



Citation for published version:

Korolev, Y 2021 'Two-layer neural networks with values in a Banach space'.

Publication date:
2021

[Link to publication](#)

University of Bath

Alternative formats

If you require this document in an alternative format, please contact:
openaccess@bath.ac.uk

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Two-layer neural networks with values in a Banach space

Yury Korolev*

Abstract

We study two-layer neural networks whose domain and range are Banach spaces with separable preduals. In addition, we assume that the image space is equipped with a partial order, i.e. it is a Riesz space. As the nonlinearity we choose the lattice operation of taking the positive part; in case of \mathbb{R}^d -valued neural networks this corresponds to the ReLU activation function. We prove inverse and direct approximation theorems with Monte-Carlo rates for a certain class of functions, extending existing results for the finite-dimensional case. In the second part of the paper we study, from the regularisation theory viewpoint, the problem of finding optimal representations of such functions via signed measures on a latent space from a finite number of noisy observations. We discuss regularity conditions known as source conditions and obtain convergence rates in a Bregman distance for the representing measure in the regime when both the noise level goes to zero and the number of samples goes to infinity at appropriate rates.

Keywords: Vector-valued neural networks, ReLU, Barron space, variation norm space, curse of dimensionality, Bregman distance

AMS Subject Classification: 68Q32, 68T07, 46E40, 41A65, 65J22

Contents

1	Introduction	1
2	Linear-nonlinear decompositions in Riesz spaces	4
3	Variation norm spaces	6
3.1	The finite-dimensional case	6
3.2	The infinite-dimensional case	7
3.2.1	Definitions	7
3.2.2	Approximation theorems	13
3.2.3	Examples	17
4	Finding an optimal representation	18
4.1	Source condition	19
4.2	Convergence rates in Bregman distance	22
5	Summary and conclusions	25
A	weak-* compactness of the unit ball in the space of linear bounded operators	28
B	Proof of Proposition 4.2	28

1 Introduction

Two-layer neural networks (also referred to as neural networks with one hidden layer) are functions $f: \mathbb{R}^d \rightarrow \mathbb{R}$ of the following form

$$f(x) = \sum_{i=1}^n a_i \sigma(\langle x, b_i \rangle + c_i), \quad x \in \mathbb{R}^d, \quad (1.1)$$

*Department of Mathematical Sciences, University of Bath, Claverton Down BA2 7AY, UK. email: ymk30@bath.ac.uk

where $\{b_i\}_{i=1}^n \subset \mathbb{R}^d$ and $\{a_i\}_{i=1}^n \subset \mathbb{R}$ are referred to as weights, $\{c_i\}_{i=1}^n \subset \mathbb{R}$ as biases, and $\sigma: \mathbb{R} \rightarrow \mathbb{R}$ as the activation function. The brackets $\langle \cdot, \cdot \rangle$ denote the scalar product in \mathbb{R}^d . Individual summands $\{\sigma(\langle x, b_i \rangle + c_i)\}_{i=1}^n$ are referred to as neurons and collectively they are referred to as the hidden layer of the network.

The famous universal approximation theorem [1–3] states that if σ is not a polynomial then any continuous function on a compact set in \mathbb{R}^d can be approximated arbitrary well (in the supremum norm) with functions of the type (1.1). However, quantitative estimates (such as the number of neurons required to achieve a certain accuracy) that can be obtained in general depend on the dimension of the input space d , with the approximation error scaling as $O(n^{-d})$. For high-dimensional inputs ($d \gg 1$) this is not satisfactory, an effect which is known as the curse of dimensionality. Barron [4] showed that for L^1 functions whose Fourier transform satisfies a certain integrability condition dimension-independent Monte-Carlo rates $O(1/\sqrt{n})$ can be obtained for the approximation error in L^2 .

The condition introduced by Barron is, however, too conservative, and fails for many functions that can be approximated by (1.1) with dimension-independent rates. More general spaces, termed *variation norm spaces*, were studied in, e.g., [5–8]. Roughly speaking, these spaces consist of functions whose expansion coefficients in a certain basis or frame are absolutely summable. This approach was further extended in [9] who replaced the basis/frame expansion with an integral over a compact latent space against a finite Radon measure. This work was continued in [10], where it was shown that these spaces are in some sense optimal: they contain all weak-* limits of (1.1) as $n \rightarrow \infty$ if (1.1). We postpone the details to Section 3.1. These spaces were called variation norm (or \mathcal{F}_1) spaces in [9] and Barron spaces in [10, 11], not to be confused with the spaces introduced by Barron in [4]. We will use the notation \mathcal{F}_1 for these spaces.

Similar results can be arrived at from the mean-field perspective [12–15], where a $\frac{1}{n}$ scaling is assumed in (1.1)

$$f(x) = \frac{1}{n} \sum_{i=1}^n a_i \sigma(\langle x, b_i \rangle + c_i), \quad x \in \mathbb{R}^d, \quad (1.2)$$

and individual neurons are interpreted as interacting particles moving in a potential determined by the loss function. Since the optimal (trained) coefficients $\{a_i\}_{i=1}^n$ depend on n (cf. Remark 3.23), the formulations (1.1) and (1.2) are, in fact, equivalent.

A related concept is that of random feature models [16] in reproducing kernel Hilbert spaces, which have the same form (1.1) but differ from \mathcal{F}_1 functions in the way the parameters $\{a_i\}_{i=1}^n$ and $\{b_i, c_i\}_{i=1}^n$ are trained. While in the case of \mathcal{F}_1 functions all parameters are trainable (which is sometimes referred to as active training), in random feature models the parameters $\{b_i, c_i\}_{i=1}^n$ are fixed (sampled from a given distribution over the latent space) and only the coefficients $\{a_i\}_{i=1}^n$ are trained (this is sometimes referred to as lazy training). Features can have a more general form than in (1.1). Corresponding function spaces were called \mathcal{F}_2 spaces in [9]. While \mathcal{F}_2 functions are significantly easier to train, their approximation properties are inferior to those of \mathcal{F}_1 functions [9, 17].

So far we have discussed neural network approximations of scalar-valued functions $f: \mathbb{R}^d \rightarrow \mathbb{R}$ with finite-dimensional inputs. However, neural networks are now being employed increasingly often in inherently infinite-dimensional settings such as inverse problems [18], mathematical imaging [19] and partial differential equations (PDEs) [20], where both the input and the output spaces may be infinite-dimensional. Whilst discretising these problems and subsequently using learning algorithms in the finite-dimensional setting is possible, there are advantages in designing learning algorithms directly in infinite dimensions such as discretisation- or mesh-invariance [21].

A first universal approximation theorem for neural networks acting between Banach spaces has been obtained in [22]. The authors show that any continuous nonlinear function mapping a compact set in a Banach space to a compact set in another one can be approximated arbitrary well with a pair of two-layer neural networks. More recently [23], this construction has been extended to deep (i.e. with more than two layers) neural networks and was termed *DeepONet*. In [24], the authors interpret the DeepONet as an encoder-decoder type of network and study its approximation properties in the case when the input space is a Hilbert space. A similar approach was taken in [25] to derive a different network architecture. Other existing approaches include Fourier Neural Operators [26] and graph kernel networks [27]. There is also a large body of literature on kernel methods in Hilbert [28] or Banach spaces [29]. Random feature models with Banach-space valued features are being used in PDEs to learn mappings between Banach or Hilbert spaces [21]. We refer the reader to the review [30] for details on vector-valued kernel methods.

Our main goal in this paper is to obtain quantitative estimates (i.e. convergence rates) for the

approximation of nonlinear operators with two-layer networks by extending the results of [9–11] from the finite-dimensional and scalar-valued case $f: \mathbb{R}^d \rightarrow \mathbb{R}$ to the vector-valued and infinite-dimensional case $f: \mathcal{X} \rightarrow \mathcal{Y}$, where \mathcal{X} and \mathcal{Y} are Banach spaces. An important difference from the approach taken in [24] is that our convergence rates are obtained under a regularity assumption on the underlying function f , while [24] relies on the properties of the measure (more precisely, its push-forward by f) on the input space \mathcal{X} that generates the training samples and underlies the Lebesgue-Bochner space in which the approximation error is studied.

Considering infinite-dimensional inputs is natural if one is interested in dimension-independent approximation rates, and indeed we show that standard Monte-Carlo rates can be obtained in this setting. The situation with vector-valued outputs is more difficult. Whilst going from functions with values in $\mathcal{Y} = \mathbb{R}$ to those with values in $\mathcal{Y} = \mathbb{R}^k$ is trivial, a generalisation to infinite-dimensional output spaces \mathcal{Y} is far from being so.

The paper is organised as follows. In Section 2 we fix our function-analytic setting and recall necessary concepts from the theory of partially ordered spaces (Riesz spaces, also known as vector lattices [31]). We then rewrite (1.1) in the appropriate abstract form using a generalised ReLU activation function – the operation of taking the positive part of an element in a Riesz space, cf. (2.1). Our results can be immediately transferred to the so-called *leaky ReLU* activation, which in the language of Riesz spaces is a combination of the positive and negative parts of its argument. More generally, our results apply to any positively homogeneous and weakly-* continuous activation.

Section 3 contains our analysis of vector-valued \mathcal{F}_1 spaces. We recall known finite-dimensional results from [9–11] in Section 3.1 before proceeding to the infinite-dimensional setting in Section 3.2. In Section 3.2.1 we introduce vector-valued \mathcal{F}_1 spaces and prove their embeddings into appropriate Lipschitz and Bochner spaces. It is important to note that, at least with current techniques, we are only able to show continuity of vector-valued \mathcal{F}_1 functions with respect to the weak-* topology in the output space \mathcal{Y} ; embeddings into Lipschitz and Bochner spaces are with respect to a particular metrisation of this topology. In Section 3.2.2 we prove direct and inverse approximation theorems and obtain Monte-Carlo $O(1/\sqrt{n})$ approximation rates, where n is the number of summands (neurons) in (1.1). In Section 3.2.3 we give examples of output spaces \mathcal{Y} for which these results hold. These are spaces where the non-linearity $\sigma: \mathcal{Y} \rightarrow \mathcal{Y}$ is weakly-* continuous. In the case of the ReLU function this turns out to be a significant restriction: among the common spaces only sequence spaces ℓ^p , $1 < p \leq \infty$, and Lipschitz spaces seem to have this property.

In Section 4 we turn our attention to training vector-valued \mathcal{F}_1 functions. It is well known (e.g., [32, 33]) that training a neural network from noisy observations is an ill-posed inverse problem and therefore regularisation is required. It may be achieved implicitly by introducing stochasticity into gradient descent dynamics and/or early stopping [14, 34–36] or explicitly by combining the empirical loss with an appropriate norm penalty [9]. The choice of this norm determines the model class in the language of approximation theory [37].

It has also been shown [38] that training sufficiently wide two-layer networks with weight decay [39] is equivalent to solving a regularised variational problem with an \mathcal{F}_1 norm penalty term. This problem boils down to the reconstruction of a Radon measure over a compact metric space from a finite number of point observations (pairings with continuous functions). It is a well-studied problem. In the context of deterministic regularisation, it has been studied in [40]. Support localisation and sparse spikes recovery under certain separation conditions have been studied in [41, 42]. In these works, the data term is considered fixed, which in our setting means that the number of data points is fixed, too. We are interested in the regime when the number of data points goes to infinity. We will also assume in our setting that the noise in the training data goes to zero, which is common in the inverse problems literature, but different from the statistical learning setting. The reason for this is that the problem of reconstructing a measure is ill-posed and therefore convergence of statistical estimators can be arbitrary slow (e.g., [43]) unless regularisation is used. For spectral regularisation, convergence rates in the statistical setting were obtained in [44]. Rates for variational regularisation with convex p -homogeneous functionals, $p > 1$, were obtained in a recent paper [45]. Our case of a 1-homogeneous functional (the Radon norm of a measure) is not covered by existing literature and studying it is beyond the scope of this paper.

Combining standard machinery of variational regularisation such as source conditions and Bregman distances [46] with Monte-Carlo integration, we obtain convergence rates in Bregman distance that hold with high probability and derive an optimal choice rule of the regularisation parameter as a function of the noise in the data and the sample size.

Main contributions. Below we summarise our main contributions.

- We extend the definition of \mathcal{F}_1 functions with the ReLU activation function to the vector-valued case and show well-posedness of this definition if the output space \mathcal{Y} is a Riesz space (vector lattice) where lattice operations are weakly-* sequentially continuous;
- we prove embeddings of the \mathcal{F}_1 space into Lipschitz and Bochner spaces defined using an appropriate metrisation of the weak-* topology on \mathcal{Y} ;
- we prove direct and inverse approximation theorems for \mathcal{F}_1 functions and obtain Monte-Carlo approximation rates;
- we analyse variational regularisation in the context of finding an optimal representation of an \mathcal{F}_1 function via an integral over a compact latent space, provide an interpretation of the source condition in terms of the support of the measure that realises the \mathcal{F}_1 function and obtain convergence rates in Bregman distance (with high probability) in the regime of vanishing noise and growing amount of data.

Notation. Throughout the paper, \mathcal{X} will be the input space of the neural network and \mathcal{Y} its image space. Both will be assumed to have separable preduals \mathcal{X}° and \mathcal{Y}° , respectively, and \mathcal{Y} will be assumed to have a partial order (i.e., \mathcal{Y} is a Riesz space). The latent space will be denoted by \mathcal{Z} . The space of linear operators $\mathcal{X} \rightarrow \mathcal{Y}$ will be denoted by $\mathcal{L}(\mathcal{X}; \mathcal{Y})$ and the space of nuclear operators by $\mathcal{N}(\mathcal{X}; \mathcal{Y})$. The unit ball in \mathcal{Y} will be denoted by $\mathbb{B}_{\mathcal{Y}}$, and ditto for other spaces. A metric that metrises the weak-* topology on the unit ball in \mathcal{Y} will be denoted by d_* and (\mathcal{Y}, d_*) will stand for the space \mathcal{Y} equipped with this weak-* metric, ditto for other spaces when appropriate. This space is not complete, but we will only work with bounded sets in \mathcal{Y} , so lack of completeness will not cause any issues.

2 Linear-nonlinear decompositions in Riesz spaces

Let \mathcal{X}, \mathcal{Y} be Banach spaces. We are interested in approximating continuous functions $f: \mathcal{X} \rightarrow \mathcal{Y}$ by two-layer neural networks. Let \mathcal{Z} be a *latent space*, to be specified later, and let $B: \mathcal{X} \rightarrow \mathcal{Z}$ and $A: \mathcal{Z} \rightarrow \mathcal{Y}$ be linear bounded operators and $c \in \mathcal{Z}$. We will consider approximations of the following form, cf. (1.1)

$$f_{A,B,c}(x) := A\sigma(Bx + c), \quad (2.1)$$

where $\sigma: \mathcal{Z} \rightarrow \mathcal{Z}$ is a non-linear function. This representation of neural networks is often referred to as linear-nonlinear decomposition. We will consider the ReLU activation function, which in finite-dimensional spaces is defined as follows

$$\text{ReLU}(z) := (z_+^1, \dots, z_+^m)^T = z_+, \quad z \in \mathbb{R}^m, \quad (2.2)$$

i.e. it takes the componentwise positive part of a vector $z \in \mathbb{R}^m$. This operation makes sense in a much more general setting of Riesz spaces [31] (also known as vector lattices).

Definition 2.1 (Riesz space). A vector space E equipped with a partial order “ \leq ” is called a Riesz space if for any $x, y, z \in E$ and any scalar $\lambda \geq 0$

- $x \leq y \implies x + z \leq y + z$;
- $x \leq y \implies \lambda x \leq \lambda y$;
- a supremum $x \vee y \in E$ exists, i.e. an element such that $x \vee y \geq x$, $x \vee y \geq y$ and

$$\forall z \text{ s.t. } z \geq x, z \geq y \implies z \geq x \vee y.$$

In a Riesz space, the positive and the negative parts of an element $x \in E$ and its absolute value are defined as follows

$$x_+ := x \vee 0, \quad x_- := (-x)_+, \quad |x| := x_+ + x_-. \quad (2.3)$$

Clearly, one has

$$x = x_+ - x_-, \quad x \in E.$$

Definition 2.2 (Banach lattice). Let E be a Riesz space equipped with a norm with the following monotonicity property

$$|x| \leq |y| \implies \|x\| \leq \|y\| \quad \forall x, y \in E. \quad (2.4)$$

If, in addition, E is norm complete, it is called a Banach lattice.

We will only come across Banach lattices in Section 3.2.3. In the rest of the paper, the monotonicity property (2.4) will not be needed. For a more detailed introduction to Riesz spaces we refer to [31].

Example 2.3. Here we give some common examples of Banach lattices and an example of a Riesz space that is *not* a Banach lattice.

- finite-dimensional spaces \mathbb{R}^m and sequence spaces ℓ^p are Banach lattices under the following componentwise ordering

$$x \geq y \iff x_i \geq y_i \quad \forall i;$$

- the space of continuous functions $C(\Omega)$ (Ω – compact) is a Banach lattice under the following pointwise ordering

$$f \geq g \iff f(t) \geq g(t) \quad \forall t \in \Omega;$$

- the spaces $L^p(\Omega)$, $1 \leq p \leq \infty$, are Banach lattices under the following ordering

$$f \geq g \iff f(t) \geq g(t) \quad \text{a.e. in } \Omega;$$

- the space of signed Radon measures $\mathcal{M}(\Omega)$ is a Banach lattice under the following ordering

$$\mu \geq \nu \iff \mu(A) \geq \nu(A) \quad \forall A \subset \Omega.$$

The positive and the negative parts of a measure are given by the Hahn-Jordan decomposition;

- let X, Y be Banach lattices. A linear operator $A: X \rightarrow Y$ is called *positive*, $A \geq 0$, if for all $x \geq 0$ one has $Ax \geq 0$ (this is not to be confused with positivity in the sense of positive semidefiniteness). A linear operator that can be written as a difference of two positive operators is called *regular*. Regular operators are always bounded. The subspace of regular operators $L^r(X; Y)$ is a Banach lattice under the following partial order

$$A \geq B \iff Ax \geq Bx \quad \forall x \geq 0;$$

- the space of Lipschitz functions $\text{Lip}(\Omega)$ with the pointwise ordering

$$f \geq g \iff f(t) \geq g(t) \quad \forall t \in \Omega$$

is a Riesz space but *not* a Banach lattice, i.e. the monotonicity property (2.4) fails [47].

