RKHS) have been a very useful tool in machine learning methods. The kernel property, together with the easy representation of the dual space due to the inner product structure, give easy access to various results from representation and optimisation theory. For example, this is the key to prove the representer theorem for random feature models [Bach, 2017b]. However, to consider full two layer neural networks E et al. [2021] has shown that one needs an infinite union of RKHS.

In this paper, we show that the Barron spaces satisfy the Reproducing Kernel Banach space (RKBS) property. These spaces were introduced by Zhang et al. [2009] to study the representation properties in machine learning, and were brought into the context of neural networks by Bartolucci et al. [2023] to prove a representer theorem. In merging these two approaches, Barron spaces and RKBS, the approximation error bounds of the Barron space and the representation results of RKBS can be combined. As an RKBS is not by definition a Hilbert space, a description of the dual space is necessary to study the eponymous reproducing kernel.

The class of RKBS used by Bartolucci et al. [2023] for neural networks have an integral representation with measure over the parameters. We show that its dual space can be identified with an RKBS which has an integral representation with measures over the data. These two spaces form an adjoint RKBS pair with a well-defined reproducing kernel. This enables the fundamental basis for learned primal-dual methods for empirical risk minimisation in the context of architecture search and targeted optimisation. Here the primal variables are the neural networks parameters and the dual variables are observations of the data. The dual view opens a door to the field of primal-dual algorithms.

As mentioned above, the Barron space can be written as an infinite union of RKHS. We show that this argument generalises to RKBS, where an RKBS with an integral representation over measures can be written as an infinite union of certain spaces with an L^p -norm, which also satisfy the RKBS property.

Surprisingly, this result is independent of the choice of p . For $p = 2$, these spaces satisfy the RKHS property, and we recover the result of E et al. [2021].

1.1 Research Context

The Reproducing Kernel Hilbert spaces have a long history, [Aronszajn, 1950], and have been developed into an indispensable concept in several fields of mathematics, such as operator theoretic approaches in complex analysis. In the last decades, they have also become very relevant in machine learning [Bach, 2017b]. Advantages of these spaces are that inputs can be explained geometrically and that many powerful results from functional analysis can be used [Scholkopf and Smola, 2018]. This has led to many effective algorithms for different learning tasks, such as support vector machines [Wahba, 1999, Evgeniou et al., 2000].

However, the Hilbert structure does not appear to be suitable for certain learning tasks. These spaces offer effectively only a single geometry, as all Hilbert spaces of the same dimension are isometrically isomorphic, cf Rudin [1991]. On the contrary, general Banach spaces offer much more flexibility, as the norm does not need to follow the parallelogram law. This is relevant when considering l^1 -regularisation, which is useful for creating sparsity [Tibshirani, 1996, Micchelli and Pontil, 2007, Song et al., 2013]. This naturally leads to the question if an equivalent theory for Reproducing Kernel Banach spaces can be developed. Such a generalisation was proposed by Canu et al. [2003], where the Reproducing Kernel Banach space is based on bounded point evaluation of functions. I suggest to rather say "to consider semi inner products in Banach spaces, which to some extent allow mimicking the Hilbert space case [Lumer, 1960, Giles, 1967] and are applied to machine learning [Zhang and Zhang, 2017, Zhang et al., 2009, García and Portal, 2013]. A good overview on the general RKBS theory can be found in Lin et al. [2022].

Neural networks are often studied in the infinite width case, as the universal approximation theorem guarantees that any continuous function can be approximated in the supremum norm, given some mild assumptions on the activation function [Cybenko, 1989, Hornik, 1991, Leshno et al., 1993]. Barron [1993] managed to prove an L^2 convergence rate based on Fourier techniques, which do not suffer from the curse of dimensionality. In honour of his work, E et al. [2020] defined the Barron spaces, which consists of functions which can be approximated by neural networks without the weights blowing up. With the corresponding norms, E et al. [2019b] were able to prove bounds of different kinds of generalisation errors.

On the other hand, Bartolucci et al. [2023] used RKBS to understand neural networks and prove a general representer theorem for these spaces. The norm of the RKBS corresponding to ReLU activation functions can be represented in terms of an inverse Radon transform of a bounded real measure [Parhi and Nowak, 2021].

Duality is a key concept in the study of functional analysis. The duality pairing for Banach spaces is the general analogue of the inner product of Hilbert spaces. This enables primal-dual algorithms, which have been very successful in the field of convex analysis. For example, total variation denoising methods have been significantly improved using a proper description of dual space of functions of bounded variation [Chambolle, 2004, Chan et al., 1999, Osher et al., 2005, Brune et al., 2011]. For the gradient descent algorithm, advances are made to incorporate duality, such dual coordinated ascent methods [Shalev-Shwartz and Zhang, 2013, Raj and Bach, 2021].

1.2 Overview

In Section 2, we first give an overview of the different, but equivalent definitions of RKHS and how they generalise to RKBS. Next, we introduce the two spaces corresponding to infinitely wide neural networks: The class of integral RKBS from Bartolucci et al. [2023] and the Barron spaces of E et al. [2019b]. We then show that the Barron space indeed satisfies the RKBS property.

In Section 3, we show that an RKBS with an integral representation over measures can be written as an infinite union of certain L^p -like spaces, which also satisfy the RKBS property.

In Section 4, we prove that an RKBS which has an integral representation with measure over the parameters, has a dual space which is an RKBS with an integral representation with measures over the data. Using this, we construct the saddle point problem for an unconstrained optimisation problem.

1.3 Notation

In the following sections, we employ the following notation conventions and assumptions: $X \subseteq \mathbb{R}^d$ denotes the set from which the data points $x \in X$ are sampled and $\Omega \subseteq \mathbb{R}^D$ denotes the set of weights $w \in \Omega$ over which we want to optimise our model. We assume that X and Ω are closed and endowed with the natural subspace topology.

We consider a bounded, continuous kernel $\varphi \in C_0(X \times \Omega)$ which vanishes at infinity. We introduce the notation $\varphi_w \in C_0(X)$ for the function $x \mapsto \varphi(x, w)$ given some $w \in \Omega$. Similarly, we write $\varphi_x \in C_0(\Omega)$ for the function $w \mapsto \varphi(x, w)$ given some $x \in X$.

We consider real-valued Radon measures, i.e. the signed, finite, regular Borel measures, defined on either X or Ω . The natural norm for a Radon measure μ is the total variation norm.

$$\|\mu\|_{\mathcal{M}(\Omega)} := |\mu|(\Omega)$$

where $|\mu|$ is defined through the Jordan decomposition: There exists two unique positive measures μ^+ and μ^- such that $\mu = \mu^+ - \mu^-$. Then $|\mu| := \mu^+ + \mu^-$. The Banach space $\mathcal{M}(\Omega)$ is defined as the set of all Radon measures on Ω with finite total variation norm. A measure $\mu \in \mathcal{M}(\Omega)$ represents the distributions of weights. Similarly, $\rho \in \mathcal{M}(X)$ represents a data distribution. The subspace of probability measures is denoted by $\mathcal{P}(\Omega), \mathcal{P}(X)$ respectively. Point measures will always be referred to by using the Greek letter δ .

The space of continuous functions vanishing at infinity, $C_0(\Omega)$, is a Banach space when endowed with the supremum norm. The Riesz representation theorem implies that $\mathcal{M}(X), \mathcal{M}(\Omega)$ can be identified as the dual spaces of $C_0(X), C_0(\Omega)$ respectively, and the duality pairings imply that the following integrals are well-defined and finite

$$\begin{aligned} \langle \rho, \varphi_w \rangle &= \int_X \varphi_w(x) d\rho(x) \\ \langle \mu, \varphi_x \rangle &= \int_\Omega \varphi_x(w) d\mu(w) \end{aligned}$$

for all $\rho \in \mathcal{M}(X), \mu \in \mathcal{M}(\Omega)$.

Given some measure $\mu \in \mathcal{M}(\Omega)$ and integrable function $h \in L^1(\mu)$, we use the notation $d\nu := h d\mu$ to define a measure $\nu \in \mathcal{M}(\Omega)$ which is absolutely continuous with respect to μ and for every Borel set $A \subseteq \Omega$

$$\nu(A) = \int_A h(w) d\mu(w)$$

Finally, when there exists an isometric isomorphism between two Banach spaces $\mathcal{B}_1, \mathcal{B}_2$, we denote this by $\mathcal{B}_1 \cong \mathcal{B}_2$.

2 Reproducing Kernel Banach Spaces

In this section, we will introduce the Reproducing Kernel Banach spaces (RKBS) similarly as Bartolucci et al. [2023] by first comparing three equivalent definitions of Reproducing Kernel Hilbert spaces (RKHS). We will then introduce a class of integral RKBS, which will be the main object of study in this paper. Finally, we show that the Barron space with ReLU functions of E et al. [2019b] satisfies the integral RKBS properties.

2.1 Definition of RKHS

Reproducing Kernel Hilbert spaces can be equivalently characterised in three ways: boundedness of point evaluation, existence of a reproducing kernel and as a quotient space of a feature space.

First, the classical definition where point evaluation is a bounded functional.

Definition 1. *A Hilbert space \mathcal{H} of functions on X is a reproducing kernel Hilbert space (RKHS) if, for all $x \in X$, there exists a constant $C_x > 0$ such that for all $f \in \mathcal{H}$*

$$|f(x)| \leq C_x \|f\|_{\mathcal{H}} \tag{1}$$

As point evaluation is a linear functional, the Riesz representation theorem guarantees the existence of the eponymous reproducing kernel

Theorem 2. *A Hilbert space \mathcal{H} of functions on X satisfies the RKHS property if and only if there exists a function $K : X \times X \mapsto \mathbb{R}$ such that for all $x \in X$*

$$\begin{aligned} K(x, \cdot) &\in \mathcal{H} \\ f(x) &= \langle f, K(x, \cdot) \rangle \quad \forall f \in \mathcal{H} \end{aligned} \tag{2}$$

This is easy to check that K is symmetric and positive definite. Each such kernel defines a unique RKHS [Aronszajn, 1950]. This property of an RKHS has proven very useful in empirical risk minimisation problems, where the minimiser can be written in terms of the reproducing kernel.

