

Function-Space Optimality of Neural Architectures with Multivariate Nonlinearities*

Rahul Parhi[†] and Michael Unser[‡]

Abstract. We investigate the function-space optimality (specifically, the Banach-space optimality) of a large class of shallow neural architectures with multivariate nonlinearities/activation functions. To that end, we construct a new family of Banach spaces defined via a regularization operator, the k -plane transform, and a sparsity-promoting norm. We prove a representer theorem that states that the solution sets to learning problems posed over these Banach spaces are completely characterized by neural architectures with multivariate nonlinearities. These optimal architectures have skip connections and are tightly connected to orthogonal weight normalization and multi-index models, both of which have received recent interest in the neural network community. Our framework is compatible with a number of classical nonlinearities including the rectified linear unit activation function, the norm activation function, and the radial basis functions found in the theory of thin-plate/polyharmonic splines. We also show that the underlying spaces are special instances of reproducing kernel Banach spaces and variation spaces. Our results shed light on the regularity of functions learned by neural networks trained on data, particularly with multivariate nonlinearities, and provide new theoretical motivation for several architectural choices found in practice.

Key words. multi-index models, multivariate nonlinearities, neural networks, regularization, representer theorem

MSC codes. 46E27, 47A52, 68T05, 82C32, 94A12

DOI. 10.1137/23M1620971

1. Introduction. In supervised machine learning, the goal is to predict an output $y \in \mathcal{Y}$ (e.g., a label or response) from an input $\mathbf{x} \in \mathcal{X}$ (e.g., a feature or example), where \mathcal{X} and \mathcal{Y} denote the domain of the inputs and outputs, respectively. One solves this task by “training” a model to fit a set of data which consists of a finite number of input-output pairs $\{(\mathbf{x}_m, y_m)\}_{m=1}^M \subset \mathcal{X} \times \mathcal{Y}$. The goal is to “learn” a function $f : \mathcal{X} \rightarrow \mathcal{Y}$ with $f(\mathbf{x}_m) \approx y_m$, $m = 1, \dots, M$, such that f can accurately predict the output $y \in \mathcal{Y}$ of a new input $\mathbf{x} \in \mathcal{X}$. This task is usually formulated as an optimization problem of the form

$$(1.1) \quad \min_{f \in \mathcal{F}} \sum_{m=1}^M \mathcal{L}(y_m, f(\mathbf{x}_m)) + \lambda \Phi(f),$$

*Received by the editors December 1, 2023; accepted for publication (in revised form) October 14, 2024; published electronically January 9, 2025.

<https://doi.org/10.1137/23M1620971>

Funding: This work was supported in part by the Swiss National Science Foundation under grant 200020_219356 / 1.

[†]Department of Electrical and Computer Engineering, University of California, San Diego, La Jolla, CA 92093 USA, and part of this work was done while the author was with the Biomedical Imaging Group, École polytechnique fédérale de Lausanne, CH-1015 Lausanne, Switzerland (rahul@ucsd.edu).

[‡]Biomedical Imaging Group, École polytechnique fédérale de Lausanne, CH-1015 Lausanne, Switzerland (michael.unser@epfl.ch).

where \mathcal{F} is a prescribed *model class* of functions that map $\mathcal{X} \rightarrow \mathcal{Y}$, $\mathcal{L}(\cdot, \cdot)$ is a *loss function*, and $\Phi: \mathcal{F} \rightarrow \mathbb{R}_{\geq 0}$ is a *regularization functional* that injects prior knowledge/regularity on the function to be learned. The hyperparameter $\lambda > 0$ controls the tradeoff between data fidelity and regularity. Without the inclusion of the regularization functional in (1.1), the problem is typically ill-posed. Indeed, in many practical scenarios the problem is *overparameterized* as the dimension of the model class \mathcal{F} greatly exceeds the number M of data. A classical choice of model class is a reproducing kernel Hilbert space (RKHS). The accompanying regularization functional is the squared Hilbert norm of the RKHS. In this scenario, the RKHS representer theorem establishes that there exists a solution to (1.1) that takes the form of a linear combination of reproducing kernels centered at the data sites [10, 29, 55, 70, 71]. This provides an exact characterization of the function-space optimality of kernel methods.