Throughout the paper, we will assume that the latent space \mathcal{Z} is a Riesz space and will understand the ReLU function as the (nonlinear) operation of taking the positive part

$$\text{ReLU}(z) := z_+, \quad z \in \mathcal{Z}. \quad (2.5)$$

Sometimes we will keep the notation $\sigma(\cdot)$ to denote the nonlinearity in (2.1). Our results also apply verbatim to the *leaky ReLU* activation

$$\text{leaky ReLU}(z) := z_+ - \lambda z_-, \quad \lambda \in (0, 1), \quad z \in \mathcal{Z}. \quad (2.6)$$

Since σ is positively one-homogeneous, (2.1) is invariant under the transformation $(A, B, c) \rightarrow (A/t, tB, tc)$, hence without loss of generality we can assume that

$$\|B\|_{\mathcal{X} \rightarrow \mathcal{Z}} \leq 1, \quad \|c\|_{\mathcal{Z}} \leq 1. \quad (2.7)$$

Remark 2.4 (Compact notation). We will use the following compact notation (cf., e.g., [11]). Let us augment the she space \mathcal{X} with \mathbb{R} , i.e. consider the Cartesian product $\mathcal{X} \times \mathbb{R}$. The pair (B, c) will be understood as a linear operator $\mathcal{X} \times \mathbb{R} \rightarrow \mathcal{Z}$ acting as $(x, \alpha) \mapsto Bx + \alpha c$. Abusing notation, we will refer to the cartesian product $\mathcal{X} \times \mathbb{R}$ as just \mathcal{X} , to (x, α) as just x and to (B, c) as B . Then (2.1) becomes

$$f_{A,B}(x) := A\sigma(Bx). \quad (2.8)$$

For inputs of the form $(x, 1)$ this is the same as (2.1).

Remark 2.5. If $\mathcal{Y} = \mathbb{R}$, then the operator A from (2.8) is acting $A: \mathcal{Z} \rightarrow \mathbb{R}$ and hence it can be identified with an element of the dual space $a \in \mathcal{Z}^*$. Equation (2.8) can then be rewritten in terms of a dual pairing

$$f_{a,B}(x) := \langle \sigma(Bx), a \rangle. \quad (2.9)$$

Example 2.6 (Countably wide two-layer neural networks [11]). Let $\mathcal{X} = \mathbb{R}^d$ and $\mathcal{Y} = \mathbb{R}$. Choose $\mathcal{Z} = \ell^\infty$ and $B: \mathbb{R}^d \rightarrow \ell^\infty$ that maps $x \mapsto (b_i^T x)_{i=1}^\infty$. The requirement $\|B\|_{\mathbb{R}^d \rightarrow \ell^\infty} \leq 1$ means that we need $\|b_i\| \leq 1$ for all i . Hence, the b_i 's lie in the unit ball of the dual space of \mathcal{X} , which in this case coincides with \mathcal{X} . In this way, we obtain the space of all countably wide two-layer neural networks

$$\hat{\mathcal{F}}_\infty := \left\{ \sum_{i=1}^\infty a_i (\sigma(Bx))_i : \sum_{i=1}^\infty |a_i| < \infty \right\} \quad (2.10)$$

with the following norm

$$\|f\|_{\hat{\mathcal{F}}_\infty} := \inf_{a,B} \{ \|a\|_{\ell^1} : f \equiv f_{a,B} \} = \inf_{a,B} \left\{ \sum_{i=1}^\infty |a_i| : f \equiv f_{a,B} \right\}. \quad (2.11)$$

Here we implicitly restrict a to be an element of the predual space ℓ^1 rather than the dual $(\ell^\infty)^*$. A generalisation to the case $\mathcal{Y} = \mathbb{R}^k$ is easily obtained by taking \mathcal{Z} to be a vector-valued sequence space $\mathcal{Z} = \ell^\infty(\mathbb{R}^k)$.

3 Variation norm spaces

3.1 The finite-dimensional case

Variation norm spaces \mathcal{F}_1 [9] (also termed Barron spaces in [10, 11]) consist of functions that can be written in the following form

$$f(x) = \int_{\mathcal{V}} \varphi_v(x) d\mu(v),$$

where φ_v are *features* parametrised by the elements of a compact topological space \mathcal{V} and μ is a signed Radon measure on \mathcal{V} . If $\mathcal{X} = \mathbb{R}^d$ and $\mathcal{Y} = \mathbb{R}$, the latent space \mathcal{V} is typically chosen as the unit sphere (or unit ball) in \mathbb{R}^d .

For features of the form $\varphi_v(x) = \sigma(\langle v, x \rangle)$ with a positively one-homogeneous activation σ , variation norm functions can be written in terms of the linear-nonlinear decomposition as shown in [11]. To this end, let us choose $\mathcal{Z} := \mathcal{C}(\mathbb{B}_{\mathbb{R}^d})$ to be the space of continuous functions on the unit ball $\mathbb{B}_{\mathbb{R}^d}$ and let

$$B: \mathbb{R}^d \rightarrow \mathcal{C}(\mathbb{B}_{\mathbb{R}^d}), \quad x \mapsto \ell_x(\cdot), \quad \text{where } \ell_x(b) := b^T x, \quad b \in \mathbb{B}_{\mathbb{R}^d}. \quad (3.1)$$

Since $\|b\|_{\mathbb{R}^d} \leq 1$, the requirement $\|B\|_{\mathcal{X} \rightarrow \mathcal{Z}} \leq 1$ is clearly met.

Remark 3.1. If $\mathcal{X} = \mathbb{R}^d$ is equipped with a p -norm with $p \neq 2$, the latent space \mathcal{Z} is defined over the unit ball in the dual space of \mathcal{X} , i.e. \mathbb{R}^d with the dual q -norm (q is the Hölder conjugate of p). Hence, the latent space can be written as the space of continuous function over the unit ball in the dual space of \mathcal{X}

$$\mathcal{Z} = \mathcal{C}(\mathbb{B}_{\mathcal{X}^*}).$$

If \mathcal{X} is infinite-dimensional, the topology on $\mathbb{B}_{\mathcal{X}^*}$ needs to be specified. In order to preserve compactness, weak-* topology on $\mathbb{B}_{\mathcal{X}^*}$ is the natural choice. We will come back to this in Section 3.2.

According to Remark 2.5, the operator A can be identified with a Radon measure $a \in C(\mathbb{B}_{\mathbb{R}^d})^* = \mathcal{M}(\mathbb{B}_{\mathbb{R}^d})$. The expression (2.9) then reads as follows

$$f_a(x) = \int_{\mathbb{B}_{\mathbb{R}^d}} \sigma(Bx) da. \quad (3.2)$$

We note the following simple fact.

Proposition 3.2. *A function f_a as defined in (3.2) is Lipschitz with constant $\|a\|_{\mathcal{M}}$, i.e. for any $x, x' \in \mathcal{X}$ we have*

$$|f_a(x) - f_a(x')| \leq \|a\|_{\mathcal{M}} \|x - x'\|_{\mathcal{X}}.$$

Proof. This follows from the fact that the operation $\sigma(z) = z_+$ is 1-Lipschitz in $\mathcal{C}(\mathbb{B}_{\mathbb{R}^d})$. \square

Since $f_a(0) = 0$, we conclude that $f_a \in \text{Lip}_0(\mathcal{X})$, the space of Lipschitz functions on \mathcal{X} that vanish at zero, and

$$\|f_a\|_{\text{Lip}_0} \leq \|a\|_{\mathcal{M}}.$$

(Recall that, according to Remark 2.4, \mathcal{X} is a shorthand notation for $\mathcal{X} \times \mathbb{R}$ and neural networks correspond to inputs of the form $(x, 1)$ and they will not vanish at a common base point. However, this does not prevent us from viewing the function (3.2) as a Lipschitz function with base point at zero in $\mathcal{X} \times \mathbb{R}$.)

It is known that not all Lipschitz functions can be written in the form (3.2). The space of functions that can be written in this form was termed the variation norm space in [9] and the Barron space in [10].

Definition 3.3 (Scalar-valued \mathcal{F}_1 functions). Let $\mathcal{X} = \mathbb{R}^d$. The space of (scalar-valued) \mathcal{F}_1 functions is defined as follows

$$\mathcal{F}_1(\mathcal{X}) := \{f \in \text{Lip}_0(\mathcal{X}) : \|f\|_{\mathcal{F}_1} < \infty\}, \quad (3.3)$$

where

$$\|f\|_{\mathcal{F}_1} := \inf_{a \in \mathcal{M}(\mathbb{B}_{\mathbb{R}^d})} \{\|a\|_{\mathcal{M}} : f(x) = f_a(x), \quad x \in \mathcal{X}\}. \quad (3.4)$$

Remark 3.4. Since Lipschitz functions grow at most linearly at infinity, \mathcal{F}_1 functions lie in the Lebesgue space $L^1_\pi(\mathcal{X})$ if π is a probability measure over \mathcal{X} with a finite first moment, cf. [11].

3.2 The infinite-dimensional case

In this section we study the vector-valued case when both X and \mathcal{Y} are (in general, infinite-dimensional) Banach spaces. As noted in Remark 3.1, if the input space \mathcal{X} is infinite-dimensional, the latent space $\mathcal{Z} = \mathcal{C}(\mathbb{B}_{\mathcal{X}^*})$ needs to be considered over the unit ball $\mathbb{B}_{\mathcal{X}^*}$ equipped with the weak-* topology in order to preserve the duality with Radon measures. In the vector-valued case $\mathcal{Y} \neq \mathbb{R}$, as we shall see, the latent space \mathcal{Z} needs to be chosen as the space of \mathcal{Y} -valued continuous functions over the unit ball in the spaces of linear bounded operators $\mathcal{L}(\mathcal{X}; \mathcal{Y})$. Understanding the dualities that are involved in this case requires results from vector-valued measures [48] and tensor products of Banach spaces [49].

Our setting is as follows. We will assume that \mathcal{X} and \mathcal{Y} are duals of separable Banach spaces \mathcal{X}° and \mathcal{Y}° (more precisely, we will assume that \mathcal{X}° and \mathcal{Y}° possess Schauder bases). We will also assume that \mathcal{Y} has a lattice structure, i.e. that it is also a Riesz space. Although we assume \mathcal{Y} to be both a Banach space and a Riesz space, we do not assume that it is a Banach lattice, i.e. the monotonicity property (2.4) is not required.

3.2.1 Definitions

Denote by $\mathcal{L}(\mathcal{X}; \mathcal{Y})$ the space of linear bounded operators $\mathcal{X} \rightarrow \mathcal{Y}$ and by $\mathcal{N}(\mathcal{X}^\circ; \mathcal{Y}^\circ)$ the space of nuclear operators $\mathcal{X}^\circ \rightarrow \mathcal{Y}^\circ$ defined as follows [49]

Definition 3.5 (Nuclear operators). Let E and F be Banach spaces. An operator $N : E \rightarrow F$ is called nuclear if it can be written in the following form

$$Nx = \sum_{i=1}^{\infty} \langle x, \varphi_i \rangle y_i \quad (3.5)$$

for some sequences $\{\varphi_i\}_{i \in \mathbb{N}} \subset E^*$ and $\{y_i\}_{i \in \mathbb{N}} \subset F$ such that $\sum_{i=1}^{\infty} \|\varphi_i\| \|y_i\| < \infty$. Such representation need not be unique. The nuclear norm of $N \in \mathcal{N}(E; F)$ is defined as follows

$$\|N\|_{\mathcal{N}} := \inf \left\{ \sum_{i=1}^{\infty} \|\varphi_i\| \|y_i\| : Nx = \sum_{i=1}^{\infty} \langle x, \varphi_i \rangle y_i \quad \forall x \right\}.$$

The sequences $\{\varphi_i\}_{i \in \mathbb{N}}$ and $\{y_i\}_{i \in \mathbb{N}}$ can also be chosen to be normalised, in which case the nuclear representation (3.5) becomes

$$Nx = \sum_{i=1}^{\infty} \lambda_i \langle x, \varphi_i \rangle y_i \quad (3.6)$$

with some coefficients $\{\lambda_i\}_{i \in \mathbb{N}} \subset \mathbb{R}$ such that $\sum_{i=1}^{\infty} |\lambda_i| < \infty$.

To proceed, we will need to establish compactness properties of the unit ball in $\mathcal{L}(\mathcal{X}; \mathcal{Y})$. We will need the following result.

Definition 3.6 (Approximation property). A Banach space E is said to have the approximation property if for any Banach space F , any operator $T: E \rightarrow F$, any compact set $M \subset E$ and any $\varepsilon > 0$ there exists a finite-rank operator $S: E \rightarrow F$ such that $\|Tx - Sx\|_F \leq \varepsilon$ for all $x \in M$.

Proposition 3.7. *Let \mathcal{X} and \mathcal{Y} be duals of separable Banach spaces \mathcal{X}° and \mathcal{Y}° which possess Schauder bases. Then $\mathcal{L}(\mathcal{X}; \mathcal{Y})$ is isometrically isomorphic to $(\mathcal{N}(\mathcal{X}^\circ; \mathcal{Y}^\circ))^*$. For any $K \in \mathcal{L}(\mathcal{X}; \mathcal{Y})$ and $N \in \mathcal{N}(\mathcal{X}^\circ; \mathcal{Y}^\circ)$ the duality pairing can be written as follows*

$$\langle N, K \rangle = \text{Tr}(KN^*) := \sum_{i=1}^{\infty} \langle \eta_i, Kx_i \rangle, \quad (3.7)$$

where $\{x_i, \eta_i\}_{i \in \mathbb{N}} \subset \mathcal{X} \times \mathcal{Y}^\circ$ is any nuclear representation (3.5) of $N \in \mathcal{N}(\mathcal{X}^\circ; \mathcal{Y}^\circ)$. Consequently, the unit ball $\mathbb{B}_{\mathcal{L}(\mathcal{X}; \mathcal{Y})}$ is weakly- $*$ compact and metrisable.

Proof. Any Banach space with a Schauder basis has the approximation property [49, Ex. 4.4]. Hence, we can use Propositions A.1-A.2 with $E = \mathcal{X}$ and $F = \mathcal{Y}$ to obtain the claim. \square

The weak- $*$ topology can be metrised by the following metric [50, Thm. V.5.1]. Let $\{N_i\}_{i=1}^{\infty}$ be a countable dense system in $\mathcal{N}(\mathcal{X}^\circ; \mathcal{Y}^\circ)$, which exists since $\mathcal{N}(\mathcal{X}^\circ; \mathcal{Y}^\circ)$ is separable. For any $K, L \in \mathbb{B}_{\mathcal{L}(\mathcal{X}; \mathcal{Y})}$, define

$$d(K, L) = \sum_{i=1}^{\infty} 2^{-i} \frac{|\langle N_i, K - L \rangle|}{1 + |\langle N_i, K - L \rangle|}.$$

If $\{N_i\}_{i=1}^{\infty}$ are normalised, we can use the following equivalent metric

$$d_*(K, L) := \sum_{i=1}^{\infty} 2^{-i} |\langle N_i, K - L \rangle|. \quad (3.8)$$

We will treat $\mathbb{B}_{\mathcal{L}(\mathcal{X}; \mathcal{Y})}$ as a compact metric space with the weak- $*$ metric (3.8) and sometimes use the notation $(\mathbb{B}_{\mathcal{L}(\mathcal{X}; \mathcal{Y})}, d_*)$ to emphasise it.

We make the following simple observation.

Proposition 3.8. *Let $x \in \mathcal{X}$ and $\eta \in \mathcal{Y}^\circ$ be arbitrary but fixed. Define the following rank-one operator $N_x^\eta: \mathcal{X}^\circ \rightarrow \mathcal{Y}^\circ$*

$$N_x^\eta(\xi) := \langle \xi, x \rangle \eta, \quad \xi \in \mathcal{X}^\circ; \quad \|N_x^\eta\| = \|x\| \|\eta\|, \quad (3.9)$$

Then for any $K \in \mathbb{B}_{\mathcal{L}(\mathcal{X}; \mathcal{Y})}$ we have

$$\langle N_x^\eta, K \rangle = \langle \eta, Kx \rangle.$$

Proof. Since the nuclear representation of N_x^η consists of only one term (3.9), the claim follows directly from (3.7). \square

Let $\{\xi_j\}_{j \in \mathbb{N}}$ and $\{\eta_i\}_{i \in \mathbb{N}}$ be Schauder bases in \mathcal{X}° and \mathcal{Y}° , respectively, such that $\|\xi_j\| = \|\eta_i\| = 1$ for all $i, j \in \mathbb{N}$. Denote the corresponding coefficient functionals by $\{x_j\}_{j \in \mathbb{N}}$ and $\{y_i\}_{i \in \mathbb{N}}$. We get the following density result.

Proposition 3.9. *The system of operators $\{N_{x_j}^{\eta_i}\}_{i, j \in \mathbb{N}}$ (cf. (3.9)) is dense in $\mathcal{N}(\mathcal{X}^\circ; \mathcal{Y}^\circ)$.*

Remark 3.10. Famously, there exist separable spaces that do not possess a basis [51], however, most known examples are rather exotic.