Finally, we introduce a definition which is popular in machine learning. An RKHS can be written as a quotient space of a feature space Ψ

Theorem 3. *A Hilbert space \mathcal{H} of functions on X satisfies the RKHS property if and only if there exists a Hilbert space Ψ and a map $\psi : X \mapsto \Psi$ such that*

$$\begin{aligned} \mathcal{H} &= \Psi / \mathcal{N}(A) \\ \|f\|_{\mathcal{H}} &= \inf_{f=A\nu} \|\nu\|_{\Psi} \end{aligned} \tag{3}$$

where A maps features in Ψ to functions on X and is defined as

$$(A\nu)(x) = \langle \psi(x), \nu \rangle \tag{4}$$

for all $x \in X$ and $\nu \in \Psi$.

The reproducing kernel K is related to the map ψ via the adjoint of the mapping $A : \Psi \mapsto \mathcal{H}$

$$\psi(x) = A^*K(x, \cdot) \tag{5}$$

2.2 Introduction to RKBS

Definition 1 immediately translates to a Banach setting.

Definition 4. *A Banach space \mathcal{B} of functions on X is a reproducing kernel Hilbert space (RHKS) if, for all $x \in X$, there exists a constant $C_x > 0$ such that for all $f \in \mathcal{B}$*

$$|f(x)| \leq C_x \|f\|_{\mathcal{B}} \tag{6}$$

A trivial example of an RKBS, besides any RKHS space, is the Banach space of bounded continuous functions equipped with the supremum norm.

Remark 5. *Any Banach space \mathcal{B} can be identified with the subspace $j(\mathcal{B})$ of its bi-dual \mathcal{B}^{**} via the canonical embedding j . Since $j(\mathcal{B})$ is a space of functions on \mathcal{B}^* and by definition satisfies (6), it is an RKBS on \mathcal{B}^* . For Hilbert spaces, this means that any Hilbert space \mathcal{H} can be identified as an RKHS of functions over \mathcal{H} .*

The canonical characterisation is, however, not very useful in practise. As in most applications, X is assumed to be relatively ‘small’ compared to \mathcal{B} . One sufficient condition for this intuitive constraint is that X is chosen such that \mathcal{B} contains non-linear functions of X .

Theorem 3 also translates to the Banach setting, by replacing the inner product with a duality pairing between Ψ and its dual space.

Theorem 6. [Bartolucci et al., 2023, Proposition 3.3] *A Banach space \mathcal{B} of functions on X satisfies the RKBS property if and only if there exists a Banach space Ψ and a map $\psi : X \mapsto \Psi^*$ such that*

$$\begin{aligned} \mathcal{B} &= \Psi / \mathcal{N}(A) \\ \|f\|_{\mathcal{B}} &= \inf_{f=A\nu} \|\nu\|_{\Psi} \end{aligned} \tag{7}$$

where the linear transformation A maps elements of the Banach space Ψ to functions on X and is defined as

$$(A\nu)(x) := \langle \psi(x), \nu \rangle \tag{8}$$

for all $x \in X$ and $\nu \in \Psi$.

Note that the null-space $\mathcal{N}(A)$ is closed, so the quotient space $\Psi/\mathcal{N}(A)$ is complete and hence a proper Banach space.

Unfortunately, Definition 2 does not generalise easily to the Banach setting due to the lack of an analogue to the Riesz-Frechet representation theorem. However, based on Definition 4, the point evaluation operator δ_x is an element of the dual space of \mathcal{B} , which allows for a definition of an adjoint RKBS with a reproducing kernel.

Definition 7. [Lin et al., 2022, Definition 2.2] Let \mathcal{B} be an RKBS on the set X . If there exists a Banach space $\mathcal{B}^\#$ of functions on a set Ω , such that it can be isomorphically embedded into \mathcal{B}^* and if there exists a function $K : X \times \Omega \mapsto \mathbb{R}$, such that $K(x, \cdot) \in \mathcal{B}^\#$ for all $x \in X$ and

$$f(x) = \langle K(x, \cdot), f \rangle \quad (9)$$

for all $x \in X$ and $f \in \mathcal{B}$, then we call K a reproducing kernel for \mathcal{B} .

If, in addition, $\mathcal{B}^\#$ is also an RKBS on Ω and it holds that $K(\cdot, w) \in \mathcal{B}$ for all $w \in \Omega$ and

$$g(w) = \langle g, K(\cdot, w) \rangle \quad (10)$$

for all $w \in \Omega$ and $g \in \mathcal{B}^\#$, then we call $\mathcal{B}^\#$ an adjoint RKBS of \mathcal{B} and call \mathcal{B} and $\mathcal{B}^\#$ an adjoint pair of RKBS. In this case $K^*(w, x) := K(x, w)$, for $x \in X$ and $w \in \Omega$, is a reproducing kernel of $\mathcal{B}^\#$.

Note that compared to the RKHS case, the reproducing kernel K is no longer symmetric. Also, it follows immediately that any RKHS is adjoint to itself. Furthermore, the adjoint RKBS needs not be unique, as it depends on the chosen set Ω and the representation of the space $\mathcal{B}^\#$.

For these adjoint pairs of RKBS, we recover again the identity

$$\psi(x) = A^* K(x, \cdot) \quad (11)$$

Lin et al. [2022, Theorem 2.4] prove that any separable RKBS admits a reproducing kernel. However, the separability conditions are not necessary, as there exists a canonical way to define the reproducing kernel for any RKBS by using that the elements of the dual space \mathcal{B}^* can be understood as functions of \mathcal{B} .

Theorem 8. Let \mathcal{B} be an RKBS on the set X and let $\delta_x \in \mathcal{B}^*$ be the point evaluation functional at $x \in X$. Then

$$\mathcal{B}^\# := \overline{\text{span}\{\delta_x | x \in X\}} \quad (12)$$

and \mathcal{B} form an adjoint pair of RKBS with a kernel $K : X \times \mathcal{B} \mapsto \mathbb{R}$

$$K(x, f) := f(x) \quad (13)$$

for all $x \in X$, $f \in \mathcal{B}$.

Proof. By definition, $\mathcal{B}^\#$ is a Banach space of functions of \mathcal{B} . Furthermore, $K(x, \cdot) = \delta_x \in \mathcal{B}^\#$ for all $x \in X$, and hence for all $f \in \mathcal{B}$

$$\langle f, K(x, \cdot) \rangle = \langle f, \delta_x \rangle = f(x)$$

Moreover, $K(\cdot, f) = f \in \mathcal{B}$ and for hence for all $g \in \mathcal{B}^\#$

$$\langle g, K(\cdot, f) \rangle = \langle g, f \rangle = g(f)$$

by definition of the duality pairing. □

Again, in most applications, we want Ω to also be 'small' or that $\mathcal{B}^\#$ also contains non-linear functions of Ω . In the rest of this paper, X and Ω are always 'small' with respect to the different RKBSs.

2.3 A Class of Integral RKBS

Next, we will describe a class of RKBS where we take the Radon measures as the feature space. As we will see in the next section, this is relevant to discuss spaces of functions defined by neural networks. We will construct this RKBS using the feature map of Definition 6.

We fix a function $\varphi \in C_0(X \times \Omega)$. For the feature space Ψ , we choose the space of Radon measures $\mathcal{M}(\Omega)$ equipped with the total variation norm. Then we define the feature map $\psi(x) := \varphi_x$, where we implicitly assume that $\varphi_x \in C_0(\Omega)$ is canonically embedded into the bi-dual of $C_0(\Omega)$.

$$(A\mu)(x) := \langle \psi(x), \mu \rangle = \int_{\Omega} \varphi_x(w) d\mu(w) = \int_{\Omega} \varphi(x, w) d\mu(w) \quad (14)$$

for all $\mu \in \mathcal{M}(\Omega)$. The RKBS $\mathcal{F}(X, \Omega)$ is then defined as

$$\mathcal{F}(X, \Omega) := \{f : X \mapsto \mathbb{R} \mid \exists \mu \in \mathcal{M}(\Omega) \text{ s.t. } f = A\mu\} \quad (15)$$

Hence, $\mathcal{F}(X, \Omega)$ can be identified as the quotient space $\mathcal{M}(\Omega)/\mathcal{N}(A)$ and its norm is given by

$$\|f\|_{\mathcal{F}(X, \Omega)} := \inf_{f=A\mu} \|\mu\|_{\mathcal{M}(\Omega)} \quad (16)$$

In plain terms, $\mathcal{F}(X, \Omega)$ contains functions f which can be represented by an integral over φ with any measure μ and the norm of f depends on how large this measure, in terms of its total variation, needs to be. Note that $\varphi \in C_0(X \times \Omega)$ in combination with dominated convergence imply that $\mathcal{F}(X, \Omega)$ is a subspace of $C_0(X)$, albeit with a different norm.

This integral RKBS $\mathcal{F}(X, \Omega)$ is equivalent to the space \mathcal{F}_1 with variation norm γ_1 used by Bach [2017a] and Chizat et al. [2020]. The subscript 1 denotes that the norm behaves similarly to a 1-norm. Similarly, this space is equivalent to the integral RKBS defined by Bartolucci et al. [2023], where $\varphi(x, w)$ is replaced by $\rho(x, \theta)\beta(\theta)$. The smoothing function β is used to ensure that the integral (14) converges for all μ . However, we opted to combine ρ and β into φ , as it produces cleaner notation in later sections. For neural networks with ReLU activation functions, Parhi and Nowak [2021] defined a norm in terms of the inverse Radon transform of the network, which is equivalent to the $\mathcal{F}(X, \Omega)$ -norm [Bartolucci et al., 2023]. We will add to this list of equivalences by proving that the Barron spaces defined by E et al. [2019b] are also an instance of an integral class of RKBS.

2.4 Barron Spaces

The Barron norm for neural networks is defined by E et al. [2019b] to quantify the size of the weights in a neural network of infinite width. In this section we will repeat this definition and show that it satisfies the RKBS property, Definition 6, by showing it is equivalent to the RKBS $\mathcal{F}(X, \Omega)$ for a certain φ .