Recently, there has been a line of work that investigates the function-space optimality of neural networks [4, 5, 41, 42, 43, 44, 46, 54, 62, 66]. Crucially, these works define and study (non-Hilbertian) Banach spaces defined by sparsity or variation. These spaces have an analytic description via the Radon transform [30, 41, 43]. The accompanying neural network representer theorems for these spaces were first established in [43] and then studied and refined by a number of authors [5, 44, 62, 66]. While these results characterize the function-space optimality of neural networks, they only consider univariate nonlinearities. We refer the reader to the recent survey [45] for an up-to-date summary of this research direction. The purpose of this paper is to further extend the existing results on the function-space optimality of neural architectures, with a particular focus on *multivariate nonlinearities*, which have gained recent interest in the neural network community [2, 21, 23, 38].

The form of a neuron with an m -variate nonlinearity, $1 \leq m \leq d$, is

$$(1.2) \quad \mathbf{x} \mapsto \rho(\mathbf{A}\mathbf{x} - \mathbf{t}), \quad \mathbf{x} \in \mathbb{R}^d,$$

where $\rho: \mathbb{R}^m \rightarrow \mathbb{R}$ is the nonlinearity (or activation function), $\mathbf{A} \in \mathbb{R}^{m \times d}$ is a weight matrix that controls the orientation of the neuron, and $\mathbf{t} \in \mathbb{R}^m$ is a bias which controls the offset of neuron. When $m = 1$, these atoms can be written as

$$(1.3) \quad \mathbf{x} \mapsto \rho(\boldsymbol{\alpha}^\top \mathbf{x} - t), \quad \mathbf{x} \in \mathbb{R}^d,$$

with $\rho: \mathbb{R} \rightarrow \mathbb{R}$, $\boldsymbol{\alpha} \in \mathbb{R}^d$, and $t \in \mathbb{R}$, which is the form of a classical neuron with a univariate nonlinearity. Neurons with m -variate nonlinearities as in (1.2) have been studied under many different names, including *m -sparse functions* [1, 18], *generalized ridge functions* [28], *$(d-m)$ -plane ridge functions* [47], and *multi-index models* [8, 9, 15, 32, 35]. Notably, multi-index models have gained recent interest from the neural network community [1, 4, 18, 48]. A shallow neural architecture with an m -variate nonlinearity $\rho: \mathbb{R}^m \rightarrow \mathbb{R}$ takes the form

$$(1.4) \quad \mathbf{x} \mapsto \sum_{n=1}^N v_n \rho(\mathbf{A}_n \mathbf{x} - \mathbf{t}_n), \quad \mathbf{x} \in \mathbb{R}^d,$$

where, for $n = 1, \dots, N$, $v_n \in \mathbb{R}$, $\mathbf{A}_n \in \mathbb{R}^{m \times d}$, and $\mathbf{t}_n \in \mathbb{R}^m$. Such architectures are sometimes called *generalized translation networks* [38, 39, 40] and are classically known to be universal

approximators if and only if¹ $\rho : \mathbb{R}^m \rightarrow \mathbb{R}$ is not a polynomial [39, Corollary 3.3]. These architectures have also recently been studied in the context of ridgelet analysis for a variety of shallow neural architectures [61]. In this paper, we characterize, for all integers m with $1 \leq m \leq d$, the function-space optimality of neural architectures with m -variate nonlinearities of the form (1.4), for a large class of nonlinearities. We show that these architectures are optimal solutions to data-fitting problems posed over (non-Hilbertian) Banach spaces defined via a sparsity-promoting norm in the domain of the k -plane transform. When $m = 1$, our framework is compatible with univariate nonlinearities and classical neural architectures, including the ReLU activation function. At the opposite extreme ($m = d$) our framework encompasses sparse kernel expansions and radial basis functions [3, 50, 63]. To the best of our knowledge, the results for $1 < m < d$ are new.