We will consider the following metric on \mathcal{Y} defined analogously to (3.8)

$$d_*(y_1, y_2) := \sum_{i=1}^{\infty} 2^{-i} |\langle \eta_i, y_1 - y_2 \rangle|. \quad (3.10)$$

By the Banach-Alaoglu theorem, the unit ball $\mathbb{B}_{\mathcal{Y}}$ is compact with respect to this metric. We will denote by (\mathcal{Y}, d_*) the space \mathcal{Y} equipped with this metric. It can be turned into a normed space by defining

$$\|y\|_{(\mathcal{Y}, d_*)} := d_*(y, 0). \quad (3.11)$$

It is easy to verify that this functional is indeed a norm, i.e. it is absolutely one-homogeneous, satisfies the triangle inequality and $\|y\|_{(\mathcal{Y}, d_*)} = 0 \iff y = 0$. The space (\mathcal{Y}, d_*) is not complete, but we will work with strongly bounded (i.e. in the norm of \mathcal{Y}) sets, so this incompleteness will not present a problem. The main point is that \mathcal{Y} and (\mathcal{Y}, d_*) have different metrics.

Since \mathcal{Y} is a Riesz space, the norm in (\mathcal{Y}, d_*) can be equivalently expressed as follows

$$\|y\|_{(\mathcal{Y}, d_*)} = d_*(y, 0) = \sum_{i=1}^{\infty} 2^{-i} |\langle \eta_i, y \rangle| = \sum_{i=1}^{\infty} 2^{-i} |\langle \eta_i, y_+ - y_- \rangle| = d_*(y_+, y_-). \quad (3.12)$$

A similar construction can be used to define a norm on the completion of the space of zero-mean Radon measures with respect to the Kantorovich-Rubinstein distance [47, 52].

In the previous section we defined the space \mathcal{Z} as the space of continuous functions on the unit ball of the dual space of \mathcal{X} , cf. Remark 3.1. A natural generalisation to the vector-valued case is to take \mathcal{Z} to be the space of \mathcal{Y} -valued functions on $\mathbb{B}_{\mathcal{L}(\mathcal{X}; \mathcal{Y})}$ that are continuous with respect to the weak-* topologies in $\mathbb{B}_{\mathcal{L}(\mathcal{X}; \mathcal{Y})}$ and \mathcal{Y} . Hence, we take $\mathcal{Z} = \mathcal{C}(\mathbb{B}_{\mathcal{L}(\mathcal{X}; \mathcal{Y})}; (\mathcal{Y}, d_*))$, where $\mathbb{B}_{\mathcal{L}(\mathcal{X}; \mathcal{Y})}$ is considered as a compact metric space with the weak-* metric (3.8), and equip it with the following norm

$$\mathcal{Z} := \mathcal{C}(\mathbb{B}_{\mathcal{L}(\mathcal{X}; \mathcal{Y})}; (\mathcal{Y}, d_*)), \quad \|f\|_{\mathcal{Z}} := \sup_{K \in \mathbb{B}_{\mathcal{L}(\mathcal{X}; \mathcal{Y})}} \|f(K)\|_{(\mathcal{Y}, d_*)}, \quad (3.13)$$

where $\|\cdot\|_{(\mathcal{Y}, d_*)}$ is as defined in (3.11).

The space $\mathcal{C}(\mathbb{B}_{\mathcal{L}(\mathcal{X}; \mathcal{Y})}; (\mathcal{Y}, d_*))$ naturally inherits the lattice structure from \mathcal{Y} . To this end, we define the following partial order on $\mathcal{C}(\mathbb{B}_{\mathcal{L}(\mathcal{X}; \mathcal{Y})}; (\mathcal{Y}, d_*))$

$$\varphi \geq \psi \iff \varphi(K) \geq_{\mathcal{Y}} \psi(K) \quad \forall K \in \mathbb{B}_{\mathcal{L}(\mathcal{X}; \mathcal{Y})}, \quad \varphi, \psi \in \mathcal{C}(\mathbb{B}_{\mathcal{L}(\mathcal{X}; \mathcal{Y})}; (\mathcal{Y}, d_*)), \quad (3.14)$$

where $\leq_{\mathcal{Y}}$ denotes the partial order in \mathcal{Y} . It can be verified that $\mathcal{C}(\mathbb{B}_{\mathcal{L}(\mathcal{X}; \mathcal{Y})}; (\mathcal{Y}, d_*))$ is a Riesz space. However, it is not a Banach lattice with the supremum norm (3.13). The reason is that the \mathcal{Y}^* -norm (3.11) lacks the monotonicity property (2.4).

Now we define the operator B analogously to Section 3.1

$$B: \mathcal{X} \rightarrow \mathcal{C}(\mathbb{B}_{\mathcal{L}(\mathcal{X}; \mathcal{Y})}; (\mathcal{Y}, d_*)), \quad x \mapsto \mathcal{L}_x(\cdot), \quad \mathcal{L}_x(K) := Kx \quad \forall K \in \mathbb{B}_{\mathcal{L}(\mathcal{X}; \mathcal{Y})}. \quad (3.15)$$

The function \mathcal{L}_x maps $\mathbb{B}_{\mathcal{L}(\mathcal{X}; \mathcal{Y})}$ into the ball $\mathbb{B}_{\mathcal{Y}, \|x\|}$ of radius $\|x\|$. We need to check that \mathcal{L}_x is continuous, i.e. that the operator B is well-defined.

Proposition 3.11. *The function $\mathcal{L}_x(\cdot)$ defined in (3.15) is sequentially continuous with respect to the weak-* topologies on $\mathbb{B}_{\mathcal{L}(\mathcal{X}; \mathcal{Y})}$ and \mathcal{Y} for any fixed $x \in \mathcal{X}$. Consequently, $\mathcal{L}_x(\cdot) \in \mathcal{C}(\mathbb{B}_{\mathcal{L}(\mathcal{X}; \mathcal{Y})}; (\mathcal{Y}, d_*))$.*

Proof. Let $\{K_n\}_{n \in \mathbb{N}} \subset \mathbb{B}_{\mathcal{L}(\mathcal{X}; \mathcal{Y})}$ be an arbitrary sequence such that $K_n \rightharpoonup^* K$ for some $K \in \mathbb{B}_{\mathcal{L}(\mathcal{X}; \mathcal{Y})}$. Fix $x \in \mathcal{X}$ and let $\eta \in \mathcal{Y}^\circ$ be arbitrary. Let N_x^η be the rank-one operator defined in (3.9). Then we have

$$\langle N_x^\eta, K_n \rangle \rightarrow \langle N_x^\eta, K \rangle \quad \text{as } n \rightarrow \infty$$

and therefore

$$\langle \eta, K_n x \rangle \rightarrow \langle \eta, Kx \rangle \quad \text{as } n \rightarrow \infty$$

by Proposition 3.8. But this means that

$$\langle \eta, \mathcal{L}_x(K_n) \rangle \rightarrow \langle \eta, \mathcal{L}_x(K) \rangle \quad \text{as } n \rightarrow \infty.$$

Since this holds for all $\eta \in \mathcal{Y}^\circ$, we have $\mathcal{L}_x(K_n) \rightharpoonup^* \mathcal{L}_x(K)$. □

Corollary 3.12. *We have that $\|B\|_{\mathcal{X} \rightarrow \mathcal{C}(\mathbb{B}_{\mathcal{L}(\mathcal{X}; \mathcal{Y})}; (\mathcal{Y}, d_*))} \leq 1$. Indeed, for any $x \in \mathcal{X}$ with $\|x\| \leq 1$ and any $K \in \mathbb{B}_{\mathcal{L}(\mathcal{X}; \mathcal{Y})}$,*

$$\|\mathcal{L}_x(K)\|_{\mathcal{Y}^*} = d_*(\mathcal{L}_x(K), 0) = \sum_{i=1}^{\infty} 2^{-i} |\langle \eta_i, Kx \rangle| \leq \sup_{i \in \mathbb{N}} |\langle \eta_i, Kx \rangle| \leq \|K\| \|x\| \leq 1,$$

which also holds for the supremum in $x \in \mathbb{B}_{\mathcal{X}}$ and $K \in \mathbb{B}_{\mathcal{L}(\mathcal{X}; \mathcal{Y})}$.

Now we need to make sure that the nonlinearity σ does not take us outside the space of weakly-* continuous functions, i.e. that $\sigma(\mathcal{L}_x(\cdot)) \in \mathcal{C}(\mathbb{B}_{\mathcal{L}(\mathcal{X};\mathcal{Y})}; (\mathcal{Y}, d_*)$).

Proposition 3.13. *Suppose that lattice operations in \mathcal{Y} are weakly-* sequentially continuous. Then the function $\sigma(\mathcal{L}_x(\cdot))$ is sequentially continuous with respect to weak-* topologies on $\mathbb{B}_{\mathcal{L}(\mathcal{X};\mathcal{Y})}$ and \mathcal{Y} for any fixed $x \in \mathcal{X}$.*

Proof. We recall that for any $K \in \mathbb{B}_{\mathcal{L}(\mathcal{X};\mathcal{Y})}$, $\sigma(\mathcal{L}_x(K)) = (Kx)_+$, where the positive part is taken in \mathcal{Y} . Let $K_n \rightharpoonup^* K \in \mathbb{B}_{\mathcal{L}(\mathcal{X};\mathcal{Y})}$. Then $\mathcal{L}_x(K_n) \rightharpoonup^* \mathcal{L}_x(K)$ in \mathcal{Y} by Proposition 3.13. Since lattice operations in \mathcal{Y} are weakly-* sequentially continuous, we also get that $(\mathcal{L}_x(K_n))_+ \rightharpoonup^* (\mathcal{L}_x(K))_+$ in \mathcal{Y} . \square

Remark 3.14. It turns out that, apart from the trivial case $\dim(\mathcal{Y}) < \infty$, lattice operations are rarely sequentially continuous in the weak or weak-* topologies. For example, lattice operations are not weakly-* sequentially continuous in all $L_p([0,1])$ spaces for $1 < p \leq \infty$ (consider the sequence $f_n(x) := \sin(2\pi nx) \rightharpoonup^* 0$, while $\langle (f_n)_+, \mathbb{1} \rangle = \int_0^1 (\sin(2\pi nx))_+ dx = 1/\pi$ for all n). However, lattice operations are weakly-* sequentially continuous in sequence spaces ℓ^p , $1 < p \leq \infty$, and in spaces of Lipschitz functions (see Section 3.2.3 for details).

For readers' convenience, we provide here a proof of sequential weak-* continuity of lattice operations in ℓ^p spaces. In fact, lattice operations can even be 1-Lipschitz with respect to the d_* metric (3.10).

Proposition 3.15. *Let $\mathcal{Y} = \ell^p$, $1 < p \leq \infty$, and let the d_* metric (3.10) be defined using the standard basis $\{e_i\}_{i=1}^\infty$ as the countable dense system in the predual. Then the lattice operations are 1-Lipschitz with respect to the d_* metric.*

Proof. Let $\eta_n \rightharpoonup^* \eta$ on \mathbb{B}_{ℓ^p} . Taking the standard basis $\{e_i\}_{i=1}^\infty$ as the countable dense system in the predual ℓ^q (q is the conjugate exponent of p), we get from (3.10)

$$d_*((\eta_n)_+, \eta_+) = \sum_{i=1}^{\infty} 2^{-i} |(\eta_n^i)_+ - \eta_+^i| \leq \sum_{i=1}^{\infty} 2^{-i} |\eta_n^i - \eta^i| = d_*(\eta_n, \eta),$$

where η_n^i and η^i denote the i -th coordinate of η_n and η , respectively. The inequality above holds because lattice operations are 1-Lipschitz in \mathbb{R} . \square

We will identify the operator A with a measure $a \in \mathcal{M}(\mathbb{B}_{\mathcal{L}(\mathcal{X};\mathcal{Y})})$. For any $f \in \mathcal{C}(\mathbb{B}_{\mathcal{L}(\mathcal{X};\mathcal{Y})}; (\mathcal{Y}, d_*))$ and $a \in \mathcal{M}(\mathbb{B}_{\mathcal{L}(\mathcal{X};\mathcal{Y})})$, we will denote by $\langle f, a \rangle$ the \mathcal{Y} -valued pairing of f and a . More precisely, for any $K \in \mathbb{B}_{\mathcal{L}(\mathcal{X};\mathcal{Y})}$ and any η_i from the countable system $\{\eta_i\}_{i \in \mathbb{N}}$ (cf. Proposition 3.9), the function

$$K \mapsto \langle \eta_i, f(K) \rangle$$

is in $\mathcal{C}(\mathbb{B}_{\mathcal{L}(\mathcal{X};\mathcal{Y})}; \mathbb{R})$ and therefore the pairing

$$\langle \langle \eta_i, f(\cdot) \rangle, a \rangle = \int_{\mathbb{B}_{\mathcal{L}}} \langle \eta_i, f(K) \rangle da(K)$$

is well-defined for any $a \in \mathcal{M}(\mathbb{B}_{\mathcal{L}(\mathcal{X};\mathcal{Y})})$. In the above formula, $\mathbb{B}_{\mathcal{L}}$ stands for $\mathbb{B}_{\mathcal{L}(\mathcal{X};\mathcal{Y})}$. Using the linearity of dual pairings, we rewrite this as follows

$$\langle \langle \eta_i, f(\cdot) \rangle, a \rangle = \left\langle \eta_i, \int_{\mathbb{B}_{\mathcal{L}}} f(K) da(K) \right\rangle = \langle \eta_i, \langle f(\cdot), a \rangle \rangle.$$

We will hence write

$$\langle f, a \rangle := \int_{\mathbb{B}_{\mathcal{L}}} f(K) da(K), \quad \langle f, a \rangle \in (\mathcal{Y}, d_*), \quad (3.16)$$

understanding this as a weak-* limit. In fact, f can be understood as the density of some (\mathcal{Y}, d_*) -valued measure [48] on $\mathbb{B}_{\mathcal{L}(\mathcal{X};\mathcal{Y})}$ and $\langle f, a \rangle$ as its pairing with the constant $\mathbb{1}$ function. With this notation, we can generalise (2.9) to the vector-valued case

$$f_a(x) := \langle \sigma(Bx), a \rangle, \quad f_a: \mathcal{X} \rightarrow (\mathcal{Y}, d_*). \quad (3.17)$$

We have the following analogue of Proposition 3.2.

Theorem 3.16. *Let $a \in \mathcal{M}(\mathbb{B}_{\mathcal{L}(\mathcal{X};\mathcal{Y})})$ and suppose that lattice operations in \mathcal{Y} are 1-Lipschitz with respect to the d_* metric. Then the function f_a defined in (3.17) is Lipschitz with constant $\|a\|_{\mathcal{M}}$ with respect to the d_* metric in \mathcal{Y} , i.e. for any $x, x' \in \mathcal{X}$ we have*

$$d_*(f_a(x), f_a(x')) \leq \|a\|_{\mathcal{M}} \|x - x'\|.$$

Consequently, $f_a \in \text{Lip}_0(\mathcal{X}; (\mathcal{Y}, d_*))$, where $\text{Lip}_0(\mathcal{X}; (\mathcal{Y}, d_*))$ is the space of Lipschitz functions with respect to the d_* metric in \mathcal{Y} vanishing at zero.

Proof. Let $x, x' \in \mathcal{X}$ be fixed and let $\{\eta_i\}_{i=1}^{\infty}$ be as defined in Proposition 3.9. The d_* distance between $f_a(x)$ and $f_a(x')$ is given by

$$d_*(f_a(x), f_a(x')) = \sum_{i=1}^{\infty} 2^{-i} |\langle \eta_i, \langle \sigma(Bx) - \sigma(Bx'), a \rangle \rangle|.$$

Using the linearity of dual pairings and remembering that $\sigma(Bx)(K) = (Kx)_+$, we get

$$\begin{aligned} |\langle \eta_i, \langle \sigma(Bx) - \sigma(Bx'), a \rangle \rangle| &= \left| \left\langle \eta_i, \int_{\mathbb{B}_{\mathcal{L}}} ((Kx)_+ - (Kx')_+) da(K) \right\rangle \right| \\ &= \left| \int_{\mathbb{B}_{\mathcal{L}}} \langle \eta_i, (Kx)_+ - (Kx')_+ \rangle da(K) \right| \\ &\leq \int_{\mathbb{B}_{\mathcal{L}}} |\langle \eta_i, (Kx)_+ - (Kx')_+ \rangle| d|a|(K), \end{aligned}$$

where $|a| \in \mathcal{M}(\mathbb{B}_{\mathcal{L}(\mathcal{X};\mathcal{Y})})$ is the total variation of a . Now, using the above estimate as well as Fubini's theorem, we get that

$$\begin{aligned} d_*(f_a(x), f_a(x')) &\leq \sum_{i=1}^{\infty} 2^{-i} \int_{\mathbb{B}_{\mathcal{L}}} |\langle \eta_i, (Kx)_+ - (Kx')_+ \rangle| d|a|(K) \\ &= \int_{\mathbb{B}_{\mathcal{L}}} \sum_{i=1}^{\infty} 2^{-i} |\langle \eta_i, (Kx)_+ - (Kx')_+ \rangle| d|a|(K) \\ &= \int_{\mathbb{B}_{\mathcal{L}}} d_*((Kx)_+, (Kx')_+) d|a|(K) \\ &\leq \int_{\mathbb{B}_{\mathcal{L}}} d_*(Kx, Kx') d|a|(K), \end{aligned}$$

where we used the fact that lattice operations in \mathcal{Y} are 1-Lipschitz with respect to the d_* metric. Now, since $d_*(y, y') \leq \|y - y'\|$ for all $y, y' \in \mathcal{Y}$, we have

$$\begin{aligned} d_*(f_a(x), f_a(x')) &\leq \int_{\mathbb{B}_{\mathcal{L}}} \|K(x - x')\| d|a|(K) \\ &\leq \|x - x'\| \int_{\mathbb{B}_{\mathcal{L}}} d|a|(K) \\ &= \|a\|_{\mathcal{M}} \|x - x'\| \end{aligned}$$

as desired. Noting that $f_a(0) = 0$ completes the proof. \square

We also have the following continuity result with respect to the weak-* topology on \mathcal{X} .