A shallow neural network or perceptron f maps data vectors $x \in X$, where we assume X is compact, to a scalar output. We consider a continuous activation function $\sigma : \mathbb{R} \mapsto \mathbb{R}$ and some weights $(a_j, v_j, b_j) \in \mathbb{R} \times \Omega \subseteq \mathbb{R}^{d+1}$, where $j \in \{1, \dots, m\}$, and functions of the form

$$f(x) = \frac{1}{m} \sum_{j=1}^m a_j \sigma(v_j^T x + b_j) \quad (17)$$

Formally transitioning to the continuous limit, we can consider functions $f : X \rightarrow \mathbb{R}$, for which there exists a probability measure $\pi \in \mathcal{P}(\mathbb{R} \times \Omega)$ such that

$$f(x) := (A\pi)(x) := \int_{\mathbb{R} \times \Omega} a \sigma(v^T x + b) d\pi(a, v, b) \quad (18)$$

The vector space of such function is called Barron space. The Barron norm of f is defined by taking the infimum of all probability measures π , which produce f , more precisely,

$$\|f\|_{\mathcal{B}_\sigma} := \inf_{f=A\pi} \int_{\mathbb{R} \times \Omega} |a|(1 + \|v\|_1 + |b|) d\pi(a, v, b) \quad (19)$$

For the Barron space to be well-defined, we either require that Ω is compact or that σ grows at most linearly, i.e. $\sigma(y) \in O(y)$. In this case there exists an $M > 0$ such that $|\sigma(y)| \leq My$, for all $y \in \mathbb{R}$ and

$$\begin{aligned} |f(x)| &\leq \int_{\mathbb{R} \times \Omega} |a| |\sigma(v^T x + b)| d\pi(a, v, b) \\ &\leq M \int_{\mathbb{R} \times \Omega} |a| (|v^T x + b|) d\pi(a, v, b) \\ &\leq M \int_{\mathbb{R} \times \Omega} |a| (\|x\|_\infty \|v\|_1 + |b|) d\pi(a, v, b) \end{aligned}$$

for all $x \in X$ and $\pi \in \mathcal{P}(\Omega)$ such that $f = A\pi$. Taking the infimum we get that $|f(x)| \leq M\|x\|_\infty \|f\|_{\mathcal{B}_\sigma}$ for all $f \in \mathcal{B}_\sigma$ and also immediately that the Barron space is an RKBS.

When σ is 1-homogeneous, like for the popular ReLU activation function, the definition of the Barron norm is slightly different

$$\|f\|_{\mathcal{B}_\sigma} := \inf_{f=A\pi} \int_{\mathbb{R} \times \Omega} |a|(\|v\|_1 + |b|) d\pi(a, v, b) \quad (20)$$

There is a related definition where the above norm is replaced by an infimum over the $L^p(\pi)$ -norm of $|a|(1 + \|v\|_1 + |b|)$ or $|a|(\|v\|_1 + |b|)$ for 1-homogeneous σ . In this latter case E et al. [2019a, Proposition 1] shows that the Barron norms for all $1 \leq p \leq \infty$ are equal.

We can show that the Barron space is an instance of the RKBS $\mathcal{F}(X, \Omega)$. First, we treat the case where σ is not 1-homogeneous.

Theorem 9. *Let $\sigma : \mathbb{R} \mapsto \mathbb{R}$ be continuous and let either Ω be compact or $\sigma(y) \in O(y)$*

Let $\mathcal{F}(X, \Omega)$ be the integral RKBS defined by $\varphi(x, w) = (1 + \|v\|_1 + |b|)^{-1} \sigma(v^T x + b)$ where $w = (v, b)$. Then $\mathcal{B}_\sigma \cong \mathcal{F}(X, \Omega)$.

Proof. Let σ be continuous. It is sufficient to show that for every $\pi \in \mathcal{P}(\mathbb{R} \times \Omega)$ there exists a Radon measure μ such that $A\pi = A\mu$ and

$$\|\mu\|_{\mathcal{M}(\Omega)} \leq \int_{\mathbb{R} \times \Omega} |a|(1 + \|v\|_1 + |b|) d\pi(a, v, b)$$

and conversely that for every Radon measure μ there exists a $\pi \in \mathcal{P}(\mathbb{R} \times \Omega)$ such that $A\mu = A\pi$ and

$$\int_{\mathbb{R} \times \Omega} |a|(1 + \|v\|_1 + |b|) d\pi(a, v, b) \leq \|\mu\|_{\mathcal{M}(\Omega)}$$

Taking the infima then gives the equality of norms. In the rest of this proof, we use the convention that $w = (v, b)$.

Define $\varphi(x, w) := (1 + \|v\|_1 + |b|)^{-1} \sigma(v^T x + b)$. By the assumptions above $\varphi \in C_0(X \times \Omega)$.

Let $\pi \in \mathcal{P}(\mathbb{R} \times \Omega)$, then define the measure $\mu \in \mathcal{M}(\Omega)$ such that for every Borel set $B \subseteq \Omega$

$$\mu(B) = \int_{\mathbb{R} \times B} a(1 + \|v\|_1 + |b|) d\pi(a, v, b)$$

Hence, we get that

$$\begin{aligned} (A\mu)(x) &= \int_{\Omega} \varphi(x, w) d\mu(w) \\ &= \int_{\mathbb{R} \times \Omega} a(1 + \|v\|_1 + |b|)(1 + \|v\|_1 + |b|)^{-1} \sigma(v^T x + b) d\pi(a, v, b) = (A\pi)(x) \end{aligned}$$

for all $x \in X$. Furthermore,

$$|\mu|(\Omega) \leq \int_{\mathbb{R} \times \Omega} |a|(1 + \|v\|_1 + |b|) d\pi(a, v, b)$$

Let $\mu \in \mathcal{M}(\Omega)$, with its Jordan decomposition $\mu = \mu^+ - \mu^-$. We define the measure π as

$$\begin{aligned} d\pi(a, v, b) &= \frac{1}{2(1 + \|v\|_1 + |b|)|\mu^+(\Omega)} d\mu^+(w) d\delta_{a=2|\mu^+(\Omega)}(a) \\ &\quad + \frac{1}{2(1 + \|v\|_1 + |b|)|\mu^-(\Omega)} d\mu^-(w) d\delta_{a=2|\mu^-(\Omega)}(a) \end{aligned}$$

Clearly, $\pi \in \mathcal{P}(\mathbb{R} \times \Omega)$. Furthermore,

$$\begin{aligned}
(A\pi)(x) &= \int_{\mathbb{R} \times \Omega} a\sigma(v^T x + b) d\pi(a, v, b) \\
&= \frac{1}{2|\mu^+(\Omega)|} \int_{\Omega} \int_{\mathbb{R}} a\varphi(x, w) d\delta_{a=2|\mu^+(\Omega)}(a) d\mu^+(w) \\
&\quad + \frac{1}{2|\mu^-(\Omega)|} \int_{\Omega} \int_{\mathbb{R}} a\varphi(x, w) d\delta_{a=2|\mu^-(\Omega)}(a) d\mu^-(w) \\
&= \int_{\Omega} \varphi(x, w) d\mu^+(w) + \int_{\Omega} \varphi(x, w) d\mu^-(w) \\
&= \int_{\Omega} \varphi(x, w) d\mu(w) = (A\mu)(x)
\end{aligned}$$

for all $x \in X$ and

$$\begin{aligned}
\int_{\mathbb{R} \times \Omega} |a|(1 + \|v\|_1 + |b|) d\pi(a, v, b) &= \frac{1}{2|\mu^+(\Omega)|} \int_{\Omega} \int_{\mathbb{R}} |a| d\delta_{a=2|\mu^+(\Omega)}(a) d\mu^+(w) \\
&\quad + \frac{1}{2|\mu^-(\Omega)|} \int_{\Omega} \int_{\mathbb{R}} |a| d\delta_{a=2|\mu^-(\Omega)}(a) d\mu^-(w) \\
&= \mu^+(\Omega) + \mu^-(\Omega) = |\mu|(\Omega)
\end{aligned}$$

Which concludes the proof. \square

Next, we show that the analogous statement holds when σ is 1-homogeneous. The proof is inspired by [E and Wojtowytsch, 2020b, Lemma A.5] and [E and Wojtowytsch, 2021]

Theorem 10. *Let σ be continuous and 1-homogeneous. and $\varphi(x, w) = \sigma(v^T x + b)$ with $w = (v, b)$. Then $\mathcal{B}_\sigma \cong \mathcal{F}(X, \Omega)$*

We employ the same technique as the proof of the Theorem 9. Let σ be continuous and 1-homogeneous, which means that for any $\lambda > 0$ and $y \in \mathbb{R}$ it holds that $\sigma(\lambda y) = \lambda\sigma(y)$. The key idea of the proof is reformulating the Barron space over Ω' where $\|v\| + |b| = 1$.

$$\Omega' := \left\{ \frac{w}{\|w\|_1} \mid w \in \Omega, w \neq 0 \right\}$$

Then, as Ω' is compact, the proof follows from Theorem 9.