1.1. Main contributions and road map. Our results shed light on the regularity of the functions learned by neural networks trained on data. They provide new theoretical motivation for several architectural choices often found in practice, particularly with multivariate nonlinearities. These results hinge on recent developments regarding the distributional extension and invertibility of the k -plane transform and its dual [47]. The main contributions and organization of this paper are summarized in the remainder of this section.

New neural network Banach spaces. We propose and study the properties of a new family of *native spaces*, defined by²

$$(1.5) \quad \mathcal{M}_L^k(\mathbb{R}^d) := \left\{ f : \mathbb{R}^d \rightarrow \mathbb{R} \text{ is measurable : } \begin{array}{l} \|K_{d-k} \mathcal{R}_k L f\|_{\mathcal{M}} < \infty, \\ \text{ess sup}_{\mathbf{x} \in \mathbb{R}^d} |f(\mathbf{x})| (1 + \|\mathbf{x}\|_2)^{-n_L} < \infty \end{array} \right\} \subset \mathcal{S}'(\mathbb{R}^d),$$

where k is an integer such that $0 \leq k < d$, $\mathcal{S}'(\mathbb{R}^d)$ denotes the space of tempered distributions, L is a k -plane-admissible pseudodifferential operator (in the sense of Definition 3.1), \mathcal{R}_k denotes the k -plane transform, and K_{d-k} is the filtering operator of computed tomography (CT) which is such that $\mathcal{R}_k^* K_{d-k} \mathcal{R}_k = \text{Id}$. The \mathcal{M} -norm denotes the total variation norm (in the sense of measures). It can be viewed as a “generalization” of the L^1 -norm that can also be applied to distributions such as the Dirac impulse. Said differently, if $f \in \mathcal{M}_L^k(\mathbb{R}^d)$ is such that $K_{d-k} \mathcal{R}_k L f$ is a bona fide function (not a distribution), then

$$(1.6) \quad \|K_{d-k} \mathcal{R}_k L f\|_{\mathcal{M}} = \|K_{d-k} \mathcal{R}_k L f\|_{L^1}.$$

The growth restriction of degree n_L plays the role of a proxy to the order of L ; more specifically, n_L is the highest polynomial degree annihilated by L . The growth restriction in the definition of the native space ensures that the null space of the operator $K_{d-k} \mathcal{R}_k L$ is finite-dimensional. In subsection 3.1, we prove that, when equipped with an appropriate direct-sum topology, $\mathcal{M}_L^k(\mathbb{R}^d)$ forms a Banach space that is isometrically isomorphic to the Cartesian product of a space of (Radon) measures with the space of polynomials of degree at most n_L . They add to the growing list of “neural Banach spaces” that are currently being actively investigated [57].

¹This equivalence holds under the global assumption that $\rho : \mathbb{R}^m \rightarrow \mathbb{R}$ does not grow faster than a polynomial, i.e., it is a *tempered* function.

²We refer to a function as *measurable* when it is measurable with respect to the Lebesgue σ -algebra.

Representer theorems for neural networks with multivariate nonlinearities. We prove a representer theorem (Theorem 3.8) that states that, under mild assumptions on the loss function and L , the solution set to the optimization problem

$$(1.7) \quad \min_{f \in \mathcal{M}_L^k(\mathbb{R}^d)} \sum_{m=1}^M \mathcal{L}(y_m, f(\mathbf{x}_m)) + \lambda \|K_{d-k} \mathcal{R}_k L f\|_{\mathcal{M}}$$

is completely characterized by shallow neural architectures with $(d-k)$ -variate activation functions matched to the operator L and widths bounded by the number M of data (independent of the dimension d of the data). This result sheds light on the role of biases, skip connections, and the use of structured weight matrices in neural architectures. Indeed, these architectures take the form