Theorem 3.17. *Let $a \in \mathcal{M}(\mathbb{B}_{\mathcal{L}(\mathcal{X};\mathcal{Y})})$ and suppose that lattice operations in \mathcal{Y} are weakly-* sequentially continuous. Then the function f_a defined in (3.17) is sequentially continuous with respect to the weak-* topologies in \mathcal{X} and \mathcal{Y} , i.e. for any sequence $x_n \rightharpoonup^* x$ in \mathcal{X} we have*

$$f_a(x_n) \rightharpoonup^* f_a(x) \quad \text{in } \mathcal{Y}.$$

Proof. Let $x_n \rightharpoonup^* x$ and let $K \in \mathbb{B}_{\mathcal{L}(\mathcal{X};\mathcal{Y})}$ be weak-*–weak-* continuous (that is, K has to be the adjoint of a bounded operator $\mathcal{Y}^\circ \rightarrow \mathcal{X}^\circ$). Then we have

$$Kx_n \rightharpoonup^* Kx$$

and, since lattice operations in \mathcal{Y} are weakly-* sequentially continuous,

$$(Kx_n)_+ \rightharpoonup^* (Kx)_+. \quad (3.18)$$

In particular, this holds for finite-rank operators $\mathcal{X} \rightarrow \mathcal{Y}$. Now let $\eta \in \mathcal{Y}^\circ$ be arbitrary and consider

$$\langle \eta, f_a(x_n) \rangle = \int_{\mathbb{B}_{\mathcal{L}}} \langle \eta, (Kx_n)_+ \rangle da(K).$$

Since \mathcal{X}° has the approximation property, by [49, Prop. 4.6(iii)] finite-rank operators are weakly-* dense in $\mathcal{L}(\mathcal{X}, \mathcal{Y})$ and we can approximate the above integral with an integral over finite-rank operators, for which (3.18) holds. Therefore, we have

$$\begin{aligned} \lim_{n \rightarrow \infty} \langle \eta, f_a(x_n) \rangle &= \lim_{n \rightarrow \infty} \int_{\mathbb{B}_{\mathcal{L}}} \langle \eta, (Kx_n)_+ \rangle da(K) = \int_{\mathbb{B}_{\mathcal{L}}} \lim_{n \rightarrow \infty} \langle \eta, (Kx_n)_+ \rangle da(K) \\ &= \int_{\mathbb{B}_{\mathcal{L}}} \langle \eta, (Kx)_+ \rangle da(K) = \langle \eta, f_a(x) \rangle \end{aligned}$$

by the dominated convergence theorem. Since $\eta \in \mathcal{Y}^\circ$ was arbitrary, this implies the claim. \square

Now we are ready to extend Definition 3.3 to the infinite-dimensional case.

Definition 3.18 (Vector-valued \mathcal{F}_1 functions). Let \mathcal{X}, \mathcal{Y} be duals of separable Banach spaces \mathcal{X}° and \mathcal{Y}° that admit Schauder bases. Let \mathcal{Y} be such that lattice operations are 1-Lipschitz with respect to the d_* metric. Define the space of \mathcal{Y} -valued \mathcal{F}_1 functions as follows

$$\mathcal{F}_1(\mathcal{X}; \mathcal{Y}) := \{f \in \text{Lip}_0(\mathcal{X}; (\mathcal{Y}, d_*)) : \|f\|_{\mathcal{F}_1} < \infty\}, \quad (3.19)$$

where $\text{Lip}_0(\mathcal{X}; (\mathcal{Y}, d_*))$ is the space of Lipschitz functions with respect to the d_* metric in \mathcal{Y} that vanish at zero (cf. Proposition 3.16) and

$$\|f\|_{\mathcal{F}_1} := \inf_{a \in \mathcal{M}(\mathbb{B}_{\mathcal{L}(\mathcal{X}; \mathcal{Y})})} \{ \|a\|_{\mathcal{M}} : f_a(x) = f(x) \forall x \in \mathcal{X} \}. \quad (3.20)$$

We set $\|f\|_{\mathcal{F}_1} = \infty$ if there is no feasible point in (3.20).

Remark 3.19. We emphasise that the target space in this definition is (\mathcal{Y}, d_*) and not \mathcal{Y} . That is, \mathcal{F}_1 functions are Lipschitz only with respect to the weak-* metric (3.10) in \mathcal{Y} , cf. Theorem 3.16.

Remark 3.20. Let $\pi \in \mathcal{P}_p(\mathcal{X})$ be a probability measure over \mathcal{X} with p finite moments. Then the variation norm space $\mathcal{F}_1(\mathcal{X}; \mathcal{Y})$ is a subspace of the Bochner space $L_\pi^p(\mathcal{X}; (\mathcal{Y}, d_*))$. Indeed, since Lipschitz functions grow at most linearly at infinity, we get from Theorem 3.16 that for any $f = \langle \sigma(B\cdot), a \rangle \in \mathcal{F}_1(\mathcal{X}; \mathcal{Y})$

$$\|f\|_{L_\pi^p}^p = \int_{\mathcal{X}} \|f\|_{(\mathcal{Y}, d_*)}^p d\pi(x) \leq \|a\|_{\mathcal{M}}^p \int_{\mathcal{X}} \|x\|_{\mathcal{X}}^p d\pi(x) = m_p(\pi) \|f\|_{\mathcal{F}_1}^p < \infty,$$

where $m_p(\pi)$ is the p -th moment of π .

The next result shows that the infimum in (3.20) is indeed attained.

Proposition 3.21. *If (3.20) is feasible, it admits a minimiser.*

Proof. Let $\{a_k\}_{k=1}^\infty$ be a minimising sequence. Since it is bounded, it contains a weakly-* convergent subsequence (that we don't relabel)

$$a_k \rightharpoonup^* a \quad \text{in } \mathcal{M}(\mathbb{B}_{\mathcal{L}(\mathcal{X}; \mathcal{Y})}).$$

Let $\{\eta_i\}_{i=1}^\infty \subset \mathcal{Y}^\circ$ be as defined in Proposition 3.9. Fix an arbitrary $x \in \mathcal{X}$. Since a_k is feasible, we have that $f_{a_k}(x) = f(x)$ for all $x \in \mathcal{X}$ and for all $k \in \mathbb{N}$. With computations similar to those in Proposition 3.16, we get that for any $i \in \mathbb{N}$

$$\begin{aligned} \langle \eta_i, f(x) \rangle &= \langle \eta_i, f_{a_k}(x) \rangle = \langle \eta_i, \langle \sigma(Bx), a_k \rangle \rangle = \int_{\mathbb{B}_{\mathcal{L}(\mathcal{X}; \mathcal{Y})}} \langle \eta_i, (Kx)_+ \rangle da_k(K) \\ &\rightarrow \int_{\mathbb{B}_{\mathcal{L}(\mathcal{X}; \mathcal{Y})}} \langle \eta_i, (Kx)_+ \rangle da(K) = \langle \eta_i, f_a(x) \rangle, \end{aligned}$$

since the function $K \mapsto \langle \eta_i, (Kx)_+ \rangle$ is continuous on $\mathbb{B}_{\mathcal{L}(\mathcal{X};\mathcal{Y})}$ by Proposition 3.13. Hence, $\langle \eta_i, f_a(x) \rangle = \langle \eta_i, f(x) \rangle$ holds for all i and therefore $\langle \sigma(Bx), a \rangle = f(x)$. This holds for all x , hence the feasible set in (3.20) is weakly-* closed. Since the Radon norm is weakly-* lower semicontinuous, we get the claim by the direct method of calculus of variations. \square

3.2.2 Approximation theorems

Our goal in this section is to show that \mathcal{Y} -valued \mathcal{F}_1 functions can be well approximated with countably wide neural networks and that, in some sense, they are the largest class of functions with this property. Similar results in the scalar-valued ($\mathcal{Y} = \mathbb{R}$) and finite-dimensional ($\mathcal{X} = \mathbb{R}^d$) case can be found in [10, 11].

Theorem 3.22 (Inverse approximation). *Let \mathcal{X}, \mathcal{Y} be duals of separable Banach spaces \mathcal{X}^\diamond and \mathcal{Y}^\diamond that possess Schauder bases and let \mathcal{Y} be such that lattice operations are 1-Lipschitz with respect to the d_* metric (3.10). Let*

$$f_n(x) := \sum_{i=1}^n \alpha_i (K_i x)_+, \quad x \in \mathcal{X},$$

where K_i are finite-rank operators $\mathcal{X} \rightarrow \mathcal{Y}$ satisfying $\|K_i\| \leq 1$ and α_i rational numbers (both K_i and α_i depend on n), and suppose that

$$\sum_{i=1}^n |\alpha_i| \leq C \quad \text{uniformly in } n.$$

Then there exists $f \in \mathcal{F}_1(\mathcal{X}; \mathcal{Y})$ with $\|f\|_{\mathcal{F}_1} \leq C$ such that

(i) for any $x \in \mathcal{X}$

$$\sum_{i=1}^n \alpha_i (K_i x)_+ \rightarrow^* f(x) \quad \text{in } \mathcal{Y} \text{ as } n \rightarrow \infty,$$

the convergence being along a subsequence that does not depend on x ;

(ii) if $\mu \in \mathcal{P}_p(\mathcal{X})$ is a probability measure on \mathcal{X} with $p \geq 1$ finite moments then convergence in the Bochner space $L_\mu^p(\mathcal{X}; (\mathcal{Y}, d_*))$ also holds along the same subsequence

$$\left\| f(x) - \sum_{i=1}^n \alpha_i (K_i x)_+ \right\|_{L_\mu^p(\mathcal{X}; (\mathcal{Y}, d_*))} \rightarrow 0 \quad \text{as } n \rightarrow \infty.$$

Remark 3.23. Since α_i depend on n , it does not matter whether we use a normalisation $\frac{1}{n} \sum_{i=1}^n \alpha_i (K_i x)_+$ or not. In fact, this is nothing but a change of notation.

Proof of Theorem 3.22. (i) We rewrite f_n using Radon measures on the unit ball in $\mathcal{L}(\mathcal{X}; \mathcal{Y})$

$$f_n(x) = \left\langle \sigma(Bx), \sum_{i=1}^n a_i \right\rangle,$$

where $a_i := \alpha_i \delta_{K_i} \in \mathcal{M}(\mathbb{B}_{\mathcal{L}(\mathcal{X};\mathcal{Y})})$. Using the condition that $\sum_{i=1}^n |\alpha_i| \leq C$ uniformly in n , with the same arguments as in Proposition 3.21 we get that

$$\sum_{i=1}^n a_i \rightarrow^* a \quad \text{in } \mathcal{M}(\mathbb{B}_{\mathcal{L}(\mathcal{X};\mathcal{Y})})$$

along a subsequence and

$$\|a\|_{\mathcal{M}} \leq \liminf_{n \rightarrow \infty} \left\| \sum_{i=1}^n a_i \right\|_{\mathcal{M}} \leq \liminf_{n \rightarrow \infty} \sum_{i=1}^n |\alpha_i| \leq C.$$

Letting

$$f(x) := \langle \sigma(Bx), a \rangle, \quad x \in \mathcal{X},$$

we observe that $\|f\|_{\mathcal{F}_1} \leq \|a\|_{\mathcal{M}} \leq C$. Now let $\eta \in \mathcal{Y}^\circ$ be arbitrary. Then

$$\langle \eta, f(x) - f_n(x) \rangle = \left\langle \eta, \left\langle a - \sum_{i=1}^n a_i, \sigma(Bx) \right\rangle \right\rangle = \left\langle \langle \eta, (\cdot x)_+ \rangle, a - \sum_{i=1}^n a_i \right\rangle \rightarrow 0$$

as $n \rightarrow \infty$, since the function $K \mapsto \langle \eta, (Kx)_+ \rangle$ is weakly-* continuous on $\mathbb{B}_{\mathcal{L}(\mathcal{X}; \mathcal{Y})}$. Since $\eta \in \mathcal{Y}^\circ$ was arbitrary, we get that $f_n(x) \xrightarrow{*} f(x)$ on \mathcal{Y} .

(ii) Let $\eta \in \mathbb{B}_{\mathcal{Y}^\circ}$ be arbitrary. Then

$$\begin{aligned} |\langle \eta, f(x) - f_n(x) \rangle| &= |\langle \eta, \langle \sigma(Bx), a - a_n \rangle \rangle| = \left| \int_{\mathbb{B}_{\mathcal{L}}} \langle \eta, (Kx)_+ \rangle d(a - a_n) \right| \\ &\leq \left| \int_{\mathbb{B}_{\mathcal{L}}} |\langle \eta, (Kx)_+ \rangle| d|a - a_n| \right| \leq \left| \int_{\mathbb{B}_{\mathcal{L}}} \|\eta\| \|(Kx)_+\| d|a - a_n| \right| \\ &\leq \|x\| \left| \int_{\mathbb{B}_{\mathcal{L}}} \|K\| d|a - a_n| \right| \leq 2C \|x\|. \end{aligned}$$

Therefore, for all $x \in \mathcal{X}$ we have

$$d_*(f(x), f_n(x)) = \sum_{i=1}^{\infty} 2^{-i} |\langle \eta_i, f(x) - f_n(x) \rangle| \leq \sup_{i=1, \dots, n} |\langle \eta_i, f(x) - f_n(x) \rangle| \leq 2C \|x\|.$$

Let $\mu \in \mathcal{P}_p(\mathcal{X})$ be a probability measure on \mathcal{X} with $p \geq 1$ finite moments. Then the function $x \mapsto \|x\|$ is in $L^p(\mathcal{X})$, hence the sequence $d_*(f(x), f_n(x))$ is dominated by an L^p -function. Since $d_*(f(x), f_n(x)) \rightarrow 0$ pointwise on \mathcal{X} , convergence in $L^p_\mu(\mathcal{X}; (\mathcal{Y}, d_*))$ follows from the dominated convergence theorem. \square

Next we prove approximation rates of functions in $\mathcal{F}_1(\mathcal{X}; \mathcal{Y})$ with finite two-layer neural networks.

Theorem 3.24 (Direct approximation). *Let \mathcal{X}, \mathcal{Y} be duals of separable Banach spaces \mathcal{X}° and \mathcal{Y}° that possess Schauder bases and let \mathcal{Y} be such that lattice operations are 1-Lipschitz with respect to the d_* metric (3.10). Let $f \in \mathcal{F}_1(\mathcal{X}; \mathcal{Y})$. Then for every $c > 2\sqrt{2}$ there exists a sequence of finite-rank operators $\{K_i\}_{i=1}^n \subset \mathcal{L}(\mathcal{X}; \mathcal{Y})$ with $\|K_i\| \leq 1$ and a sequence of rational numbers $\{\alpha_i\}_{i=1}^n$ such that*

(i) for any $x \in \mathcal{X}$

$$d_* \left(\sum_{i=1}^n \alpha_i (K_i x)_+, f(x) \right) \leq \frac{c \|f\|_{\mathcal{F}_1} \|x\|}{\sqrt{n}};$$

(ii) if $\mu \in \mathcal{P}_p(\mathcal{X})$ is a probability measure on \mathcal{X} with $p \geq 1$ finite moments then the rate also holds in the Bochner space $L^p_\mu(\mathcal{X}; (\mathcal{Y}, d_*))$

$$\left\| f(x) - \sum_{i=1}^n \alpha_i (K_i x)_+ \right\|_{L^p_\mu(\mathcal{X}; (\mathcal{Y}, d_*))} \leq \frac{c \|f\|_{\mathcal{F}_1} (m_p(\mu))^{\frac{1}{p}}}{\sqrt{n}},$$

where $m_p(\mu) < \infty$ is the p -th moment of μ .

Proof. (i) Since $f \in \mathcal{F}_1(\mathcal{X}, \mathcal{Y})$, it can be written as follows, cf. (3.17)

$$f(x) = \langle \sigma(Bx), a \rangle = \int_{\mathbb{B}_{\mathcal{L}}} (Kx)_+ da(K),$$

where $a \in \mathcal{M}(\mathbb{B}_{\mathcal{L}(\mathcal{X}; \mathcal{Y})})$ is as in Proposition 3.21. Using the Hahn-Jordan decomposition $a = a_+ - a_-$ and normalising, we get that

$$f(x) = \|a_+\|_{\mathcal{M}} \langle \sigma(Bx), \rho^+ \rangle - \|a_-\|_{\mathcal{M}} \langle \sigma(Bx), \rho^- \rangle := f^+(x) - f^-(x), \quad (3.21)$$

where $\rho^\pm := \frac{a_\pm}{\|a_\pm\|_{\mathcal{M}}}$ are probability measures on $\mathbb{B}_{\mathcal{L}(\mathcal{X};\mathcal{Y})}$. Let $K_i^\pm \sim \rho^\pm$ be random i.i.d. samples from ρ^\pm , $i = 1, \dots, n$, and let $\rho_i^\pm := \frac{1}{n} \delta_{K_i^\pm}$. Define

$$f_n(x) := \|a_+\|_{\mathcal{M}} \left\langle \sigma(Bx), \sum_{i=1}^n \rho_i^+ \right\rangle - \|a_-\|_{\mathcal{M}} \left\langle \sigma(Bx), \sum_{i=1}^n \rho_i^- \right\rangle := f_n^+(x) - f_n^-(x).$$

Since finite-rank operators are weakly-* dense in $\mathbb{B}_{\mathcal{L}(\mathcal{X};\mathcal{Y})}$, we can assume without loss of generality that K_i^\pm are finite-rank.