Let $\pi \in \mathcal{P}(\mathbb{R} \times \Omega)$. For any Borel set $B' \subseteq \mathbb{R} \times \Omega'$ we define the Borel set $B \subseteq \mathbb{R} \times \Omega$

$$B := \left\{ (a, w) \in \mathbb{R} \times \Omega \mid \frac{w}{\|w\|_1} \in B' \right\}$$

We define $\pi' \in \mathcal{P}(\mathbb{R} \times \Omega')$ such that

$$\pi(B') = \int_B (\|v\|_1 + |b|) d\pi(a, v, b)$$

Hence, we get by the change of variables formula that for every function $g : \mathbb{R} \times \Omega \mapsto \mathbb{R}$

$$\int_B g\left(a, \frac{w}{\|w\|_1}\right) (\|v\|_1 + |b|) d\pi(a, v, b) = \int_{B'} g'(a, w) d\pi'(a, v, b)$$

where $g'(a, w) = g\left(a, \frac{w}{\|w\|_1}\right)$ for all $(a, w) \in \mathbb{R} \times \Omega$. Then we get that

$$\begin{aligned}
(A\pi)(x) &= \int_{\mathbb{R} \times \Omega} a\sigma(v^T x + b) d\pi(a, v, b) \\
&= \int_{\mathbb{R} \times \Omega} a(\|v\|_1 + |b|) \sigma\left(\frac{v^T x + b}{\|v\|_1 + |b|}\right) d\pi(a, v, b) \\
&= \int_{\mathbb{R} \times \Omega'} a\sigma(v^T x + b) d\pi'(a, v, b) \\
&= (A\pi')(x)
\end{aligned}$$

due to the fact that σ is 1-homogeneous. Furthermore,

$$\int_{\mathbb{R} \times \Omega} |a|(\|v\|_1 + |b|)d\pi(a, v, b) = \int_{\mathbb{R} \times \Omega'} |a|d\pi'(a, v, b) = \int_{\mathbb{R} \times \Omega'} |a|(\|v\|_1 + |b|)d\pi'(a, v, b)$$

For the converse, unfortunately we can not say that any $\pi' \in \mathcal{P}(\mathbb{R} \times \Omega')$ is also a probability measure on $\mathbb{R} \times \Omega$, as Ω' is not necessarily a subset of Ω . Hence, we need to define a function $g : \Omega' \mapsto \mathbb{R}$ such that $g(w)w \in \Omega$, i.e. a scale factor that scales Ω' into Ω . Then, given a $\pi' \in \mathcal{P}(\mathbb{R} \times \Omega')$, it is easy to see that π pushed forward by g is the corresponding measure on $\mathbb{R} \times \Omega$. \square

3 Integral RKBS as a Union of p-Norm RKBS

In this section, we extend Proposition 3 of E et al. [2021] that the Barron space can be written as a union of RKHS. We show that the integral RKBS $\mathcal{F}(X, \Omega)$ can be written as a union of certain RKBS which are quotient spaces of the L^p -spaces. Remarkably, this is independent of the choice p . The proof reveals an interesting fact about these RKBS: The curvature of the functions is not relevant for the norm, only the average absolute value of the function.

3.1 Definitions of p-Norm RKBS

We first need to properly define the p -norm RKBS and the concept of a union of Banach spaces.

Given a $\pi \in \mathcal{P}(\Omega)$, a $1 \leq p < \infty$ and a map $\varphi : X \times \Omega \mapsto \mathbb{R}$ such that $\varphi_x \in L^q(\pi)$ for all $x \in X$ and where q is such that $\frac{1}{p} + \frac{1}{q} = 1$.

We define the p -norm RKBS $\mathcal{L}^p(\pi)$ using again the feature map of Definition 6. We choose the feature space to be $\Psi = L^p(\pi)$ with its dual space $\Psi^* = L^q(\pi)$, where $\frac{1}{p} + \frac{1}{q} = 1$ and the map $\psi : X \mapsto L^q(\pi)$ as $\psi(x) = \varphi_x$. Effectively, this means that for every $f \in \mathcal{L}^p(\pi)$ there exists an $h \in L^p(\pi)$ such that

$$f(x) = (Ah)(x) = \langle \varphi_x, h \rangle = \int_{\Omega} \varphi(x, w)h(w)d\pi \quad (21)$$

for all $x \in X$. The norm of f is given by the infimum of the p -norm of all h which satisfy the above relation.

$$\|f\|_{\mathcal{L}^p(\pi)} := \inf_{f=Ah} \|h\|_{L^p(\pi)} \quad (22)$$

Given an infinite family of normed vector spaces \mathcal{Z}_I , indexed by some set I , we define the union of these spaces as

$$\mathcal{Z}_{\cup} := \bigcup_{i \in I} \mathcal{Z}_i := \{z \mid \exists i \in I \text{ such that } z \in \mathcal{Z}_i\} \quad (23)$$

\mathcal{Z}_{\cup} is a normed vector space with the norm

$$\|z\|_{\mathcal{Z}_{\cup}} := \inf_{z \in \mathcal{Z}_i} \|z\|_{\mathcal{Z}_i} \quad (24)$$

3.2 Unions of p-Norm RKBS

We want to show that the RKBS $\mathcal{F}(X, \Omega)$ can be written as a union of $\mathcal{L}^p(\pi)$ for all $\pi \in \mathcal{P}(\Omega)$. We prove this in two steps, first we show the statement holds for $p = 1$ and then we show how this extends to any $p > 1$.

For this section, we use the assumptions of Section 1.3. The assumption that $\varphi \in C_0(X \times \Omega)$ implies that $\varphi_x \in L^q(\pi)$ for any $x \in X$, $\pi \in \mathcal{P}(\Omega)$ and $1 \leq q \leq \infty$. Hence, for any $1 \leq p < \infty$, and $\pi \in \mathcal{P}(\Omega)$ the RKBS $\mathcal{L}^p(\pi)$ is well-defined and so is $\mathcal{F}(X, \Omega)$.

Theorem 11.

$$\mathcal{F}(X, \Omega) \cong \bigcup_{\pi \in \mathcal{P}(\Omega)} \mathcal{L}^1(\pi) \quad (25)$$

Proof. We will prove this theorem by showing that for every Radon measure $\mu \in \mathcal{M}(\Omega)$ there exists a $\pi \in \mathcal{P}(\Omega)$ and function $h \in L^1(\pi)$ such that $Ah = A\mu$ and $\|h\|_{L^1(\pi)} \leq \|\mu\|_{\mathcal{M}(\Omega)}$ and conversely that for every $\pi \in \mathcal{P}(\Omega)$ and function $h \in L^1(\pi)$ there exists a Radon measure $\mu \in \mathcal{M}(\Omega)$ such that $A\mu = Ah$ and $\|\mu\|_{\mathcal{M}(\Omega)} \leq \|h\|_{L^1(\pi)}$. The equality of the norms then follows from taking the corresponding infima.

Let $\mu \in \mathcal{M}(\Omega)$ and define the $\pi \in \mathcal{P}(\Omega)$ as

$$\pi := \frac{1}{|\mu|(\Omega)} |\mu|$$

As $\mu \ll |\mu| \ll \pi$, we define h as the Radon-Nikodym derivative $h := \frac{d\mu}{d\pi}$. Then

$$Ah = \int_{\Omega} \varphi(\cdot, w) h(w) d\pi = \int_{\Omega} \varphi(\cdot, w) d\mu = A\mu$$

and

$$\|h\|_{L^1(\pi)} = \int_{\Omega} |h| d\pi = \int_{\Omega} \text{sgn}(h) h d\pi = \int_{\Omega} \text{sgn}(h) d\mu \leq |\mu|(\Omega)$$

where $\text{sgn}(h)$ denotes the sign of h .

Conversely, let $\pi \in \mathcal{P}$ and $h \in L^1(\pi)$. Define the positive measures μ^+, μ^- such that $d\mu^+ := \max\{h, 0\} d\pi$ and $d\mu^- := \min\{h, 0\} d\pi$ and define $\mu := \mu^+ - \mu^-$. Clearly, $d\mu = h d\pi$ and

$$A\mu = \int_{\Omega} \varphi(\cdot, w) d\mu = \int_{\Omega} \varphi(\cdot, w) h(w) d\pi = Ah$$

as well as

$$|\mu|(\Omega) = \int_{\Omega} d\mu^+ + \int_{\Omega} d\mu^- = \int_{\Omega} |h| d\pi = \|h\|_{L^1(\pi)}$$

□

Next, we show that $\mathcal{L}^1(\pi)$ can be identified as the union of the spaces $\mathcal{L}^p(\tau)$ over all $\tau \in \mathcal{P}(\Omega)$ which are absolutely continuous with respect to $\pi \in \mathcal{P}(\Omega)$. Notably, this result is independent of the value of p . This is the case, because for every function $h \in \mathcal{L}^1(\pi)$ we can construct a $\tau \in \mathcal{P}(\Omega)$ such that h becomes constant and for constant functions all p -norms are equal.

Theorem 12. *Let $\pi \in \mathcal{P}(\Omega)$ and let $1 \leq p < \infty$, then*

$$\mathcal{L}^1(\pi) \cong \bigcup_{\tau \ll \pi} \mathcal{L}^p(\tau) \tag{26}$$

where $\tau \ll \pi$ denotes all $\tau \in \mathcal{P}(\Omega)$ which are absolutely continuous with respect to π .

Proof. The proof is similarly structured as the proof of the previous theorem. We will show that for every $h \in L^1(\pi)$ there exists a $\tau \in \mathcal{P}(\Omega)$ and function $h_p \in L^p(\tau)$ such that $Ah = Ah_p$ and $\|h_p\|_{L^p(\tau)} \leq \|h\|_{L^1(\pi)}$. Conversely, for every $\tau \in \mathcal{P}(\Omega)$ and function $h_p \in L^p(\tau)$ there exists a function $h \in L^1(\pi)$ such that $Ah_p = Ah$ and $\|h\|_{L^1(\pi)} \leq \|h_p\|_{L^p(\tau)}$. Again, taking infima gives the equality of the norms.

Let $h \in L^1(\pi)$. If $h = 0$, then choose $h_p = 0$. Otherwise define the measure τ such that

$$d\tau = \frac{|h|}{\|h\|_{L^1(\pi)}} d\pi$$

Clearly, $\tau \ll \pi$ and a probability measure as

$$\tau(\Omega) = \frac{1}{\|h\|_{L^1(\pi)}} \int_{\Omega} |h(w)| d\pi(w) = 1$$

Next, we define the function

$$h_p(w) = \begin{cases} \|h\|_{L^1(\pi)} \frac{h(w)}{|h(w)|} & h(w) \neq 0 \\ 0 & h(w) = 0 \end{cases}$$

τ and g are constructed such that $h d\pi = h_p d\tau$, hence

$$Ah = \int_{\Omega} \varphi(\cdot, w) h(w) d\pi(w) = \int_{\Omega} \varphi(\cdot, w) h_p(w) d\tau(w) = Ah_p$$

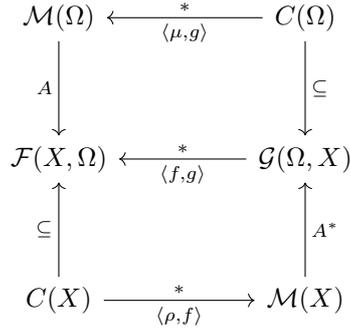


Figure 1: A diagram depicting the different relations between the Banach spaces of continuous functions $C(X), C(\Omega)$, of Radon measures $\mathcal{M}(X), \mathcal{M}(\Omega)$ and the RKBS $\mathcal{F}(X, \Omega), \mathcal{G}(\Omega, X)$ which are defined in (15) and (28) respectively. The star $*$ denotes a duality relation with the corresponding representation of the duality pairing as used in Lemma 14. The A and A^* denote the embeddings (14) and (29) respectively. The \subseteq relation denotes that one space is a subspace of the other, but the norm is not necessarily inherited.