$$(1.8) \quad \mathbf{x} \mapsto c(\mathbf{x}) + \sum_{n=1}^N v_n \rho_L(\mathbf{A}_n \mathbf{x} - \mathbf{t}_n)$$

with $N \leq M$, where, for $n = 1, \dots, N$, $\mathbf{A}_n \in \mathbb{R}^{(d-k) \times d}$ is such that $\mathbf{A}_n \mathbf{A}_n^\top = \mathbf{I}_{d-k}$ (identity matrix), $\mathbf{t}_n \in \mathbb{R}^{d-k}$, and $v_n \in \mathbb{R} \setminus \{0\}$. The function c is a polynomial of degree at most n_L and the function $\rho_L: \mathbb{R}^{d-k} \rightarrow \mathbb{R}$ is a $(d-k)$ -variate nonlinearity matched to the operator L . Finally, the regularization cost of (1.8) is $\sum_{n=1}^N |v_n| = \|\mathbf{v}\|_1$. The term $\mathbf{x} \mapsto c(\mathbf{x})$ that appears in (1.8) can be viewed as a (generalized) skip connection in neural network parlance. Note that (1.8) is exactly a sparse combination of multi-index models with learnable orientations \mathbf{A}_n , offsets \mathbf{t}_n , and fixed profiles specified by the multivariate nonlinearity ρ_L as well as a generalized translation network as in (1.4). Thus, if the data lie on a low-dimensional subspace (or union of subspaces), the neural architecture could automatically adapt to this structure and avoid the curse of dimensionality.

Connections to prior work. In section 4, we instantiate our results on the function-space optimality of neural architectures. First, we discuss implications of our representer theorem to the training and regularization of neural networks. These results provide new insight into the role of overparameterization and the use of *orthogonal weight normalization* in network architectures, which corresponds to the property that $\mathbf{A}_n \mathbf{A}_n^\top = \mathbf{I}_{d-k}$ in (1.8) [2, 24, 25, 33]. This property has been shown to increase the stability [2] and generalization properties [25] of neural architectures. We then discuss specific examples of neural architectures that are compatible with our framework. These architectures include the popular ReLU [19], the norm activation function/nonlinearity [23], and the radial basis functions found in the theory of thin-plate/polyharmonic splines [12, 71]. In particular, our theory provides a way to interpolate between the completely anisotropic atoms found in neural architectures with univariate nonlinearities ($k = (d-1)$) to the completely isotropic atoms found in the theory of sparse kernel expansions and radial basis functions ($k = 0$) in a similar vein to how α -molecules interpolate between ridgelets (anisotropic) and wavelets (isotropic) [22].

In section 5, we discuss how the native space $\mathcal{M}_L^k(\mathbb{R}^d)$ can be viewed as an example of a reproducing kernel Banach space (RKBS) [5, 34, 62, 72] as well as an example of a variation space [4, 11, 31, 37, 56, 58, 57]. These are classical approaches used for the understanding of neural networks through the lens of functional analysis and approximation theory. Thus, any abstract result for RKBSs or variation spaces immediately applies to $\mathcal{M}_L^k(\mathbb{R}^d)$.

2. Mathematical preliminaries and notation. The Schwartz space of smooth and rapidly decreasing functions on \mathbb{R}^d is denoted by $\mathcal{S}(\mathbb{R}^d)$. Its continuous dual is the space $\mathcal{S}'(\mathbb{R}^d)$ of tempered distributions. We let $L^p(\mathbb{R}^d)$ denote the Lebesgue space for $1 \leq p \leq \infty$ and define the weighted L^∞ -space

$$(2.1) \quad L_{-\alpha}^\infty(\mathbb{R}^d) := \left\{ f : \mathbb{R}^d \rightarrow \mathbb{R} \text{ is measurable} : \|f\|_{L_{-\alpha}^\infty} := \operatorname{ess\,sup}_{\mathbf{x} \in \mathbb{R}^d} |f(\mathbf{x})|(1 + \|\mathbf{x}\|_2)^{-\alpha} < \infty \right\}.$$