Let $\{\eta_j\}_{j \in \mathbb{N}}$ be a normalised dense system in \mathcal{Y}° as defined in Proposition 3.9. Fix $x \in \mathcal{X}$ and consider the following functions $\varphi_j^x \in \mathcal{C}(\mathbb{B}_{\mathcal{L}(\mathcal{X};\mathcal{Y})})$

$$\varphi_j^x(K) := \langle \eta_j, (Kx)_+ \rangle, \quad j \in \mathbb{N},$$

where continuity is with respect to the weak-* topology on $\mathbb{B}_{\mathcal{L}(\mathcal{X};\mathcal{Y})}$. With this notation, we have that $\langle \eta_j, \langle \sigma(Bx), \rho^\pm \rangle \rangle = \langle \varphi_j^x, \rho^\pm \rangle$ and hence

$$d_*(f^\pm(x), f_n^\pm(x)) = \|a_\pm\|_{\mathcal{M}} \sum_{j=1}^{\infty} 2^{-j} \left| \left\langle \varphi_j^x, \rho^\pm - \sum_{i=1}^n \rho_i^\pm \right\rangle \right| = \|a_\pm\|_{\mathcal{M}} \mathbb{E}_j \left| \left\langle \varphi_j^x, \rho^\pm - \sum_{i=1}^n \rho_i^\pm \right\rangle \right|,$$

where we denoted by \mathbb{E}_j the integration over \mathbb{N} with respect to the discrete measure $j \mapsto 2^{-j}$.

Since ρ_i^\pm are random, $d_*^2(f^\pm(x), f_n^\pm(x))$ is a random variable. Taking its mean over the random samples ρ_i^\pm and using Jensen's inequality, we get the following estimate

$$\begin{aligned} \frac{1}{\|a_\pm\|_{\mathcal{M}}^2} \mathbb{E}_{\rho^\pm} (d_*^2(f^\pm(x), f_n^\pm(x))) &= \mathbb{E}_{\rho^\pm} \left(\mathbb{E}_j \left| \left\langle \varphi_j^x, \rho^\pm - \sum_{i=1}^n \rho_i^\pm \right\rangle \right|^2 \right) \\ &\leq \mathbb{E}_{\rho^\pm} \mathbb{E}_j \left| \left\langle \varphi_j^x, \rho^\pm - \sum_{i=1}^n \rho_i^\pm \right\rangle \right|^2 \\ &= \mathbb{E}_j \mathbb{E}_{\rho^\pm} \left| \left\langle \varphi_j^x, \rho^\pm - \sum_{i=1}^n \rho_i^\pm \right\rangle \right|^2. \end{aligned}$$

The central limit theorem implies that

$$\mathbb{E}_{\rho^\pm} \left| \left\langle \varphi_j^x, \rho^\pm - \sum_{i=1}^n \rho_i^\pm \right\rangle \right|^2 \leq \frac{\text{var}_{\rho^\pm}(\varphi_j^x)}{n},$$

where $\text{var}_{\rho^\pm}(\varphi_j^x)$ denotes the variance of φ_j^x with respect to ρ^\pm . This variance can be estimated as follows

$$\begin{aligned} \text{var}_{\rho^\pm}(\varphi_j^x) &:= \int_{\mathbb{B}_{\mathcal{L}}} (\varphi_j^x(K))^2 d\rho^\pm(K) - \left(\int_{\mathbb{B}_{\mathcal{L}}} \varphi_j^x(K) d\rho^\pm(K) \right)^2 \\ &= \int_{\mathbb{B}_{\mathcal{L}}} |\langle \eta_j, (Kx)_+ \rangle|^2 d\rho^\pm(K) - \left(\int_{\mathbb{B}_{\mathcal{L}}} \langle \eta_j, (Kx)_+ \rangle d\rho^\pm(K) \right)^2 \\ &\leq \int_{\mathbb{B}_{\mathcal{L}}} |\langle \eta_j, (Kx)_+ \rangle|^2 d\rho^\pm(K) \\ &\leq \int_{\mathbb{B}_{\mathcal{L}}} |\langle \eta_j, Kx \rangle|^2 d\rho^\pm(K) \\ &\leq \int_{\mathbb{B}_{\mathcal{L}}} \|K\|^2 \|x\|^2 d\rho^\pm(K) \\ &\leq \|x\|^2, \end{aligned}$$

since ρ^\pm is a probability measure. This estimate is uniform in j , hence we get that

$$\mathbb{E}_{\rho^\pm} (d_*^2(f^\pm(x), f_n^\pm(x))) \leq \frac{\|x\|^2 \|a_\pm\|_{\mathcal{M}}^2}{n}.$$

Jensen's inequality then implies that

$$\mathbb{E}_{\rho^\pm}(d_*) = \sqrt{(\mathbb{E}_{\rho^\pm}(d_*)^2)} \leq \sqrt{\mathbb{E}_{\rho^\pm}(d_*^2)} \leq \frac{\|x\| \|a_\pm\|_{\mathcal{M}}}{\sqrt{n}},$$

where d_* stands for $d_*(f^\pm(x), f_n^\pm(x))$. The variance of d_* can be estimated similarly:

$$\text{var}_{\rho^\pm}(d_*) = \mathbb{E}_{\rho^\pm}(d_*^2) - (\mathbb{E}_{\rho^\pm}(d_*))^2 \leq \mathbb{E}_{\rho^\pm}(d_*^2) \leq \frac{\|x\|^2 \|a_\pm\|_{\mathcal{M}}^2}{n}.$$

Chebyshev's inequality implies that for any $\varepsilon > 1$

$$P\left(|d_* - \mathbb{E}_{\rho^\pm}(d_*)| \geq \varepsilon \sqrt{\text{var}_{\rho^\pm}(d_*)}\right) \leq \frac{1}{\varepsilon^2}$$

and therefore

$$P\left(|d_* - \mathbb{E}_{\rho^\pm}(d_*)| < \varepsilon \sqrt{\text{var}_{\rho^\pm}(d_*)}\right) \geq 1 - \frac{1}{\varepsilon^2} > 0$$

for any $\varepsilon > 1$. (We note that the empirical measures ρ_i^+ and ρ_i^- arise from samples two different probability distributions ρ^\pm , which even have disjoint supports by the properties of the Hahn-Jordan decomposition.) Since the above probability is strictly positive, there exist samples \hat{K}_i^\pm and empirical measures $\hat{\rho}_i^\pm$ such that the following estimate holds for $\hat{f}_n^\pm := \|a_\pm\|_{\mathcal{M}} \langle \sigma(Bx), \sum_{i=1}^n \hat{\rho}_i^\pm \rangle$

$$\left|d_*(f^\pm(x), \hat{f}_n^\pm(x)) - \mathbb{E}_{\rho^\pm}(d_*(f^\pm(x), f_n^\pm(x)))\right| < \varepsilon \sqrt{\text{var}(d_*(f^\pm(x), f_n^\pm(x)))}$$

and

$$\begin{aligned} d_*(f^\pm(x), \hat{f}_n^\pm(x)) &< \mathbb{E}_{\rho^\pm}(d_*(f^\pm(x), f_n^\pm(x))) + \varepsilon \sqrt{\text{var}(d_*(f^\pm(x), f_n^\pm(x)))} \\ &\leq (1 + \varepsilon) \frac{\|x\| \|a_\pm\|_{\mathcal{M}}}{\sqrt{n}}. \end{aligned}$$

Denoting

$$\hat{f}_n := \|a_+\|_{\mathcal{M}} \left\langle \sigma(Bx), \sum_{i=1}^n \hat{\rho}_i^+ \right\rangle - \|a_-\|_{\mathcal{M}} \left\langle \sigma(Bx), \sum_{i=1}^n \hat{\rho}_i^- \right\rangle$$

and using the triangle inequality, we finally get the following estimate

$$\begin{aligned} d_*(f(x), \hat{f}_n(x)) &\leq d_*(f^+(x), \hat{f}_n^+(x)) + d_*(f^-(x), \hat{f}_n^-(x)) \\ &\leq (1 + \varepsilon) \frac{\|x\| (\|a_+\|_{\mathcal{M}} + \|a_-\|_{\mathcal{M}})}{\sqrt{n}} \\ &= (1 + \varepsilon) \frac{\|x\| \|a\|_{\mathcal{M}}}{\sqrt{n}} \\ &= (1 + \varepsilon) \frac{\|x\| \|f\|_{\mathcal{F}_1}}{\sqrt{n}}. \end{aligned}$$

The function $\hat{f}_n(x)$, in fact, contains $2n$ terms. Replacing n with $n/2$ and letting $c := (1 + \varepsilon)\sqrt{2} > 2\sqrt{2}$, we get the desired estimate.

(ii) Convergence rates in Bochner spaces can be obtained similarly to Theorem 3.22. □

Remark 3.25. The fact that the operators K_i in Theorems 3.22 and 3.24 are finite-rank is important from the computational point of view, since any operator we can realise on a computer will be finite-rank.

Remark 3.26. Monte-Carlo rates in Lebesgue spaces $L_\pi^2(\mathcal{X})$ were obtained in [10, 11] and Theorem 3.24 extends them to Bochner spaces $L_\pi^2(\mathcal{X}; (\mathcal{Y}, d_*))$. In the case $2 < p < \infty$, suboptimal rates $O(n^{-1/p})$ are obtained using an interpolation argument in [11] (discussion immediately after Theorem 3.8); our Theorem 3.24 improves on these results even in the scalar-valued case by providing Monte-Carlo rates for all $p \in [1, \infty)$.

3.2.3 Examples

As already pointed out in Remark 3.14, lattice operations are rarely sequentially continuous in the weak or weak-* topologies. Therefore, Theorems 3.22 and 3.24 make a rather strong assumption on the output space \mathcal{Y} .

Atomic Banach lattices. If \mathcal{Y} is a Banach lattice with an order continuous norm then lattice operations are weakly sequentially continuous if and only if \mathcal{Y} is purely atomic [53, Cor. 2.3].

Definition 3.27 (Purely atomic Banach lattice). Let E be a Banach lattice. An element $0 \neq a \in E_+$ is called an atom if for any $b \in E$ such that $0 \leq b \leq a$ it holds $b = \lambda a$ for some $\lambda \in \mathbb{R}$. A Banach lattice is called *purely atomic* if the only element that is disjoint from all atoms is the zero element.

The situation with sequential weak-* continuity of lattice operations is an even more delicate matter, and purely atomic dual spaces may fail to have weakly-* sequentially continuous lattice operations [54, §3].

Clearly, sequence spaces ℓ^p , $1 < p < \infty$, are purely atomic and reflexive, hence lattice operations are weakly-* sequentially continuous. Sequential weak-* continuity (and, indeed, 1-Lipschitzness under a certain choice of the basis in the predual) of lattice operations in ℓ^∞ was shown in Proposition 3.15. Lattice operations in ℓ^1 (considered as the dual of the space c of convergent sequences) are not sequentially continuous in the weak-* topology [54, §3].

Another example of purely atomic Banach lattices are $\mathcal{L}^p(\mu)$ spaces, $1 < p < \infty$, where the measure μ is purely atomic.

Lipschitz spaces. If \mathcal{Y} is a Riesz space but not a Banach lattice then the result of [53, Cor. 2.3] does not apply and atomicity is not required.

Let Ω be a complete metric space of finite diameter and let $\mathcal{Y} = \text{Lip}(\Omega)$ be the space of Lipschitz functions on Ω with the following norm

$$\|f\|_{\text{Lip}} := \max\{\text{Lip}(f), \|f\|_\infty\},$$

where $\text{Lip}(f)$ denotes the Lipschitz constant of f , and the pointwise partial order

$$f \geq g \iff f(x) \geq g(x) \quad \text{for all } x \in \Omega.$$

Theorem 3.28 ([47, Thm. 3.23 and Cor. 3.4]). *Let Ω be a complete metric space of finite diameter. Then the space $\text{Lip}(\Omega)$ has a unique predual. On bounded subsets of $\text{Lip}(\Omega)$ its weak-* topology coincides with the topology of pointwise convergence.*

The next result shows that lattice operations are sequentially weakly-* continuous in Lipschitz spaces.

Theorem 3.29. *Let Ω be a complete metric space of finite diameter. Then lattice operations in $\text{Lip}(\Omega)$ are sequentially weakly-* continuous.*

Proof. Let $f_n \rightharpoonup^* f$ weakly-* in $\text{Lip}(\Omega)$. Since weakly-* convergent sequences are bounded, we get from Theorem 3.28 that convergence is pointwise, i.e.

$$f_n(x) \rightarrow f(x) \quad \text{for all } x \in \Omega.$$

Consequently,

$$(f_n)_+(x) \rightarrow f_+(x) \quad \text{for all } x \in \Omega.$$

Since $\|f_+\|_{\text{Lip}} \leq \|f\|_{\text{Lip}}$ for any $f \in \text{Lip}(\Omega)$ [47], the sequence $(f_n)_+$ is bounded and therefore $(f_n)_+ \rightharpoonup^* f_+$ weakly-* in $\text{Lip}(\Omega)$. \square

Remark 3.30. As we can see from the proof, lattice operations are weakly-* continuous in $\text{Lip}(\Omega)$ because they are continuous in \mathbb{R} . With an appropriate choice of the dense system in the predual (cf. Proposition 3.9) we will also get that lattice operations are 1-Lipschitz with respect to the d_* metric (3.10) in $\text{Lip}(\Omega)$.

We conclude that the choice $\mathcal{Y} = \text{Lip}(\Omega)$ satisfies the conditions of Theorems 3.22 and 3.24.

4 Finding an optimal representation

It has been shown (e.g. [38]) that training a sufficiently wide (scalar-valued) two-layer neural network with weight decay [39] results in functions which are optimal with respect to the \mathcal{F}_1 norm. These functions solve the following variational problem

$$\min_{g \in \mathcal{F}_1} \frac{1}{pm} \sum_{i=1}^m |g(x_i) - f(x_i)|^p + \lambda \|g\|_{\mathcal{F}_1}, \quad (4.1)$$

where $\{(x_i, f(x_i))\}_{i=1}^m$ are the training data (possibly, corrupted by noise) and $\lambda > 0$ is a regularisation parameter. Using the representation $g = \langle \sigma(B \cdot), a \rangle$ valid for any \mathcal{F}_1 function, one can rewrite (4.1) as an optimisation problem over the space of measures, cf. (4.5). For a fixed data fitting term (i.e., fixed m), this is a well-studied problem, e.g. [40]. The contribution of this section compared to [40] is to study the consistency of minimisers as the number of samples $m \rightarrow \infty$. We refer to Remark 4.3 for a discussion of our setting in the context of existing literature.

Given a function $f \in \mathcal{F}_1(\mathcal{X}; \mathcal{Y})$ (whose explicit form may be unknown), consider the problem of finding the best neural network representation of f , i.e. a representation with the smallest variation norm. Define the following operator T acting on $\mathcal{M}(\mathbb{B}_{\mathcal{L}})$

$$Ta := \langle \sigma(B \cdot), a \rangle, \quad a \in \mathcal{M}(\mathbb{B}_{\mathcal{L}}). \quad (4.2)$$

By Remark 3.20 we have that T is continuous as an operator $\mathcal{M}(\mathbb{B}_{\mathcal{L}}) \rightarrow L^p_{\pi}(\mathcal{X}; (\mathcal{Y}, d_*))$, where $\pi \in \mathcal{P}_p(\mathcal{X})$ is a probability measure with $p \geq 1$ finite moments.

Finding the best representation (2.9) of f amounts to finding the minimum norm solution of the following inverse problem

$$Ta = f. \quad (4.3)$$

Definition 4.1 (Minimum-norm solution). A solution a^\dagger of (4.3) is called a minimum-norm solution if

$$\|a^\dagger\|_{\mathcal{M}} \leq \|a\|_{\mathcal{M}} \quad \text{for all } a \text{ such that } Ta = f.$$

The following result shows that T is a compact operator $\mathcal{M}(\mathbb{B}_{\mathcal{L}}) \rightarrow L^p_{\pi}(\mathcal{X}; (\mathcal{Y}, d_*))$.

Proposition 4.2. *The operator $T: a \mapsto \langle \sigma(B \cdot), a \rangle$ is compact as an operator $\mathcal{M}(\mathbb{B}_{\mathcal{L}}) \rightarrow L^p_{\pi}(\mathcal{X}; (\mathcal{Y}, d_*))$.*

Proof. The proof can be found in Appendix B. □

Since the inverse of a compact operator is unbounded, (4.3) is ill-posed and regularisation is required. In this section, we will study variational regularisation of (4.3). In particular, we will investigate the role of a further regularity condition on the function $f \in \mathcal{F}_1(\mathcal{X}; \mathcal{Y})$, known as the *source condition*, and its relation to convergence in Bregman distance and support recovery.

We will assume that $\{x_i\}_{i=1}^m \subset \mathcal{X}$ are i.i.d. samples $x_i \sim \pi$ from a probability measure π on \mathcal{X} . We will also assume the exact data $\{f(x_i)\}_{i=1}^m$ are corrupted by additive noise $\{e_i\}_{i=1}^m \subset \mathcal{Y}$ such that

$$\mathbb{E}_{\pi} \frac{1}{m} \sum_{i=1}^m \|e_i\|_{(\mathcal{Y}, d_*)}^p \leq \varepsilon^p \quad \text{and} \quad \mathbb{E}_{\pi} \frac{1}{m} \sum_{i=1}^m \|e_i\|_{(\mathcal{Y}, d_*)}^{2p} \leq \sigma^2 < \infty \quad (4.4)$$

uniformly in m for some $\varepsilon, \sigma > 0$. We end up with the following variational problem, cf. (4.1)

$$\min_{a \in \mathcal{M}(\mathbb{B}_{\mathcal{L}})} \frac{1}{pm} \sum_{i=1}^m \|Ta(x_i) - f(x_i) + e_i\|_{(\mathcal{Y}, d_*)}^p + \lambda \|a\|_{\mathcal{M}}. \quad (4.5)$$

Existence of a minimiser for any fixed $\lambda > 0$ follows from standard arguments, e.g., [46]. We denote this minimiser by $a_{m, \varepsilon}$. Since (4.5) is not strictly convex, this minimiser is not unique. All consequent results will hold for any minimiser of (4.5).