Furthermore as $|h_p|$ is constant where $h \neq 0$

$$\|h_p\|_{L^p(\tau)}^p = \int_{\Omega} |h_p(w)|^p d\tau(w) = \|h\|_{L^1(\pi)}^{p-1} \int_{\Omega} |h(w)| d\pi(w) = \|h\|_{L^1(\pi)}^p$$

Conversely, let $\tau \in \mathcal{P}(\Omega)$ such that $\tau \ll \pi$ and $h_p \in L^p(\tau)$. Define the function $h := h_p \frac{d\tau}{d\pi}$. Then

$$Ah_p = \int_{\Omega} \varphi(\cdot, w) h_p(w) d\tau(w) = \int_{\Omega} \varphi(\cdot, w) h(w) d\pi(w) = Ah$$

and as τ is a finite measure

$$\|h\|_{L^1(\pi)} = \int_{\Omega} |h(w)| d\pi(w) = \int_{\Omega} |h_p(w)| d\tau(w) = \|h_p\|_{L^1(\tau)} \leq \|h_p\|_{L^p(\tau)}$$

□

Corollary 13.

$$\mathcal{F}(X, \Omega) \cong \bigcup_{\pi \in \mathcal{P}(\Omega)} \mathcal{L}^p(\pi) \quad (27)$$

Note that due to Theorem 9 and the fact that $\mathcal{L}^2(\pi)$ is an RKHS, we recover the result of E et al. [2021].

4 Duality and Adjoints for RKBS

In this section, we characterise the dual space of the class of integral RKBS $\mathcal{F}(X, \Omega)$. Interestingly, the dual space can be interpreted as the continuous functions of the weights which can be described by an integral over φ with measures over the data instead of the weights. This space is also an RKBS and hence forms an adjoint pair of RKBS together with $\mathcal{F}(X, \Omega)$. Using these adjoints, we formulate a primal-dual optimisation problem for both infinite and finite data. See Figure 1 for an overview of the different spaces and the duality pairings.

4.1 Duality Pairing

We define the following Banach space $\mathcal{G}(\Omega, X)$ as a subset of the continuous functions $C_0(\Omega)$ which can be described using Radon measures $\mathcal{M}(X)$ over the data:

$$\begin{aligned}
\mathcal{G}(\Omega, X) &:= \{g \in C_0(\Omega) \mid \exists \rho \in \mathcal{M}(X) \text{ s.t. } A^* \rho = g\} \\
\|g\|_{\mathcal{G}(\Omega, X)} &:= \|g\|_{C_0(\Omega)}
\end{aligned} \quad (28)$$

Here, the A^* maps data measures $\rho \in \mathcal{M}(X)$ to continuous functions over the weights $C_0(\Omega)$.

$$(A^* \rho)(w) := \int_X \varphi(x, w) d\rho(x) \quad (29)$$

for all $w \in \Omega$ and $\rho \in \mathcal{M}(X)$. The assumption that $\varphi \in C_0(\mathcal{X} \times \Omega)$ in combination with the dominated convergence theorem imply that for all $\rho \in \mathcal{M}(X)$, $A^*\rho \in C_0(\Omega)$ and hence the domain of A^* is the full space of measures $\mathcal{M}(X)$. Also, note that by Definition 4, $\mathcal{G}(\Omega, X)$ is an RKBS, as point evaluation is bounded by the supremum norm.

We will show that $\mathcal{F}(X, \Omega)$ can be identified with the dual space of $\mathcal{G}(\Omega, X)$. The converse is not true, which is a consequence of the fact that the space of continuous functions $C_0(\Omega)$ is not a reflexive. First, we will show that there exists a well-defined pairing between $\mathcal{G}(\Omega, X)$ and $\mathcal{F}(X, \Omega)$.

Lemma 14 (Data Weight Pairing). *Consider the bi-linear map defined as*

$$\langle f, g \rangle := \langle \rho, f \rangle = \langle \mu, g \rangle = \langle \rho \times \mu, \varphi \rangle = \int_{X \times \Omega} \varphi(x, w) d(\rho \times \mu)(x, w) \quad (30)$$

for all $g \in \mathcal{G}(\Omega, X)$, $f \in \mathcal{F}(X, \Omega)$ and where $\rho \in \mathcal{M}(X)$, $\mu \in \mathcal{M}(\Omega)$ such that $g = A^*\rho$, $f = A\mu$. Here $\langle \rho, f \rangle$ is a pairing between $C_0(X)$ and $\mathcal{M}(X)$, $\langle \mu, g \rangle$ between $C_0(\Omega)$ and $\mathcal{M}(\Omega)$ and $\langle \rho \times \mu, \varphi \rangle$ between $C_0(X \times \Omega)$ and $\mathcal{M}(X \times \Omega)$. (See also Figure 1)

This map is well-defined and bounded

$$|\langle f, g \rangle| \leq \|f\|_{\mathcal{F}(X, \Omega)} \|g\|_{\mathcal{G}(\Omega, X)} \quad (31)$$

Proof. First, we show that (30) is well-defined. Let $g \in \mathcal{G}(\Omega, X)$, $f \in \mathcal{F}(X, \Omega)$ and let $\rho \in \mathcal{M}(X)$, $\mu \in \mathcal{M}(\Omega)$ such that $g = A^*\rho$, $f = A\mu$. By the assumptions in Section 1.3, $\varphi \in C_0(X \times \Omega)$ and the product of two Radon measures is again a Radon measure, hence

$$\int_{X \times \Omega} \varphi(x, w) d(\rho \times \mu)(x, w) = \langle \rho \times \mu, \varphi \rangle$$

is well-defined and finite, where the pairing is between $C_0(X \times \Omega)$ and $\mathcal{M}(X \times \Omega)$. This means we can use Fubini's theorem

$$\int_{X \times \Omega} \varphi(x, w) d(\rho \times \mu)(x, w) = \int_X \int_{\Omega} \varphi(x, w) d\mu(w) d\rho(x) = \int_X f(x) d\rho(x) = \langle \rho, f \rangle$$

where this last pairing is between $C_0(X)$ and $\mathcal{M}(X)$, as well as

$$\int_{X \times \Omega} \varphi(x, w) d(\rho \times \mu)(x, w) = \int_{\Omega} \int_X \varphi(x, w) d\rho(x) d\mu(w) = \int_{\Omega} g(w) d\mu(w) = \langle \mu, g \rangle$$

where this last pairing is between $C_0(\Omega)$ and $\mathcal{M}(\Omega)$.

We can show that $\langle f, g \rangle$ is independent on the choice of μ and ρ . If $\rho' \in \mathcal{M}(X)$, $\mu' \in \mathcal{M}(\Omega)$ such that $g = A^*\rho'$, $f = A\mu'$ then

$$\begin{aligned} \int_{X \times \Omega} \varphi(x, w) d(\rho' \times \mu')(x, w) &= \langle \rho', f \rangle = \int_{X \times \Omega} \varphi(x, w) d(\rho' \times \mu)(x, w) \\ &= \langle \mu, g \rangle = \int_{X \times \Omega} \varphi(x, w) d(\rho \times \mu)(x, w) \end{aligned}$$

Hence, the map is well-defined. Furthermore, the pairing is bounded as

$$|\langle f, g \rangle| = |\langle \mu, g \rangle| \leq \|\mu\|_{\mathcal{M}(\Omega)} \|g\|_{C_0(\Omega)}$$

for all $\mu \in \mathcal{M}(\Omega)$ such that $f = A\mu$. Hence, by taking the corresponding infimum we get the inequality

$$|\langle f, g \rangle| \leq \|f\|_{\mathcal{F}(X, \Omega)} \|g\|_{\mathcal{G}(\Omega, X)}$$

□

This lemma also implies that A^* can be understood as an adjoint of A as for all $\rho \in \mathcal{M}(X)$, $\mu \in \mathcal{M}(\Omega)$

$$\langle \rho, A\mu \rangle = \langle \mu, A^*\rho \rangle = \langle A^*\rho, A\mu \rangle \quad (32)$$

Here the first pairing is between $C_0(X)$ and $\mathcal{M}(X)$, the second between $C_0(\Omega)$ and $\mathcal{M}(\Omega)$ and the third between $\mathcal{G}(\Omega, X)$ and $\mathcal{F}(X, \Omega)$.

With this pairing, we can show that $\mathcal{F}(X, \Omega)$ can indeed be identified as the dual of $\mathcal{G}(\Omega, X)$. We prove this by leveraging the fact that the dual of a subspace can be identified as the quotient space with respect to its annihilator. As $\mathcal{F}(X, \Omega)$ is isomorphic to the quotient space $\mathcal{M}(\Omega)/\mathcal{N}(A)$ by definition, the proof consists of showing that $\mathcal{N}(A)$ is indeed the annihilator of $\mathcal{G}(\Omega, X)$.

Theorem 15. *Let $\mathcal{G}^*(\Omega, X)$ be the dual space of $\mathcal{G}(\Omega, X)$, then $\mathcal{G}^*(\Omega, X) \cong \mathcal{F}(X, \Omega)$*

Proof. The dual of a normed subspace \mathcal{Y} of a Banach space \mathcal{Z} can be identified as a quotient space of the dual space of the full space \mathcal{Z}^* over the annihilator of the subspace \mathcal{Y}^\perp [Rudin, 1991, Theorem 4.8]. Here, the annihilator \mathcal{Y}^\perp is defined as

$$\mathcal{Y}^\perp := \{z^* \in \mathcal{Z}^* \mid \forall y \in \mathcal{Y}, \langle z^*, y \rangle = 0\} \quad (33)$$

where the pairing is between the spaces \mathcal{Z} and \mathcal{Z}^* . Remember, that $\mathcal{F}(X, \Omega)$ is defined as the quotient space $\mathcal{M}(\Omega)/\mathcal{N}(A)$. So it only remains to prove that $\mathcal{G}(\Omega, X)^\perp \cong \mathcal{N}(A)$, where

$$\mathcal{G}(\Omega, X)^\perp := \{\mu \in \mathcal{M}(\Omega) \mid \forall g \in \mathcal{G}(\Omega, X), \langle \mu, g \rangle = 0\} \quad (34)$$

The isometry is trivial as both spaces use the total variation norm of $\mathcal{M}(\Omega)$. We prove that these spaces are isomorphic in two steps.

Let $\mu \in \mathcal{G}(\Omega, X)^\perp$, $\rho \in \mathcal{M}(X)$ and define the function $g := A^*\rho \in \mathcal{G}(\Omega, X)$. Using Lemma 14

$$0 = \langle \mu, g \rangle = \langle A\mu, \rho \rangle$$

As this holds for all $\rho \in \mathcal{M}(X)$ we get that $A\mu \equiv 0$ and hence $\mu \in \mathcal{N}(A)$.

Conversely, let $\mu \in \mathcal{N}(A)$ and $g \in \mathcal{G}(\Omega, X)$. Furthermore, let $\rho \in \mathcal{M}(X)$ such that $A^*\rho = g$. Again by Lemma 14

$$0 = \langle A\mu, \rho \rangle = \langle \mu, g \rangle$$

Hence, $\mu \in \mathcal{G}(\Omega, X)^\perp$ □

As both $\mathcal{F}(X, \Omega)$ and $\mathcal{G}(\Omega, X)$ are RKBSs, we can show that they form an adjoint pair of RKBSs and φ is the reproducing kernel in the sense of Definition 7, as $\varphi(x, w) = \langle A\delta_w, A^*\delta_x \rangle$.

Corollary 16. *$\mathcal{F}(X, \Omega)$ and $\mathcal{G}(\Omega, X)$ form an adjoint pair of RKBS with the reproducing kernel $K = \varphi$*

Proof. Define $K := \varphi$. Then

$$K(x, \cdot) = \varphi_x = \int_X \varphi(x', \cdot) d\delta_x(x') = A^*\delta_x \in \mathcal{G}(\Omega, X)$$

as $\delta_x \in \mathcal{M}(X)$ for all $x \in X$, i.e. all point measures are Radon measures. Furthermore, let $f \in \mathcal{F}(X, \Omega)$, then by Lemma 14

$$\langle f, K(x, \cdot) \rangle = \langle f, A^*\delta_x \rangle = \langle \delta_x, f \rangle = f(x)$$

for all $x \in X$.

For the adjoint we find that

$$K(\cdot, w) = \varphi_w = \int_\Omega \varphi(\cdot, w') d\delta_w(w') = A\delta_w \in \mathcal{F}(X, \Omega)$$

as $\delta_w \in \mathcal{M}(\Omega)$ for all $w \in \Omega$. Furthermore, let $g \in \mathcal{G}(\Omega, X)$ then again by Lemma 14

$$\langle K(\cdot, w), g \rangle = \langle A\delta_w, g \rangle = \langle \delta_w, g \rangle = g(w)$$

for all $w \in \Omega$. □

4.2 Dual Formulation of ERM with Infinite Data

The description of the dual space allows a dual formulation of the corresponding empirical risk optimisation (ERM) problem for integral RKBS. This opens the door to many primal-dual type algorithms to be used for these problems. First, we will consider the infinite data case, before sampling to finite data.

Given data pairs $(x, y(x))$ for all $x \in X$ and a probability measure $\nu \in \mathcal{M}(X)$ where it is assumed that $y \in L^2(\nu)$, we define the ERM as

$$\inf_{\mu \in \mathcal{M}(\Omega)} J(A\mu) + R(\mu) \quad (35)$$

where the data fidelity term $J : \mathcal{F}(X, \Omega) \mapsto \mathbb{R}$ and the regularisation term $R : \mathcal{M}(\Omega)$ are defined by

$$\begin{aligned} J(A\mu) &:= \frac{1}{2} \|A\mu - y\|_{L^2(\nu)}^2 = \frac{1}{2} \int_X ((A\mu)(x) - y(x))^2 d\nu(x) \\ R(\mu) &:= \|\mu\|_{\mathcal{M}(\Omega)} = |\mu|(\Omega) \end{aligned} \quad (36)$$

Note that for the infinite setting, the optimisation problem is convex. The dual problem can be constructed in terms of the convex conjugates of J and R , also referred to as the Fenchel–Legendre transform. However, as $\mathcal{M}(\Omega)$ and $\mathcal{F}(X, \Omega)$ are already dual spaces of a non-reflexive Banach space, we define the convex conjugates $J^* : \mathcal{M}(X) \mapsto \mathbb{R}$ and $R^* : \mathcal{G}(\Omega, X) \mapsto \mathbb{R}$ in terms of the pre-dual

$$\begin{aligned} J^*(\rho) &:= \sup_{f \in \mathcal{F}(X, \Omega)} (\langle \rho, f \rangle - J(f)) \\ R^*(g) &:= \sup_{\mu \in \mathcal{M}(\Omega)} (\langle \mu, g \rangle - R(\mu)) \end{aligned} \quad (37)$$

where $\rho \in \mathcal{M}(X)$ and $g \in C_0(\Omega)$. We can write J^* as a function of $\rho \in \mathcal{M}(X)$ instead of $g \in \mathcal{G}(\Omega, X)$ as $\langle f, A^*\rho \rangle = \langle \rho, f \rangle$ by Lemma 14.

Lemma 17. *If $\mathcal{F}(X, \Omega)$ is dense in $L^2(\nu)$, then the convex conjugates of J and R are given by*

$$\begin{aligned} J^*(\rho) &= \begin{cases} \int_X \frac{1}{2} \frac{d\rho}{d\nu}(x) + y(x) d\rho(x) & \rho \ll \nu \\ \infty & \text{otherwise} \end{cases} \\ R^*(g) &= \begin{cases} 0 & \|g\|_{C_0(\Omega)} \leq 1 \\ \infty & \text{otherwise} \end{cases} \end{aligned} \quad (38)$$

for all $\rho \in \mathcal{M}(X)$ and $g \in C_0(X)$

Proof. First, we prove the identity for J^* . Let $\rho \in \mathcal{M}(X)$. We consider $\rho \ll \nu$. Then for some $f \in \mathcal{F}(X, \Omega)$

$$\begin{aligned} \langle \rho, f \rangle - J(f) &= \int_X f(x) d\rho(x) - \frac{1}{2} \int_X (f(x) - y(x))^2 d\nu(x) \\ &= \int_X f(x) \frac{d\rho}{d\nu}(x) - \frac{1}{2} (f(x) - y(x))^2 d\nu(x) \\ &= \int_X -\frac{1}{2} \left(f(x) - y(x) - \frac{d\rho}{d\nu}(x) \right)^2 + \frac{d\rho}{d\nu}(x) \left(\frac{1}{2} \frac{d\rho}{d\nu}(x) + y(x) \right) d\nu(x) \end{aligned}$$

where we have completed the square. By the assumption that $\mathcal{F}(X, \Omega)$ is dense $L^2(\nu)$, the square vanishes in the supremum, and we get

$$\int_X \frac{d\rho}{d\nu}(x) \left(\frac{1}{2} \frac{d\rho}{d\nu}(x) + y(x) \right) d\nu(x) = \int_X \frac{1}{2} \frac{d\rho}{d\nu}(x) + y(x) d\rho(x)$$

Next, when ρ is not absolutely continuous with respect to ν , then there exists a Borel set B of X such that $\nu(B) = 0$ but $\rho(B) \neq 0$. Then for a sequence of $f_n \in \mathcal{F}(X, \Omega)$, $n \in \mathbb{N}$, converging to the indicator function of B

$$\lim_{c \rightarrow \infty} \lim_{n \rightarrow \infty} (\langle \rho, cf_n \rangle - J(cf_n)) = \lim_{c \rightarrow \infty} c\rho(B) = \infty$$

Finally, we prove the identity for R^* , which is a standard argument. Let $g \in C_0(\Omega)$. First, we consider the case when g is in the unit ball, i.e. $\|g\|_{C_0(\Omega)} \leq 1$. Then for any $\mu \in \mathcal{M}(\Omega)$ we find that

$$\langle \mu, g \rangle - R(\mu) \leq (\|g\|_{C_0(\Omega)} - 1) \|\mu\|_{\mathcal{M}(\Omega)} \leq 0$$

Hence, the supremum is achieved when μ is the zero measure. Next, when $\|g\|_{C_0(\Omega)} > 1$, there exists a $w \in \Omega$ such that $|g(w)| > 1$. Taking the point measures $c \frac{|g(w)|}{g(w)} \delta_w$ then

$$\lim_{c \rightarrow \infty} \left(\left\langle c \frac{|g(w)|}{g(w)} \delta_w, g \right\rangle - R \left(c \frac{|g(w)|}{g(w)} \delta_w \right) \right) = \lim_{c \rightarrow \infty} c(|g(w)| - 1) = \infty$$

□

By Fenchel's duality theorem, we get the strong duality equality for (35).

Theorem 18. *The strong duality equality holds*

$$\inf_{\mu \in \mathcal{M}(\Omega)} J(A\mu) + R(\mu) = \sup_{\rho \in \mathcal{M}(X)} -J^*(-\rho) - R^*(A^*\rho) \quad (39)$$

The perturbation function for this problem is given by

$$F(\mu, f) = J(A\mu - f) + R(\mu) \quad (40)$$

where $\mu \in \mathcal{M}(\Omega)$ and $f \in \mathcal{F}(X, \Omega)$.