Remark 4.3. We will study the behaviour of minimisers in the limit as $\varepsilon \rightarrow 0$ and $m \rightarrow \infty$. This is standard in inverse problems literature [46, 55], but different from the setting of statistical learning, where ε is assumed fixed and only $m \rightarrow \infty$. Due to the ill-posedness of (4.3), convergence of statistical

estimators can be arbitrary slow (e.g., [43]) unless regularisation is used. (We emphasise that this applies to estimators of the representing measure a^\dagger , not the function f itself.) For spectral regularisation, convergence rates in the statistical setting were obtained in [44]. Rates for variational regularisation with convex p -homogeneous functionals, $p > 1$, were obtained in a recent paper [45]. Our case of a 1-homogeneous functional $\|\cdot\|_{\mathcal{M}}$ is not covered by existing literature. Obtaining convergence rates for solutions of (4.5) with fixed ε in the spirit of [44, 45] is beyond the scope of the present paper.

Proposition 4.4. *The following statements about a minimiser $a_{m,\varepsilon}$ of (4.5) hold*

$$(i) \quad \mathbb{E}_\pi \|a_{m,\varepsilon}\|_{\mathcal{M}} \leq \frac{\varepsilon^p}{p\lambda} + \|a^\dagger\|_{\mathcal{M}};$$

$$(ii) \quad \mathbb{E}_\pi \|a_{m,\varepsilon}\|_{\mathcal{M}}^2 \leq 2 \left(\frac{\sigma^2}{p^2\lambda^2} + \|a^\dagger\|_{\mathcal{M}}^2 \right);$$

(iii) for any $0 < \delta < 1$,

$$\|a_{m,\varepsilon}\|_{\mathcal{M}} \leq \frac{\varepsilon^p}{p\lambda} + \|a^\dagger\|_{\mathcal{M}} + \frac{\sqrt{2}}{\delta} \left(\frac{\sigma}{p\lambda} + \|a^\dagger\|_{\mathcal{M}} \right)$$

with probability at least $1 - \delta^2$.

Proof.

(i) Using the optimality of $a_{m,\varepsilon}$ in (4.5) and comparing the value of the objective with the value at any minimum norm solution \hat{a} , we get (recall that $T\hat{a} = f$)

$$\frac{1}{pm} \sum_{i=1}^m \|Ta_{m,\varepsilon}(x_i) - f(x_i) + e_i\|_{(\mathcal{Y}, d_*)}^p + \lambda \|a_{m,\varepsilon}\|_{\mathcal{M}} \leq \frac{1}{pm} \sum_{i=1}^m \|e_i\|_{(\mathcal{Y}, d_*)}^p + \lambda \|\hat{a}\|_{\mathcal{M}}, \quad (4.6)$$

which also holds if we take the expected value over $\{x_i\}_{i=1}^m \sim \pi$. Therefore, we get

$$\mathbb{E}_\pi \|a_{m,\varepsilon}\|_{\mathcal{M}} \leq \frac{1}{p\lambda} \mathbb{E}_\pi \frac{1}{m} \sum_{i=1}^m \|e_i\|_{(\mathcal{Y}, d_*)}^p + \|\hat{a}\|_{\mathcal{M}} \leq \frac{\varepsilon^p}{p\lambda} + \|\hat{a}\|_{\mathcal{M}}. \quad (4.7)$$

(ii) Similarly, from (4.6) we get

$$\|a_{m,\varepsilon}\|_{\mathcal{M}}^2 \leq \left(\frac{1}{p\lambda} \frac{1}{m} \sum_{i=1}^m \|e_i\|_{(\mathcal{Y}, d_*)}^p + \|\hat{a}\|_{\mathcal{M}} \right)^2 \leq 2 \left(\frac{1}{p^2\lambda^2} \frac{1}{m} \sum_{i=1}^m \|e_i\|_{(\mathcal{Y}, d_*)}^{2p} + \|a^\dagger\|_{\mathcal{M}}^2 \right),$$

where we used Young's and Jensen's inequalities. Taking the expectation, we obtain the claim.

(iii) This follows from Chebyshev's inequality. □

4.1 Source condition

In order to characterise the convergence of the solutions of (4.5) in terms of convergence rates, one needs to make additional assumptions on the regularity of f . We are particularly interested in a *structural* characterisation of this convergence in terms of the convergence of the support of the measure, which is usually done using Bregman distances. We will now introduce some necessary concepts.

Definition 4.5 (Subdifferential). Let E be a Banach space and $\mathcal{J}: E \rightarrow \mathbb{R} \cup \{+\infty\}$ a proper convex lower semicontinuous functional. An element $\zeta \in E^*$ is called a subgradient of \mathcal{J} at $x_0 \in E$ if

$$\mathcal{J}(x) \geq \mathcal{J}(x_0) + \langle \zeta, x - x_0 \rangle \quad \forall x \in E.$$

The collection of all subgradients at x_0 is called the subdifferential of \mathcal{J} at x_0 and denoted by

$$\partial\mathcal{J}(x_0) := \{\zeta \in E^* : \zeta \text{ is a subgradient of } \mathcal{J} \text{ at } x_0\}.$$

Example 4.6 (Subgradients of the Radon norm). Consider $\mathcal{J}(\mu) := \|\mu\|_{\mathcal{M}}$ for $\mu \in \mathcal{M}(\Omega)$, where Ω is compact. Let $\mu_0 \in \mathcal{M}(\Omega)$. It is well known (e.g., [40]) that $\zeta_0 \in \partial \|\cdot\|_{\mathcal{M}}(\mu_0)$ if and only if

$$\|\zeta_0\|_{\mathcal{M}^*} \leq 1 \quad \text{and} \quad \|\mu_0\|_{\mathcal{M}} = \langle \zeta_0, \mu_0 \rangle. \quad (4.8)$$

Applying the Cauchy-Schwarz inequality, we get that

$$\|\mu_0\|_{\mathcal{M}} \leq \|\zeta_0\|_{\mathcal{M}^*} \|\mu_0\|_{\mathcal{M}} \leq \|\mu_0\|_{\mathcal{M}}$$

and hence the Cauchy-Schwarz inequality holds as an equality. For a sufficiently regular subgradient $\zeta_0 \in \mathcal{C}(\Omega)$, this implies that

$$\zeta_0 = \text{sign}(\mu_0) \quad \text{on } \text{supp}(\mu_0).$$

Since (4.3) is ill-posed, convergence of the minimisers of (4.5) to a minimum-norm solution may be arbitrary slow unless a further regularity assumption on a^\dagger , called the *source condition*, is made. Depending on the situation, different variants of the source condition exist. We use the definition from [56].

Assumption 1 (Source condition). The minimum-norm solution a^\dagger of (4.3) satisfies the source condition, i.e. there exists an $\omega^\dagger \in (L_\pi^q(\mathcal{X}; (\mathcal{Y}, d_*)^*))^* = L_\pi^q(\mathcal{X}; ((\mathcal{Y}, d_*)^*))^*$ such that

$$T^* \omega^\dagger \in \partial \|\cdot\|_{\mathcal{M}}(a^\dagger),$$

where $T^*: L_\pi^q(\mathcal{X}; ((\mathcal{Y}, d_*)^*))^* \mapsto \mathcal{M}^*(\mathbb{B}_{\mathcal{L}(\mathcal{X}; \mathcal{Y})})$ is the adjoint of T .

To understand the implications of Assumption 1, we need to study the adjoint operator T^* . It is easy to show using Fubini's theorem that for any $v \in L_\pi^q(\mathcal{X}; ((\mathcal{Y}, d_*)^*))^*$ (q is the conjugate exponent of p)

$$T^* v = \int_{\mathcal{X}} \langle \sigma(Bx), v(x) \rangle d\pi(x) \in \mathcal{M}^*(\mathbb{B}_{\mathcal{L}(\mathcal{X}; \mathcal{Y})}), \quad (4.9)$$

where the pairing is between (\mathcal{Y}, d_*) and $((\mathcal{Y}, d_*)^*)^*$, and for any $K \in \mathbb{B}_{\mathcal{L}(\mathcal{X}; \mathcal{Y})}$ we have

$$(T^* v)(K) = \int_{\mathcal{X}} \langle (Kx)_+, v(x) \rangle d\pi(x), \quad (4.10)$$

where $(\cdot)_+$ is the operation of taking the positive part in \mathcal{Y} . At this point, (4.10) is a formal expression, since we don't know whether $T^* v$ is continuous on $\mathbb{B}_{\mathcal{L}(\mathcal{X}; \mathcal{Y})}$ with respect to weak-* convergence. The next result shows that this is the case. Hence, as in the previous section, T^* maps into a more regular space than $\mathcal{M}^*(\mathbb{B}_{\mathcal{L}(\mathcal{X}; \mathcal{Y})})$ – its double predual $\mathcal{C}(\mathbb{B}_{\mathcal{L}(\mathcal{X}; \mathcal{Y})})$.

Proposition 4.7. *Let $\pi \in \mathcal{P}_p(\mathcal{X})$ have $p \geq 1$ finite moments. Suppose that lattice operations in \mathcal{Y} are 1-Lipschitz with respect to the d_* metric. Then $\forall v \in L_\pi^q(\mathcal{X}; ((\mathcal{Y}, d_*)^*))^*$*

$$T^* v \in \mathcal{C}(\mathbb{B}_{\mathcal{L}(\mathcal{X}; \mathcal{Y})}), \quad (4.11)$$

where continuity is with respect to the weak-* convergence in $\mathbb{B}_{\mathcal{L}(\mathcal{X}; \mathcal{Y})}$. Moreover, if we consider the unit ball $\mathbb{B}_{\mathcal{L}(\mathcal{X}; \mathcal{Y})}$ with the strong (norm-based) metric, then $T^* v$ is Lipschitz and its Lipschitz constant is bounded by

$$\text{Lip}(T^* v) \leq (m_p(\pi))^{\frac{1}{p}} \|v\|_{L_\pi^q},$$

where $m_p(\pi)$ is the p -th moment of π .

Proof. Let $K_n \xrightarrow{*} K$ weakly-* in $\mathbb{B}_{\mathcal{L}(\mathcal{X}; \mathcal{Y})}$. By (4.10) we have

$$\begin{aligned} |(T^* v)(K_n) - (T^* v)(K)| &= \left| \int_{\mathcal{X}} \langle (K_n x)_+ - (Kx)_+, v(x) \rangle d\pi(x) \right| \\ &\leq \int_{\mathcal{X}} \|(K_n x)_+ - (Kx)_+\|_{(\mathcal{Y}, d_*)} \|v(x)\|_{((\mathcal{Y}, d_*)^*)^*} d\pi(x) \\ &\leq \|(K_n \cdot)_+ - (K \cdot)_+\|_{L_\pi^p} \|v\|_{L_\pi^q}. \end{aligned}$$

Now, for any $x \in \mathcal{X}$ we have

$$\|(K_n x)_+ - (Kx)_+\|_{(\mathcal{Y}, d_*)} = d_*((K_n x)_+ - (Kx)_+, 0) = d_*((K_n x)_+, (Kx)_+) \leq d_*(K_n x, Kx),$$

since lattice operations are 1-Lipschitz with respect to the d_* metric in \mathcal{Y} .

We also have that for any $K_1, K_2 \in \mathbb{B}_{\mathcal{L}(\mathcal{X};\mathcal{Y})}$ and any $x \in \mathcal{X}$

$$d_*(K_1x, K_2x) \leq \|K_1x - K_2x\|_{\mathcal{Y}} \leq \|K_1 - K_2\|_{\mathcal{L}(\mathcal{X};\mathcal{Y})} \|x\|_{\mathcal{X}}$$

and therefore

$$d_*(K_nx, Kx)^p \leq \|K_n - K\|^p \|x\|^p \leq 2^p \|x\|^p$$

for all n . Since π has p finite moments, the right-hand side is integrable and, using the reverse Fatou lemma, we get that

$$\begin{aligned} \limsup_{n \rightarrow \infty} \|(K_n \cdot)_+ - (K \cdot)_+\|_{L^p_{\pi}}^p &= \limsup_{n \rightarrow \infty} \int_{\mathcal{X}} \|(K_nx)_+ - (Kx)_+\|_{(\mathcal{Y}, d_*)}^p d\pi(x) \\ &\leq \limsup_{n \rightarrow \infty} \int_{\mathcal{X}} (d_*(K_nx, Kx))^p d\pi(x) \\ &\leq \int_{\mathcal{X}} \limsup_{n \rightarrow \infty} (d_*(K_nx, Kx))^p d\pi(x) \\ &= 0, \end{aligned}$$

where we used the fact that $K_nx \xrightarrow{*} Kx$ in \mathcal{Y} for all $x \in \mathcal{X}$ if $K_n \xrightarrow{*} K$ in $\mathbb{B}_{\mathcal{L}(\mathcal{X};\mathcal{Y})}$. Therefore, $|(T^*v)(K_n) - (T^*v)(K)| \rightarrow 0$ and consequently $T^*v \in \mathcal{C}(\mathbb{B}_{\mathcal{L}(\mathcal{X};\mathcal{Y})})$.

To prove the second claim, we observe that

$$\int_{\mathcal{X}} (d_*(K_1x, K_2x))^p d\pi(x) \leq \|K_1 - K_2\|_{\mathcal{L}(\mathcal{X};\mathcal{Y})}^p \int_{\mathcal{X}} \|x\|_{\mathcal{X}}^p d\pi(x) = \|K_1 - K_2\|_{\mathcal{L}(\mathcal{X};\mathcal{Y})}^p m_p(\pi)$$

and

$$|(T^*v)(K_1) - (T^*v)(K_2)| \leq (m_p(\pi))^{\frac{1}{p}} \|v\|_{L^q_{\pi}} \|K_1 - K_2\|_{\mathcal{L}(\mathcal{X};\mathcal{Y})}.$$

□

Now we can give an intuitive interpretation of the source condition (Assumption 1). Let a_+^{\dagger} and a_-^{\dagger} be the positive and negative parts of a^{\dagger} in the sense of the Hahn-Jordan decomposition. From Example 4.6 we know that for any subgradient $\zeta \in \partial \|\cdot\|_{\mathcal{M}}(a^{\dagger}) \cap \mathcal{C}(\mathbb{B}_{\mathcal{L}(\mathcal{X};\mathcal{Y})})$

$$\zeta|_{\text{supp}(a_+^{\dagger})} \equiv 1 \quad \text{and} \quad \zeta|_{\text{supp}(a_-^{\dagger})} \equiv -1. \quad (4.12)$$

If a^{\dagger} satisfies Assumption 1, the subgradient $T^*\omega^{\dagger}$ is continuous with respect to the d_* metric, hence it is not allowed to jump from -1 to 1 without passing all values inbetween. In other words, a^{\dagger} is not allowed to change sign without staying zero on some interval.

Since with respect to the strong (norm-based) metric $T^*\omega^{\dagger}$ is Lipschitz on $\mathbb{B}_{\mathcal{L}(\mathcal{X};\mathcal{Y})}$, we have a lower bound on the length of this interval (in the norm-induced metric on $\mathbb{B}_{\mathcal{L}(\mathcal{X};\mathcal{Y})}$).

Proposition 4.8. *Let a^{\dagger} satisfy Assumption 1 and denote by a_+^{\dagger} and a_-^{\dagger} the positive and negative parts of a^{\dagger} (in the sense of the Hahn-Jordan decomposition). Then for any $K \in \text{supp}(a_+^{\dagger})$ and $K' \in \text{supp}(a_-^{\dagger})$ we have*

$$\|K - K'\|_{\mathcal{L}(\mathcal{X};\mathcal{Y})} \geq \frac{2}{(m_p(\pi))^{\frac{1}{p}} \|\omega^{\dagger}\|_{L^q_{\pi}}},$$

where ω^{\dagger} is as in Assumption 1 and $m_p(\pi)$ is the p -th moment of π .

Proof. From Example 4.6 we know that

$$T^*\omega^{\dagger} \equiv 1 \quad \text{on } \text{supp}(a_+^{\dagger}) \quad \text{and} \quad T^*\omega^{\dagger} \equiv -1 \quad \text{on } \text{supp}(a_-^{\dagger}).$$

Taking any $K \in \text{supp}(a_+^{\dagger})$ and $K' \in \text{supp}(a_-^{\dagger})$, we get that

$$|T^*\omega^{\dagger}(K) - T^*\omega^{\dagger}(K')| = 2.$$

Using the upper bound on the Lipschitz constant of $T^*\omega^{\dagger}$ from Proposition 4.7, we get

$$(m_p(\pi))^{\frac{1}{p}} \|\omega^{\dagger}\|_{L^q_{\pi}} \geq \text{Lip}(T^*\omega) \geq \frac{|T^*\omega^{\dagger}(K) - T^*\omega^{\dagger}(K')|}{\|K - K'\|_{\mathcal{L}(\mathcal{X};\mathcal{Y})}} = \frac{2}{\|K - K'\|_{\mathcal{L}(\mathcal{X};\mathcal{Y})}},$$

which implies the claim. □

In terms of the representation (2.9) this means that the network is not allowed to have two neurons with opposite signs and weights that are too close to each other, the distance being controlled by the source element ω^\dagger . Hence, intuitively, \mathcal{F}_1 functions satisfying the source condition are more stable in the sense that cancellations are avoided, and this effect is the stronger the smaller the norm of the source element.

Remark 4.9. The inequality (4.12) is a separation condition for the supports of the positive and the negative parts of the measure a^\dagger in terms of the strong (norm-based) metric on $\mathbb{B}_{\mathcal{L}(\mathcal{X};\mathcal{Y})}$. No such separation can be obtained in the weak-* metric (3.8). However, other metrics might admit a separation condition. For example, [42] argues that the Fisher-Rao metric is a good choice for studying support separation. Such questions are outside the scope of the present paper.

4.2 Convergence rates in Bregman distance

In modern variational regularisation, (generalised) Bregman distances are typically used to study convergence rates of solutions [46].

Definition 4.10 (Bregman distance). Let E be a Banach space and $\mathcal{J}: E \rightarrow \mathbb{R} \cup \{+\infty\}$ be a proper convex functional. The generalised Bregman distance between $u, u' \in E$ corresponding to the subgradient $\zeta' \in \partial\mathcal{J}(u')$ is defined as follows

$$D_{\mathcal{J}}^{\zeta'}(u, u') := \mathcal{J}(u) - \mathcal{J}(u') - \langle \zeta', u - u' \rangle,$$

where $\partial\mathcal{J}(u')$ denotes the subdifferential of \mathcal{J} at $u' \in E$. The symmetric Bregman distance between u and u' corresponding to $\zeta \in \partial\mathcal{J}(u)$ and $p' \in \partial\mathcal{J}(u')$ is defined as follows

$$D_{\mathcal{J}}^{\text{symm}}(u, u') := D_{\mathcal{J}}^{\zeta'}(u, u') + D_{\mathcal{J}}^{\zeta}(u', u) = \langle \zeta - \zeta', u - u' \rangle.$$

Bregman distances do not define a metric, since they do not satisfy the triangle inequality and $D_{\mathcal{J}}^{\text{symm}}(u, u') = 0$ does not imply $u = u'$ in general.

Example 4.11 (Bregman distance of the Radon norm). Consider $\mathcal{J}(\mu) := \|\mu\|_{\mathcal{M}}$ for $\mu \in \mathcal{M}(\Omega)$, where Ω is compact. Let $\mu_0 \in \mathcal{M}(\Omega)$ and $\zeta_0 \in \partial\|\cdot\|_{\mathcal{M}}(\mu_0)$. Then for any $\mu \in \mathcal{M}(\Omega)$, the Bregman distance between μ_0 and μ corresponding to the subgradient ζ_0 is given by

$$D_{\|\cdot\|_{\mathcal{M}}}^{\zeta_0}(\mu, \mu_0) = \|\mu\|_{\mathcal{M}} - \langle \zeta_0, \mu \rangle.$$

Zero Bregman distance would imply that $\zeta_0 \in \partial\|\cdot\|_{\mathcal{M}}(\mu)$. In light of Example 4.6, this means that $\zeta_0 = \text{sign}(\mu)$ on $\text{supp}(\mu)$. Hence, roughly speaking, $D_{\|\cdot\|_{\mathcal{M}}}^{\zeta_0}(\mu, \mu_0)$ measures the deviation of $\text{supp}(\mu)$ from $\text{supp}(\mu_0)$.

We will need the following result.

Lemma 4.12. Let $\pi \in \mathcal{P}_{2p}(\mathcal{X})$ be a probability measure with $2p$ finite moments and $f \in \mathcal{F}_1(\mathcal{X};\mathcal{Y})$. Then for any $0 < \delta < 1$

$$\left| \frac{1}{m} \sum_{i=1}^m \|f(x_i)\|_{(\mathcal{Y}, d_*)}^p - \|f\|_{L_{\pi}^p}^p \right| \leq \frac{1}{\delta} \frac{\|f\|_{\mathcal{F}_1}^p \sqrt{m_{2p}(\pi)}}{\sqrt{m}},$$

with probability at least $1 - \delta^2$, where $m_{2p}(\pi) < \infty$ is the $2p$ -th moment of π .

Proof. Since f is Lipschitz (with Lipschitz constant bounded by $\|f\|_{\mathcal{F}_1}$, see Proposition 3.2), it grows at most linearly at infinity, and we get

$$\begin{aligned} \text{var}_{\pi} \left(\frac{1}{m} \sum_{i=1}^m \|f(x_i)\|_{(\mathcal{Y}, d_*)}^p \right) &= \frac{1}{m^2} \sum_{i=1}^m \text{var}_{\pi} \left(\|f(x_i)\|_{(\mathcal{Y}, d_*)}^p \right) \leq \frac{1}{m} \mathbb{E}_{\pi} \left[\|f(x_1)\|_{(\mathcal{Y}, d_*)}^{2p} \right] \\ &= \frac{1}{m} \int_{\mathcal{X}} \|f(x)\|_{(\mathcal{Y}, d_*)}^{2p} d\pi(x) \leq \frac{\|f\|_{\mathcal{F}_1}^{2p}}{m} \int_{\mathcal{X}} \|x\|^{2p} d\pi(x) = \frac{\|f\|_{\mathcal{F}_1}^{2p} m_{2p}(\pi)}{m}. \end{aligned}$$

The claim now follows from Chebyshev's inequality. \square

Theorem 4.13. *Let $\pi \in \mathcal{P}_{2p}(\mathcal{X})$ be a probability measure with $2p$ finite moments and suppose that a minimum-norm solution a^\dagger satisfies the source condition (Assumption 1). Then, up to multiplicative constants and higher-order terms in λ and ε , the following estimate holds for any $0 < \delta < 1$ with probability at least $1 - 3\delta^2$*

$$D_{\|\cdot\|_{\mathcal{M}}}^{\zeta^\dagger}(a_{m,\varepsilon}, a^\dagger) \leq \begin{cases} \frac{\varepsilon^p}{\lambda} + \frac{1}{\delta^p \lambda^{p+1} \sqrt{m}} + \lambda^{q-1} \|\omega^\dagger\|_{L_\pi^q}^q, & \text{if } p > 1, \\ \frac{\varepsilon^p}{\lambda} + \frac{1}{\delta^p \lambda^{p+1} \sqrt{m}}, & \text{if } p = 1 \text{ and } \lambda \leq \frac{1}{\|\omega^\dagger\|}, \end{cases}$$

where $a_{m,\varepsilon}$ is a solution of (4.5), q is the Hölder conjugate of p , ω^\dagger is the source element from Assumption 1, and $\zeta^\dagger := T^* \omega^\dagger$.

Proof. The proof is similar to, e.g., [57, Thm. 3.5]. Subtracting $\lambda \langle \zeta^\dagger, a_{m,\varepsilon} \rangle$ from both sides of (4.6) and rearranging terms, we get

$$\lambda D_{\|\cdot\|_{\mathcal{M}}}^{\zeta^\dagger}(a_{m,\varepsilon}, a^\dagger) \leq \frac{1}{pm} \sum_{i=1}^m \|e_i\|_{(\mathcal{Y}, d_*)}^p - \frac{1}{pm} \sum_{i=1}^m \|Ta_{m,\varepsilon}(x_i) - f(x_i) + e_i\|_{(\mathcal{Y}, d_*)}^p + \lambda \|\hat{a}\|_{\mathcal{M}} - \lambda \langle \zeta^\dagger, a_{m,\varepsilon} \rangle.$$

Using the inequality $\|g - g'\|^p \leq 2^{p-1} (\|g\|^p + \|g'\|^p)$, which is valid for the convex and absolutely p -homogeneous function $\|\cdot\|^p$ with $p \geq 1$, we get that

$$2^{1-p} \|Ta_{m,\varepsilon}(x_i) - f(x_i)\|_{(\mathcal{Y}, d_*)}^p \leq \|Ta_{m,\varepsilon}(x_i) - f(x_i) + e_i\|_{(\mathcal{Y}, d_*)}^p + \|e_i\|_{(\mathcal{Y}, d_*)}^p.$$

Therefore, we have the following estimate

$$\lambda D_{\|\cdot\|_{\mathcal{M}}}^{\zeta^\dagger}(a_{m,\varepsilon}, a^\dagger) \leq \frac{2}{pm} \sum_{i=1}^m \|e_i\|_{(\mathcal{Y}, d_*)}^p - \frac{2^{1-p}}{pm} \sum_{i=1}^m \|Ta_{m,\varepsilon}(x_i) - f(x_i)\|_{(\mathcal{Y}, d_*)}^p + \lambda \langle \omega^\dagger, f - Ta_{m,\varepsilon} \rangle.$$

By Lemma 4.12 we have that

$$\begin{aligned} -\frac{1}{pm} \sum_{i=1}^m \|Ta_{m,\varepsilon}(x_i) - f(x_i)\|_{(\mathcal{Y}, d_*)}^p &\leq -\frac{1}{p} \|Ta_{m,\varepsilon} - f\|_{L_\pi^p}^p + \frac{1}{\delta} \frac{(\|a_{m,\varepsilon}\|_{\mathcal{M}} + \|a^\dagger\|_{\mathcal{M}})^p \sqrt{m_{2p}(\pi)}}{\sqrt{m}} \\ &\leq -\frac{1}{p} \|Ta_{m,\varepsilon} - f\|_{L_\pi^p}^p + \frac{2^{p-1}}{\delta} \frac{(\|a_{m,\varepsilon}\|_{\mathcal{M}}^p + \|a^\dagger\|_{\mathcal{M}}^p) \sqrt{m_{2p}(\pi)}}{\sqrt{m}} \end{aligned}$$

with probability at least $1 - \delta^2$. Therefore,

$$\begin{aligned} -\frac{2^{1-p}}{pm} \sum_{i=1}^m \|Ta_{m,\varepsilon}(x_i) - f(x_i)\|_{(\mathcal{Y}, d_*)}^p + \lambda \langle \omega^\dagger, f - Ta_{m,\varepsilon} \rangle &\leq \frac{1}{\delta} \frac{(\|a_{m,\varepsilon}\|_{\mathcal{M}}^p + \|a^\dagger\|_{\mathcal{M}}^p) \sqrt{m_{2p}(\pi)}}{\sqrt{m}} \\ &\quad + 2^{1-p} \left[\langle 2^{p-1} \lambda \omega^\dagger, f - Ta_{m,\varepsilon} \rangle - \frac{1}{p} \|f - Ta_{m,\varepsilon}\|_{L_\pi^p}^p \right] \end{aligned}$$

with the same probability. By the Fenchel-Young inequality we have that

$$\langle 2^{p-1} \lambda \omega^\dagger, f - Ta_{m,\varepsilon} \rangle - \frac{1}{p} \|f - Ta_{m,\varepsilon}\|_{L_\pi^p}^p \leq \begin{cases} \frac{1}{q} \|2^{p-1} \lambda \omega^\dagger\|_{L_\pi^q}^q, & p > 1, \\ \chi_{\|\cdot\|_{L_\pi^\infty} \leq 1}(\lambda \omega^\dagger), & p = 1, \end{cases}$$

where q is the conjugate exponent of p and $\chi_{\|\cdot\|_{L_\pi^\infty} \leq 1}$ is the indicator function of the unit ball in L_π^∞ .

By Proposition 4.4 we have that for any $0 < \delta < 1$

$$\|a_{m,\varepsilon}\|_{\mathcal{M}}^p \leq \left(\frac{\varepsilon^p}{p\lambda} + \|a^\dagger\|_{\mathcal{M}} + \frac{\sqrt{2}}{\delta} \left(\frac{\sigma}{p\lambda} + \|a^\dagger\|_{\mathcal{M}} \right) \right)^p \leq C(p) \left(\frac{\varepsilon^{2p}}{\lambda^p} + \|a^\dagger\|_{\mathcal{M}}^p \left(1 + \frac{1}{\delta^p} \right) + \frac{\sigma^p}{\delta^p \lambda^p} \right)$$

with probability at least $1 - \delta^2$, where $C(p)$ is a constant that depends only on p .

Furthermore, from (4.4) and Chebyshev's inequality, we get for any $0 < \delta < 1$

$$\frac{1}{m} \sum_{i=1}^m \|e_i\|_{(\mathcal{Y}, d_*)}^p \leq \varepsilon^p + \frac{1}{\delta} \frac{\sigma}{\sqrt{m}}$$

with probability at least $1 - \delta^2$.

Combining everything, we get that, up to multiplicative constants and higher-order terms in λ and ε , with probability at least $1 - 3\delta^2$

$$\lambda D_{\|\cdot\|_{\mathcal{M}}}^{\zeta^\dagger}(a_{m,\varepsilon}, a^\dagger) \leq \varepsilon^p + \frac{1}{\delta^p \lambda^p \sqrt{m}} + \begin{cases} \|\lambda \omega^\dagger\|_{L^q}^q, & p > 1, \\ \chi_{\|\cdot\|_{L^\infty} \leq 1}(\lambda \omega^\dagger), & p = 1 \end{cases} \quad (4.13)$$

A division by λ yields the claim. \square

Remark 4.14. The estimate in the above theorem suggests the following optimal choice of the regularisation parameter λ and the number of samples m

$$\begin{aligned} \lambda &\sim \varepsilon^{p-1}, \quad m \sim \varepsilon^{-2p^2}, & p > 1, \\ \lambda &= \text{const} \leq \frac{1}{\|\omega^\dagger\|}, \quad m \sim \varepsilon^{-2}, & p = 1, \end{aligned}$$

under which we have

$$D_{\|\cdot\|_{\mathcal{M}}}^{\zeta^\dagger}(a_{m,\varepsilon}, a^\dagger) = O(\varepsilon).$$

Remark 4.15. Theorem 4.13 shows that under a structural assumption on the representing measure (the support separation condition, see the discussion immediately after Proposition 4.7) one can beat the curse of dimensionality not only in terms of approximating the \mathcal{F}_1 function itself, but also in terms of recovering the representing measure, albeit in a weaker metric – the Bregman distance. As discussed in Example 4.11, this corresponds to recovering the support of the representing measure.

We will finish with a few remarks.

Remark 4.16 (Debiasing). As pointed out in Example 4.11, convergence in Bregman distance for the Radon norm corresponds to the convergence of the support of the regularised solution $a_{m,\varepsilon}$ to the support of the minimum-norm solution a^\dagger . However, it does not say anything about the convergence of the density in regions where the support has been correctly identified. In fact, it is well known that variational regularisation introduces a bias, i.e. the value of the regularisation functional at the regularised solution $\|a_{m,\varepsilon}\|_{\mathcal{M}}$ is smaller than the value at the minimum-norm solution $\|a^\dagger\|_{\mathcal{M}}$. In the mathematical imaging literature, the following procedure has been proposed to compensate for this bias [58, 59].

Let $\zeta_{m,\varepsilon} \in \partial \|\cdot\|_{\mathcal{M}}(a_{m,\varepsilon})$ be a subgradient of the Radon norm at the regularised solution. Often it is available as a by-product of primal-dual optimisation methods applied to (4.5). Consider the following problem

$$\min_{a \in \mathcal{M}(\mathbb{B}_{\mathcal{L}})} \frac{1}{pm} \sum_{i=1}^m \|Ta(x_i) - f(x_i) + e_i\|_{(\mathcal{Y}, d_*)}^p \quad \text{s.t.} \quad D_{\|\cdot\|_{\mathcal{M}}}^{\zeta_{m,\varepsilon}}(a, a_{m,\varepsilon}) = 0. \quad (4.14)$$

Since the Bregman distance $D_{\|\cdot\|_{\mathcal{M}}}^{\zeta_{m,\varepsilon}}(\cdot, a_{m,\varepsilon})$ is a non-negative convex function of its first argument, (4.14) is a convex problem because the constraint can be equivalently written as follows

$$D_{\|\cdot\|_{\mathcal{M}}}^{\zeta_{m,\varepsilon}}(a, a_{m,\varepsilon}) = 0 \iff \|a\|_{\mathcal{M}} - \langle a, \zeta_{m,\varepsilon} \rangle \leq 0.$$

Solutions of (4.14) have the same support as $a_{m,\varepsilon}$, but better fit the data. This procedure is referred to as debiasing.

Remark 4.17 (Connection to existence theorems). For $p > 1$, the data term in (4.5) is strictly convex. Since a is a scalar-valued measure (all the vector-valuedness is hidden in the data term), we can use standard *representer theorems* [60, Thm. 2 and Sect. 4.1], see also [61, 62], to conclude that (4.5) admits a sparse solution consisting of at most m Dirac deltas

$$a_{m,\varepsilon} = \sum_{i=1}^m \alpha_i \delta_{K_i},$$

where $\{K_i\}_{i=1}^m \subset \mathbb{B}_{\mathcal{L}(\mathcal{X}; \mathcal{Y})}$ are linear operators and $\{\alpha_i\}_{i=1}^m \subset \mathbb{R}$ are real numbers. Sparsity of solutions of (4.5) has been analysed in detail in [63]. We would also like to mention the recent paper [38], where

scalar-valued variation norm spaces with a ReLU activation are shown to be reproducing kernel Banach spaces and a representer theorem is obtained.

We cannot immediately conclude, however, that this sparse solution is the same as in Theorem 3.24, for which Monte-Carlo rates in Bochner spaces L^p_π hold. The connection between empirically trained neural networks (i.e. those obtained by solving (4.5)) and the optimal ones featuring in existence-type theorems like Theorem 3.24 is, to the best of our knowledge, an open question.

5 Summary and conclusions

In the first part of the paper, we studied approximation properties of two-layer neural networks as non-linear operators acting between infinite-dimensional Banach spaces, extending existing finite-dimensional results. Perhaps, the most surprising result of this part is that we only obtain continuity of vector-valued \mathcal{F}_1 functions with respect to the weak-* topology on the output space \mathcal{Y} . This is unavoidable with our current techniques, but perhaps could be remedied with a different approach. Under this assumption, however, are were able to reproduce existing scalar-valued results on approximation rates. Our Monte-Carlo rates in L^p_π spaces for $1 \leq p < \infty$ improve, in the scalar-valued case, on known ones from [11] for all $p \neq 2$ (and coincide with them for $p = 2$).