Proof. By Fenchel's duality theorem, strong duality holds if there exists a $\mu \in \mathcal{M}(\Omega)$ such that $R(\mu)$ is finite and J is finite and continuous at $A\mu$. Clearly, $R(\mu)$ is finite for any $\mu \in \mathcal{M}(\Omega)$ and J is continuous by dominated convergence as $A\mu \in C_0(X)$. □

4.3 Dual Formulation of ERM with Finite Data

If we have finite data, i.e. $X = \{x_1, \dots, x_n\}$, then we can rewrite the dual problem in simpler terms. When X is finite, then the dimension of $\mathcal{M}(X)$ is also finite and consequently so are $\mathcal{G}(\Omega, X)$ and $\mathcal{F}(X, \Omega)$. Due to the fact that the point measures are the extreme points of the unit ball of $\mathcal{M}(\Omega)$, the infimum in the definition of the norm of $\mathcal{F}(X, \Omega)$ (16) has a minimiser which is a linear combination of point measures δ_{w_i} , for certain $w_i \in \Omega$ and $i \in \{1, \dots, n\}$. This is the basis for the representer theorems of Bartolucci et al. [2023], Bredies and Carioni [2019]

Theorem 19. [Bartolucci et al., 2023, Theorem 3.9] *Let $X = \{x_1, \dots, x_n\}$, for some $n \in \mathbb{N}$. Then the ERM (35) admits a minimiser μ^* of the form*

$$\mu^* = \sum_{i=1}^n \alpha_i \delta_{w_i} \quad (41)$$

for some $\alpha_i \in \mathbb{R}, w_i \in \Omega$ where $i \in \{1, \dots, n\}$

Note that $f^* := A\mu^*$ can be written in terms of the reproducing kernel $K = \varphi$

$$f^*(x) := (A\mu^*)(x) = \sum_{i=1}^n \alpha_i \varphi(x, w_i)$$

The trade-off for the representer theorem is that finding the w_i is a non-convex problem. So we have traded the convexity of the ERM for finite dimensionality, i.e. we only need to find $2n$ parameters. This is specific to RKBS, in contrast to RKHS, where the finite dimensional problem remains convex.

As X is finite, we have finite data pairs (x_i, y_i) and the probability measure $\nu := \sum_{i=1}^n \gamma_i \delta_{x_i}$, where $\gamma_i > 0$ for $i \in \{1, \dots, n\}$ and $\sum_{i=1}^n \gamma_i = 1$. Similarly, we can write any measure $\rho \in \mathcal{M}(X)$ as $\rho := \sum_{i=1}^n \lambda_i \delta_{x_i}$, where $\lambda_i \in \mathbb{R}$ for $i \in \{1, \dots, n\}$. The dual ERM problem then becomes

Theorem 20. *Let X have cardinality $n < \infty$, then we can rewrite the primal ERM (35) using*

$$\begin{aligned} J(\alpha, \mathbf{w}) &= \sum_{i,j=1}^n \gamma_i (\alpha_j \varphi(x_i, w_j) - y_i)^2 \\ R(\alpha) &= \|\alpha\|_1 \end{aligned} \quad (42)$$

for $\alpha = (\alpha_1, \dots, \alpha_n) \in \mathbb{R}^n$, $\mathbf{w} = (w_1, \dots, w_n) \in \Omega^n$.

The dual problem can be written using

$$\begin{aligned} J^*(\lambda) &= \sum_{i=1}^n \lambda_i \left(\frac{\lambda_i}{2\gamma_i} + y_i \right) \\ R^*(\lambda) &= \begin{cases} 0 & -1 \leq \sum_{i=1}^n \lambda_i \varphi(x_i, w) \leq 1 \quad \forall w \in \Omega \\ \infty & \text{otherwise} \end{cases} \end{aligned} \quad (43)$$

for $\lambda = (\lambda_1, \dots, \lambda_n) \in \mathbb{R}^n$.

The strong duality equality holds

$$\inf_{\alpha \in \mathbb{R}^n, \mathbf{w} \in \Omega^n} J(\alpha, \mathbf{w}) + R(\alpha) = \sup_{\lambda \in \mathbb{R}^n} -J^*(-\lambda) - R^*(\lambda)$$

Proof. The statement follows directly from Lemma 17 and Theorems 18 and 19. Note that by the assumption that $\gamma_i > 0$ for $i \in \{1, \dots, n\}$, all measures $\rho \in \mathcal{M}(X)$ are absolutely continuous with respect to ν and

$$\frac{d\rho}{d\nu}(x_i) = \frac{\lambda_i}{\gamma_i}$$

for every $i \in \{1, \dots, n\}$. □

5 Conclusion and Discussion

In this paper, we have developed duality theory for neural networks, including the Barron spaces of E et al. [2019b] and the integral RKBS of Bartolucci et al. [2023]. A key conceptual result is that the weights and the data are dual concepts. On one side, we have neural networks which represent functions of the data parameterised by some weights, where on the other side we have the dual networks which are functions of the weights parameterised by some data points. These dual networks can be thought of as representing the range of functions which can be represented given some sample of the data, which is closely related to complexity measures, such as the Rademacher complexity.

We have considered RKBS instead of more commonly used RKHS. We have shown that spaces for neural networks such as the Barron spaces can be properly described by RKBS. However, a common problem which involves RKBS is that, when considering finite data samples, the ERM optimisation problem becomes non-convex. One of the main advantages of RKHS-based methods, such as support vector machines, is that the ERM is convex and thus relatively easy to solve. This is in contrast to neural networks, where gradient descent based methods are used. The class of integral RKBS we have investigated, however, is a much more flexible function space, as it is a union of a large class of different RKHS.

We have shown that the Barron space \mathcal{B}_σ can be identified with the integral RKBS $\mathcal{F}(X, \Omega)$. This theorem holds for the most commonly used activation functions σ , like the ReLU, sigmoid, inverse tangent and sine activation functions are valid. For activation functions which are higher powers of the ReLU function, special considerations are necessary. However, with the proper definition of the Barron norm for these activation functions, we conjecture that a similar statement holds.

The dual viewpoint has multiple use cases. It can be used to bring a foundation to architecture search methods for neural networks [Bungert et al., 2021]. These methods are based on Bregman iteration, which are closely related to augmented Lagrangian techniques [Brune et al., 2011]. Another use case is related to sampling. The question of which data points to sample is often referred to as active learning or sampling in experimental design [Settles, 2009]. The functions in the dual space enable a way to address this problem for neural networks. On the other hand, adversarial attacks try to find input data which gives an undesired result [Madry et al., 2019]. In the language of this paper, this means, given slightly different data measures, leads to significant differences in the dual space. Therefore, this work can give a new viewpoint on quantifying adversarial robustness.

One of the key advances in the field of neural networks was the change from shallow to deep neural networks. Most literature dealing with the theoretical study of neural networks has so far been restricted

to shallow neural networks [E et al., 2019b, Parhi and Nowak, 2021]. Recently, Unser [2019] proved a representer theorem for certain deep networks with the ReLU activation function and E and Wojtowysch [2020a] developed generalised Barron spaces which deal with deep neural networks. In principle, there are no barriers to considering deep neural networks in our integral RKBS framework (15), by taking the kernel φ to be a concatenation of activation functions σ and affine functions based on parameters w^j for layers $j \in \{1, \dots, L\}$. Then we can still take measures $\mu \in \mathcal{M}(\Omega)$ where the parameters are from $\Omega = \Omega_1 \times \dots \times \Omega_L$. We are not sure if the resulting norm of $\mathcal{F}(X, \Omega)$ is appropriate for deep neural networks, as it is then mostly dependent on the size of weights in the last (linear) layer. However, we conjecture that the ideas of the proof 9 can be extended to prove a similar result for generalised Barron spaces.

Acknowledgements

C.B. acknowledges support by the European Union’s Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 777826 (NoMADS). L.S. thanks Stephan van Gils for his support and encouragement for taking the research in this direction.

References

- N. Aronszajn. Theory of reproducing kernels. *Transactions of the American Mathematical Society*, 68(3):337–404, 1950. ISSN 0002-9947, 1088-6850. doi: 10.1090/S0002-9947-1950-0051437-7. URL <https://www.ams.org/tran/1950-068-03/S0002-9947-1950-0051437-7/>.
- Francis Bach. Breaking the Curse of Dimensionality with Convex Neural Networks. *Journal of Machine Learning Research*, 18:1–53, April 2017a.
- Francis Bach. On the Equivalence between Kernel Quadrature Rules and Random Feature Expansions. *Journal of Machine Learning Research*, 18:1–38, April 2017b.
- A.R. Barron. Universal approximation bounds for superpositions of a sigmoidal function. *IEEE Transactions on Information Theory*, 39(3):930–945, May 1993. ISSN 1557-9654. doi: 10.1109/18.256500. Conference Name: IEEE Transactions on Information Theory.
- Francesca Bartolucci, Ernesto De Vito, Lorenzo Rosasco, and Stefano Vigogna. Understanding neural networks with reproducing kernel Banach spaces. *Applied and Computational Harmonic Analysis*, 62:194–236, January 2023. ISSN 1063-5203. doi: 10.1016/j.acha.2022.08.006. URL <https://www.sciencedirect.com/science/article/pii/S1063520322000768>.
- Kristian Bredies and Marcello Carioni. Sparsity of solutions for variational inverse problems with finite-dimensional data. *Calculus of Variations and Partial Differential Equations*, 59(1):14, December 2019. ISSN 1432-0835. doi: 10.1007/s00526-019-1658-1. URL <https://doi.org/10.1007/s00526-019-1658-1>.
- Christoph Brune, Alex Sawatzky, and Martin Burger. Primal and Dual Bregman Methods with Application to Optical Nanoscopy. *International Journal of Computer Vision*, 92(2): 211–229, April 2011. ISSN 0920-5691, 1573-1405. doi: 10.1007/s11263-010-0339-5. URL <http://link.springer.com/10.1007/s11263-010-0339-5>.
- Leon Bungert, Tim Roith, Daniel Tenbrinck, and Martin Burger. Neural Architecture Search via Bregman Iterations. *arXiv:2106.02479 [cs, math]*, June 2021. URL <http://arxiv.org/abs/2106.02479>. arXiv: 2106.02479.
- Stephane Canu, Xavier Mary, and Alain Rakotomamonjy. Functional learning through kernels. *Advances in Learning Theory: Methods, Models and Application*, pages 89–110, October 2003. URL <https://hal.archives-ouvertes.fr/hal-00422240>. Publisher: IOS Press.
- Antonin Chambolle. An Algorithm for Total Variation Minimization and Applications. *Journal of Mathematical Imaging and Vision*, 20(1):89–97, January 2004. ISSN 1573-7683. doi: 10.1023/B:JMIV.0000011325.36760.1e. URL <https://doi.org/10.1023/B:JMIV.0000011325.36760.1e>.
- Tony F. Chan, Gene H. Golub, and Pep Mulet. A Nonlinear Primal-Dual Method for Total Variation-Based Image Restoration. *SIAM Journal on Scientific Computing*, 20(6):