In the second part of the paper, we studied the problem of finding the best representing measure of a variation norm function from a finite number of samples. We used variational regularisation and studied the behaviour of the regularised solution in the regime when both the number of training samples goes to infinity and the amount of noise in the samples vanishes. We obtained convergence rates in Bregman distance, which hold with high probability under the source condition, which in our case is a condition on the support of the measure that realises an \mathcal{F}_1 function. Informally, the smaller the norm of the source element, the more stable is the neural network representation of an \mathcal{F}_1 function in the sense that cancellations are avoided.

Acknowledgements

The author acknowledges financial support of the EPSRC (Fellowship EP/V003615/1), the Cantab Capital Institute for the Mathematics of Information at the University of Cambridge and the National Physical Laboratory. The author is grateful to Nik Weaver (Washington University in St. Louis) for providing the proof of Theorem 3.29. The author also thanks Jonas Latz (Heriot-Watt University) as well as two anonymous referees for useful remarks about the paper.

References

- [1] G. Cybenko. “Approximation by superpositions of a sigmoidal function”. In: *Mathematics of Control, Signals and Systems 2.4* (1989), pp. 303–314.
- [2] K. Hornik, M. Stinchcombe, and H. White. “Multilayer feedforward networks are universal approximators”. In: *Neural Networks 2.5* (1989), pp. 359–366.
- [3] M. Leshno, V. Y. Lin, A. Pinkus, and S. Schocken. “Multilayer feedforward networks with a nonpolynomial activation function can approximate any function”. In: *Neural Networks 6.6* (1993), pp. 861–867.
- [4] A. Barron. “Universal approximation bounds for superpositions of a sigmoidal function”. In: *IEEE Transactions on Information Theory 39.3* (1993), pp. 930–945.
- [5] R. A. DeVore. “Nonlinear approximation”. In: *Acta Numerica 7* (1998), 51–150.
- [6] V. Kurkova and M. Sanguinetti. “Bounds on rates of variable-basis and neural-network approximation”. In: *IEEE Transactions on Information Theory 47.6* (2001), pp. 2659–2665.
- [7] V. Kurkova and M. Sanguinetti. “Comparison of worst case errors in linear and neural network approximation”. In: *IEEE Transactions on Information Theory 48.1* (2002), pp. 264–275.
- [8] A. R. Barron, A. Cohen, W. Dahmen, and R. A. DeVore. “Approximation and learning by greedy algorithms”. In: *The Annals of Statistics 36.1* (2008), pp. 64–94.

- [9] F. Bach. “Breaking the Curse of Dimensionality with Convex Neural Networks”. In: *Journal of Machine Learning Research* 18.19 (2017), pp. 1–53.
- [10] W. E, C. Ma, and L. Wu. “Barron Spaces and the Compositional Function Spaces for Neural Network Models”. arxiv:1906.08039. 2019.
- [11] W. E and S. Wojtowytsch. “Representation formulas and pointwise properties for Barron functions”. arxiv:2006.05982. 2020.
- [12] G. Rotskoff and E. Vanden-Eijnden. “Parameters as interacting particles: long time convergence and asymptotic error scaling of neural networks”. In: *Advances in Neural Information Processing Systems*. Ed. by S. Bengio et al. Vol. 31. Curran Associates, Inc., 2018.
- [13] S. Mei, A. Montanari, and P.-M. Nguyen. “A mean field view of the landscape of two-layer neural networks”. In: *Proceedings of the National Academy of Sciences* 115.33 (2018), E7665–E7671.
- [14] L. Chizat and F. Bach. “On the Global Convergence of Gradient Descent for Over-parameterized Models using Optimal Transport”. In: *Advances in Neural Information Processing Systems*. Ed. by S. Bengio et al. Vol. 31. Curran Associates, Inc., 2018.
- [15] J. Sirignano and K. Spiliopoulos. “Mean Field Analysis of Neural Networks: A Law of Large Numbers”. In: *SIAM Journal on Applied Mathematics* 80.2 (2020), pp. 725–752.
- [16] A. Rahimi and B. Recht. “Random Features for Large-Scale Kernel Machines”. In: *Advances in Neural Information Processing Systems*. Ed. by J. Platt, D. Koller, Y. Singer, and S. Roweis. Vol. 20. Curran Associates, Inc., 2008.
- [17] C. Domingo-Enrich, A. Bietti, E. Vanden-Eijnden, and J. Bruna. “On Energy-Based Models with Overparametrized Shallow Neural Networks”. arXiv:2104.07531. 2021.
- [18] S. R. Arridge, P. Maass, O. Öktem, and C.-B. Schönlieb. “Solving inverse problems using data-driven models”. In: *Acta Numerica* 28 (2019), pp. 1–174.
- [19] M. T. McCann, K. H. Jin, and M. Unser. “Convolutional Neural Networks for Inverse Problems in Imaging: A Review”. In: *IEEE Signal Processing Magazine* 34.6 (2017), pp. 85–95.
- [20] G. Kutyniok, P. Petersen, M. Raslan, and R. Schneider. “A Theoretical Analysis of Deep Neural Networks and Parametric PDEs”. arXiv:1904.00377. 2019.
- [21] N. H. Nelsen and A. M. Stuart. “The Random Feature Model for Input-Output Maps between Banach Spaces”. arXiv:2005.10224. 2020.
- [22] T. Chen and H. Chen. “Universal approximation to nonlinear operators by neural networks with arbitrary activation functions and its application to dynamical systems”. In: *IEEE Transactions on Neural Networks* 6.4 (1995), pp. 911–917.
- [23] L. Lu et al. “Learning nonlinear operators via DeepONet based on the universal approximation theorem of operators”. In: *Nature Machine Intelligence* 3.3 (2021), pp. 218–229.
- [24] S. Lanthaler, S. Mishra, and G. E. Karniadakis. “Error estimates for DeepOnets: A deep learning framework in infinite dimensions”. arXiv:2102.09618. 2021.
- [25] K. Bhattacharya, B. Hosseini, N. B. Kovachki, and A. M. Stuart. “Model Reduction and Neural Networks for Parametric PDEs”. arXiv:2005.03180. 2020.
- [26] Z. Li et al. “Fourier Neural Operator for Parametric Partial Differential Equations”. arXiv:2010.08895. 2020.
- [27] Z. Li et al. “Neural Operator: Graph Kernel Network for Partial Differential Equations”. arXiv:2003.03485. 2020.
- [28] C. A. Micchelli and M. Pontil. “On Learning Vector-Valued Functions”. In: *Neural Computation* 17.1 (2005), pp. 177–204.
- [29] H. Zhang and J. Zhang. “Vector-valued reproducing kernel Banach spaces with applications to multi-task learning”. In: *Journal of Complexity* 29.2 (2013), pp. 195–215.
- [30] M. A. Álvarez, L. Rosasco, and N. D. Lawrence. “Kernels for Vector-Valued Functions: A Review”. In: *Foundations and Trends in Machine Learning* 4.3 (2012), pp. 195–266.
- [31] P. Meyer-Nieberg. *Banach Lattices*. Springer, 1991.

- [32] M. Burger and H. W. Engl. “Training neural networks with noisy data as an ill-posed problem”. In: *Advances in Computational Mathematics* 13.4 (2000), pp. 335–354.
- [33] E. D. Vito et al. “Learning from Examples as an Inverse Problem”. In: *Journal of Machine Learning Research* 6.30 (2005), pp. 883–904.
- [34] D. Soudry et al. “The Implicit Bias of Gradient Descent on Separable Data”. In: *Journal of Machine Learning Research* 19.70 (2018), pp. 1–57.
- [35] B. Jin and X. Lu. “On the regularizing property of stochastic gradient descent”. In: *Inverse Problems* 35.1 (2018), p. 015004.
- [36] L. Chizat and F. Bach. “Implicit Bias of Gradient Descent for Wide Two-layer Neural Networks Trained with the Logistic Loss”. In: *Proceedings of Thirty Third Conference on Learning Theory*. Ed. by J. Abernethy and S. Agarwal. Vol. 125. Proceedings of Machine Learning Research. PMLR, 2020, pp. 1305–1338.
- [37] R. DeVore, B. Hanin, and G. Petrova. “Neural Network Approximation”. arXiv:2012.14501. 2021.
- [38] R. Parhi and R. D. Nowak. “Banach Space Representer Theorems for Neural Networks and Ridge Splines”. In: *Journal of Machine Learning Research* 22.43 (2021), pp. 1–40.
- [39] A. Krogh and J. Hertz. “A Simple Weight Decay Can Improve Generalization”. In: *Advances in Neural Information Processing Systems*. Ed. by J. Moody, S. Hanson, and R. Lippmann. Vol. 4. Morgan-Kaufmann, 1991.
- [40] K. Bredies and H. K. Pikkarainen. “Inverse problems in spaces of measures”. In: *ESAIM: Control, Optimisation and Calculus of Variations* 19.1 (2013), 190–218.
- [41] V. Duval and G. Peyré. “Exact Support Recovery for Sparse Spikes Deconvolution”. In: *Foundations of Computational Mathematics* 15.5 (2015), pp. 1315–1355.
- [42] C. Poon, N. Keriven, and G. Peyré. “Support Localization and the Fisher Metric for off-the-grid Sparse Regularization”. In: *Proceedings of the Twenty-Second International Conference on Artificial Intelligence and Statistics*. Ed. by K. Chaudhuri and M. Sugiyama. Vol. 89. Proceedings of Machine Learning Research. PMLR, 2019, pp. 1341–1350.
- [43] L. Györfi, M. Kohler, A. Krzyzak, and H. Walk. *A Distribution-Free Theory of Nonparametric Regression*. New York: Springer, 2002.
- [44] G. Blanchard and N. Mücke. “Optimal Rates for Regularization of Statistical Inverse Learning Problems”. In: *Foundations of Computational Mathematics* 18.4 (2018), pp. 971–1013.
- [45] T. Bubba, M. Burger, T. Helin, and L. Ratti. “Convex regularization in statistical inverse learning problems”. arxiv:2102.09526. 2021.
- [46] M. Benning and M. Burger. “Modern Regularization Methods for Inverse Problems”. In: *Acta Numerica* 27 (2018), pp. 1–111.
- [47] N. Weaver. *Lipschitz algebras*. 2nd ed. World Scientific Publishing, 2018.
- [48] J. Diestel and J. J. Uhl. *Vector measures*. American Mathematical Society, 1977.
- [49] R. A. Ryan. *Introduction to Tensor Products of Banach Spaces*. Springer, 2002.
- [50] N. Dunford and J. T. Schwartz. *Linear Operators, Part I General Theory*. Hoboken, NJ: Interscience Publishers, 1958.
- [51] P. Enflo. “A counterexample to the approximation problem in Banach spaces”. In: *Acta Math.* 130 (1973), pp. 309–317.
- [52] V. Bogachev. *Measure theory*. Vol. 2. Springer, 2007.
- [53] Z. Chen and A. Wickstead. “Relative weak compactness of solid hulls in Banach lattices”. In: *Indagationes Mathematicae* 9.2 (1998), pp. 187–196.
- [54] Z. L. Chen and A. W. Wickstead. “Some Applications of Rademacher Sequences in Banach Lattices”. In: *Positivity* 2.2 (1998), pp. 171–191.
- [55] O. Scherzer et al. *Variational Methods in Imaging*. Springer, 2009.
- [56] M. Burger and S. Osher. “Convergence rates of convex variational regularization”. In: *Inverse Problems* 20.5 (2004), p. 1411.

- [57] L. Bungert, M. Burger, Y. Korolev, and C.-B. Schönlieb. “Variational regularisation for inverse problems with imperfect forward operators and general noise models”. In: *Inverse Problems* 36.12 (2020), p. 125014.
- [58] C.-A. Deledalle, N. Papadakis, and J. Salmon. “On Debiasing Restoration Algorithms: Applications to Total-Variation and Nonlocal-Means”. In: *Scale Space and Variational Methods in Computer Vision*. Ed. by J.-F. Aujol, M. Nikolova, and N. Papadakis. Cham: Springer International Publishing, 2015, pp. 129–141.
- [59] E.-M. Brinkmann, M. Burger, J. Rasch, and C. Sutour. “Bias Reduction in Variational Regularization”. In: *Journal of Mathematical Imaging and Vision* 59.3 (2017), pp. 534–566.
- [60] M. Unser. “A Unifying Representer Theorem for Inverse Problems and Machine Learning”. In: *Foundations of Computational Mathematics* (2020).
- [61] C. Boyer et al. “On Representer Theorems and Convex Regularization”. In: *SIAM Journal on Optimization* 29.2 (2019), pp. 1260–1281.
- [62] K. Bredies and M. Carioni. “Sparsity of solutions for variational inverse problems with finite-dimensional data”. In: *Calculus of Variations and Partial Differential Equations* 59.1 (2019), p. 14.
- [63] J. de Dios and J. Bruna. “On Sparsity in Overparametrised Shallow ReLU Networks”. arxiv:2006.10225. 2020.

A weak-* compactness of the unit ball in the space of linear bounded operators

Proposition A.1 ([49, Sec.4.1, Cor. 4.8]). *Let $E = (E^\diamond)^*$ and $F = (F^\diamond)^*$ be Banach spaces such that either E or F^\diamond has the approximation property. Then there exists an isometric isomorphism between $(\mathcal{N}(E^\diamond; F^\diamond))^*$ and $\mathcal{L}(E; F)$. For any $K \in \mathcal{L}(E; F)$ and $N \in \mathcal{N}(E^\diamond; F^\diamond)$, the duality pairing can be written as follows*

$$\langle N, K \rangle = \text{Tr}(KN^*) := \sum_{i=1}^{\infty} \langle \eta_i, Kx_i \rangle, \quad (\text{A.1})$$

where $\{x_i, \eta_i\}_{i \in \mathbb{N}} \subset E \times F^\diamond$ is any nuclear representation (3.5) of $N \in \mathcal{N}(E^\diamond; F^\diamond)$.

Using the Banach-Alaoglu theorem, we obtain

Proposition A.2. *Suppose that E, F are the duals of separable spaces E^\diamond and F^\diamond and that either E or F^\diamond has the approximation property. Then the unit ball $\mathbb{B}_{\mathcal{L}(E; F)}$ is weakly-* compact and metrisable.*

Proof. Since every nuclear operator is a limit of finite-rank operators, $\mathcal{N}(E^\diamond; F^\diamond)$ is separable and by Proposition A.1 we have that $\mathcal{L}(E; F)$ is the dual of a separable Banach space. The (sequential) Banach-Alaoglu theorem implies that the unit ball $\mathbb{B}_{\mathcal{L}(E; F)}$ is weakly-* compact and metrisable. \square

B Proof of Proposition 4.2

Proof. Consider the unit ball $\mathbb{B}_{\mathcal{M}}$ in $\mathcal{M}(\mathbb{B}_{\mathcal{L}(\mathcal{X}; \mathcal{Y})})$ and its image $T\mathbb{B}_{\mathcal{M}}$. Take an arbitrary sequence $\{f_k\}_{k \in \mathbb{N}} \subset T\mathbb{B}_{\mathcal{M}}$. It corresponds to a sequence $\{a_k\}_{k \in \mathbb{N}} \subset \mathbb{B}_{\mathcal{M}}$ such that $f_k = \langle \sigma(B \cdot), a_k \rangle$. Since $\{a_k\}_{k \in \mathbb{N}}$ is bounded, it contains a weakly-* converging subsequence (which we don’t relabel)

$$a_k \rightharpoonup^* a \quad \text{in } \mathbb{B}_{\mathcal{M}}.$$

Fix $x \in \mathcal{X}$ and consider

$$\|f(x) - f_k(x)\|_{(\mathcal{Y}, d_*)} = \sum_{i=1}^{\infty} 2^{-i} |\langle \eta_i, \langle \sigma(Bx), a - a_k \rangle \rangle| = \sum_{i=1}^{\infty} 2^{-i} |\langle \langle \eta_i, \sigma(Bx) \rangle, a - a_k \rangle|,$$

where $\{\eta_i\}_{i=1}^{\infty} \subset \mathcal{Y}^\diamond$ is a countable dense system in \mathcal{Y}^\diamond as defined in Proposition 3.9. Since the functions $\langle \eta_i, \sigma(Bx) \rangle$ are continuous on $(\mathbb{B}_{\mathcal{L}(\mathcal{X}; \mathcal{Y})}, d_*)$, we have $\langle \langle \eta_i, \sigma(Bx) \rangle, a - a_k \rangle \rightarrow 0$ for all i . We also have the following bound

$$\sum_{i=1}^{\infty} 2^{-i} |\langle \langle \eta_i, \sigma(Bx) \rangle, a - a_k \rangle| \leq \sum_{i=1}^{\infty} 2^{-i} \|a - a_k\|_{\mathcal{M}} \|x\| \leq C \|x\|, \quad (\text{B.1})$$

where we used the facts that $\|\eta_i\| = 1$, $\|B\| \leq 1$ and $\|a - a_k\|_{\mathcal{M}} \leq C$ since weakly- $*$ convergent sequences are bounded. By the dominated convergence theorem we obtain that $\|f(x) - f_k(x)\|_{(\mathcal{Y}, d_*)} \rightarrow 0$ pointwise on \mathcal{X} . Since π has $p \geq 1$ finite moments, the function $x \mapsto \|x\|$ is in L^p_π and by (B.1) and the dominated convergence theorem we get that $f_k \rightarrow f$ strongly in $L^p_\pi(\mathcal{X}; (\mathcal{Y}, d_*))$. Since $\{f_k\}_{k \in \mathbb{N}}$ was arbitrary, every sequence in $T\mathbb{B}_{\mathcal{M}}$ has a convergent sequence and $T\mathbb{B}_{\mathcal{M}}$ is (sequentially) compact, hence T is a compact operator. \square