- 1964–1977, January 1999. ISSN 1064-8275. doi: 10.1137/S1064827596299767. URL <https://epubs.siam.org/doi/abs/10.1137/S1064827596299767>. Publisher: Society for Industrial and Applied Mathematics.
- Lenaic Chizat, Lenaic Chizat, and Universite-Paris-Saclay Fr. Implicit Bias of Gradient Descent for Wide Two-layer Neural Networks Trained with the Logistic Loss. *Proceedings of Machine Learning Research*, 125:1–34, 2020.
- G. Cybenko. Approximation by superpositions of a sigmoidal function. *Mathematics of Control, Signals and Systems*, 2(4):303–314, December 1989. ISSN 1435-568X. doi: 10.1007/BF02551274. URL <https://doi.org/10.1007/BF02551274>.
- Weinan E and Stephan Wojtowytsch. On the Banach spaces associated with multi-layer ReLU networks: Function representation, approximation theory and gradient descent dynamics. *arXiv:2007.15623 [cs, math, stat]*, July 2020a. URL <http://arxiv.org/abs/2007.15623>. arXiv: 2007.15623.
- Weinan E and Stephan Wojtowytsch. Kolmogorov Width Decay and Poor Approximators in Machine Learning: Shallow Neural Networks, Random Feature Models and Neural Tangent Kernels. *arXiv:2005.10807 [cs, math, stat]*, October 2020b. URL <http://arxiv.org/abs/2005.10807>. arXiv: 2005.10807.
- Weinan E and Stephan Wojtowytsch. Representation formulas and pointwise properties for Barron functions. *arXiv:2006.05982 [cs, math, stat]*, June 2021. URL <http://arxiv.org/abs/2006.05982>. arXiv: 2006.05982.
- Weinan E, Chao Ma, and Lei Wu. Barron Spaces and the Compositional Function Spaces for Neural Network Models. *arXiv:1906.08039 [cs, math, stat]*, June 2019a. URL <http://arxiv.org/abs/1906.08039>. arXiv: 1906.08039.
- Weinan E, Chao Ma, and Lei Wu. A Priori Estimates of the Population Risk for Two-layer Neural Networks. *Communications in Mathematical Sciences*, 17(5):1407–1425, 2019b. ISSN 15396746, 19450796. doi: 10.4310/CMS.2019.v17.n5.a11. URL <http://arxiv.org/abs/1810.06397>. arXiv: 1810.06397.
- Weinan E, Chao Ma, Stephan Wojtowytsch, and Lei Wu. Towards a Mathematical Understanding of Neural Network-Based Machine Learning: what we know and what we don’t. *arXiv:2009.10713 [cs, math, stat]*, October 2020. URL <http://arxiv.org/abs/2009.10713>. arXiv: 2009.10713.
- Weinan E, Chao Ma, and Lei Wu. The Barron Space and the Flow-induced Function Spaces for Neural Network Models. *arXiv:1906.08039 [cs, math, stat]*, March 2021. URL <http://arxiv.org/abs/1906.08039>. arXiv: 1906.08039.
- Theodoros Evgeniou, Massimiliano Pontil, and Tomaso Poggio. Regularization Networks and Support Vector Machines. *Advances in Computational Mathematics*, 13(1):1–50, April 2000. ISSN 1572-9044. doi: 10.1023/A:1018946025316. URL <https://doi.org/10.1023/A:1018946025316>.
- Antonio G. García and Alberto Portal. Sampling in Reproducing Kernel Banach Spaces. *Mediterranean Journal of Mathematics*, 10(3):1401–1417, August 2013. ISSN 1660-5454. doi: 10.1007/s00009-012-0234-0. URL <https://doi.org/10.1007/s00009-012-0234-0>.
- J. R. Giles. Classes of Semi-Inner-Product Spaces. *Transactions of the American Mathematical Society*, 129(3):436–446, 1967. ISSN 0002-9947. doi: 10.2307/1994599. URL <http://www.jstor.org/stable/1994599>. Publisher: American Mathematical Society.
- Kurt Hornik. Approximation capabilities of multilayer feedforward networks. *Neural Networks*, 4(2):251–257, January 1991. ISSN 0893-6080. doi: 10.1016/0893-6080(91)90009-T. URL <https://www.sciencedirect.com/science/article/pii/089360809190009T>.
- 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, January 1993. ISSN 0893-6080. doi: 10.1016/S0893-6080(05)80131-5. URL <https://www.sciencedirect.com/science/article/pii/S0893608005801315>.
- Rong Rong Lin, Hai Zhang Zhang, and Jun Zhang. On Reproducing Kernel Banach Spaces: Generic Definitions and Unified Framework of Constructions. *Acta Mathematica Sinica, English Series*, June 2022. ISSN 1439-7617. doi: 10.1007/s10114-022-1397-7. URL <https://doi.org/10.1007/s10114-022-1397-7>.

- G Lumer. Semi-Inner-Product Spaces. *Transactions of the American Mathematical Society*, page 15, October 1960.
- Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards Deep Learning Models Resistant to Adversarial Attacks, September 2019. URL <http://arxiv.org/abs/1706.06083>. arXiv:1706.06083 [cs, stat].
- Charles A. Micchelli and Massimiliano Pontil. Feature space perspectives for learning the kernel. *Machine Learning*, 66(2):297–319, March 2007. ISSN 1573-0565. doi: 10.1007/s10994-006-0679-0. URL <https://doi.org/10.1007/s10994-006-0679-0>.
- Stanley Osher, Martin Burger, Donald Goldfarb, Jinjun Xu, and Wotao Yin. An Iterative Regularization Method for Total Variation-Based Image Restoration. *Multiscale Modeling & Simulation*, 4(2):460–489, January 2005. ISSN 1540-3459, 1540-3467. doi: 10.1137/040605412. URL <http://epubs.siam.org/doi/10.1137/040605412>.
- Rahul Parhi and Robert D Nowak. Banach Space Representer Theorems for Neural Networks and Ridge Splines. *Journal of Machine Learning Research*, 22(1):40, February 2021.
- Anant Raj and Francis Bach. Explicit Regularization of Stochastic Gradient Methods through Duality. In *PMLR*, volume 130, page 11, San Diego, California, USA., 2021.
- Walter Rudin. *Functional Analysis*. McGraw-Hill, New York, USA, 2nd edition; mcgraw-hill international edition edition, 1991. ISBN 978-0-07-100944-7.
- Bernhard Scholkopf and Alexander J. Smola. *Learning with Kernels: Support Vector Machines, Regularization, Optimization, and Beyond*. MIT Press, June 2018. ISBN 978-0-262-53657-8.
- Burr Settles. Active Learning Literature Survey. Technical Report, University of Wisconsin-Madison Department of Computer Sciences, 2009. URL <https://minds.wisconsin.edu/handle/1793/60660>. Accepted: 2012-03-15T17:23:56Z.
- Shai Shalev-Shwartz and Tong Zhang. Stochastic Dual Coordinate Ascent Methods for Regularized Loss Minimization. *Journal of Machine Learning Research*, 14:567–599, February 2013.
- Guohui Song, Haizhang Zhang, and Fred J. Hickernell. Reproducing kernel Banach spaces with the l_1 norm. *Applied and Computational Harmonic Analysis*, 34(1):96–116, January 2013. ISSN 1063-5203. doi: 10.1016/j.acha.2012.03.009. URL <https://www.sciencedirect.com/science/article/pii/S1063520312000486>.
- Robert Tibshirani. Regression Shrinkage and Selection Via the Lasso. *Journal of the Royal Statistical Society: Series B (Methodological)*, 58(1):267–288, 1996. ISSN 2517-6161. doi: 10.1111/j.2517-6161.1996.tb02080.x. URL <https://onlinelibrary.wiley.com/doi/abs/10.1111/j.2517-6161.1996.tb02080.x>. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.2517-6161.1996.tb02080.x>.
- Michael Unser. A Representer Theorem for Deep Neural Networks. *Journal of Machine Learning Research*, 20:1–30, July 2019.
- Grace Wahba. Support vector machines, reproducing kernel Hilbert spaces, and randomized GACV. In *Advances in kernel methods: support vector learning*, pages 69–88. Departement of Statistics, University of Wisconsin-Madison, 1999.
- Haizhang Zhang and Jun Zhang. Learning with Reproducing Kernel Banach Spaces. In Pei Dang, Min Ku, Tao Qian, and Luigi G. Rodino, editors, *New Trends in Analysis and Interdisciplinary Applications*, Trends in Mathematics, pages 417–423, Cham, 2017. Springer International Publishing. ISBN 978-3-319-48812-7. doi: 10.1007/978-3-319-48812-7.
- Haizhang Zhang, Yuesheng Xu, and Jun Zhang. Reproducing kernel Banach spaces for machine learning. In *2009 International Joint Conference on Neural Networks*, pages 3520–3527, Atlanta, Ga, USA, June 2009. IEEE. ISBN 978-1-4244-3548-7. doi: 10.1109/IJCNN.2009.5179093. URL <http://ieeexplore.ieee.org/document/5179093/>.