This is the space of growth-restricted functions with rate $\alpha \in \mathbb{R}$. It is a Banach space that can be identified as the continuous dual of the weighted L^1 -space

$$(2.2) \quad L_\alpha^1(\mathbb{R}^d) := \left\{ f : \mathbb{R}^d \rightarrow \mathbb{R} \text{ is measurable} : \|f\|_{L_\alpha^1} := \int_{\mathbb{R}^d} |f(\mathbf{x})|(1 + \|\mathbf{x}\|_2)^\alpha \, d\mathbf{x} < \infty \right\}.$$

The Banach space of continuous functions vanishing at $\pm\infty$ on \mathbb{R}^d equipped with the L^∞ -norm is denoted by $C_0(\mathbb{R}^d)$. By the Riesz–Markov–Kakutani representation theorem [14, Chapter 7], its continuous dual can be identified with the Banach space of finite Radon measures, denoted $\mathcal{M}(\mathbb{R}^d)$. Since $\mathcal{S}(\mathbb{R}^d)$ is dense in $C_0(\mathbb{R}^d)$, we have, by duality, that $\mathcal{M}(\mathbb{R}^d)$ is continuously embedded in $\mathcal{S}'(\mathbb{R}^d)$. Given a space \mathcal{X} and a norm $\|\cdot\|$, the completion of \mathcal{X} in $\|\cdot\|$ is a Banach space, denoted by $(\overline{\mathcal{X}}, \|\cdot\|)$. For example, we have, for $1 \leq p < \infty$, that $L^p(\mathbb{R}^d) = (\overline{\mathcal{S}(\mathbb{R}^d)}, \|\cdot\|_{L^p})$, and $C_0(\mathbb{R}^d) = (\overline{\mathcal{S}(\mathbb{R}^d)}, \|\cdot\|_{L^\infty})$.

The Fourier transform of $\varphi \in \mathcal{S}(\mathbb{R}^d)$ is defined as

$$(2.3) \quad \widehat{\varphi}(\boldsymbol{\xi}) := \mathcal{F}\{\varphi\}(\boldsymbol{\xi}) = \int_{\mathbb{R}^d} \varphi(\mathbf{x}) e^{-i\boldsymbol{\xi}^\top \mathbf{x}} \, d\mathbf{x}, \quad \boldsymbol{\xi} \in \mathbb{R}^d,$$

where $i^2 = -1$. Consequently, the inverse Fourier transform of $\widehat{\varphi} \in \mathcal{S}(\mathbb{R}^d)$ is given by

$$(2.4) \quad \mathcal{F}^{-1}\{\widehat{\varphi}\}(\mathbf{x}) = \frac{1}{(2\pi)^d} \int_{\mathbb{R}^d} \widehat{\varphi}(\boldsymbol{\xi}) e^{i\boldsymbol{\xi}^\top \mathbf{x}} \, d\boldsymbol{\xi}, \quad \mathbf{x} \in \mathbb{R}^d.$$

These operators are extended to act on $\mathcal{S}'(\mathbb{R}^d)$ by duality.

Any continuous linear shift-invariant (LSI) operator $L : \mathcal{S}(\mathbb{R}^d) \rightarrow \mathcal{S}'(\mathbb{R}^d)$ is a convolution operator specified by a unique kernel $h \in \mathcal{S}'(\mathbb{R}^d)$ such that $L\varphi = h * \varphi$. Such operators can also be specified in the Fourier domain by

$$(2.5) \quad L\varphi = \mathcal{F}^{-1}\{\widehat{L}\widehat{\varphi}\},$$

where $\widehat{L} \in \mathcal{S}'(\mathbb{R}^d)$ is the Fourier transform of the kernel $h \in \mathcal{S}'(\mathbb{R}^d)$. The tempered distribution h is the *impulse response* of L and the tempered distribution \widehat{L} is the Fourier symbol or *frequency response* of L . We shall generally use upright, roman letters for LSI operators and use the italic variant with a hat to denote its frequency response.

2.1. The k -plane transform. We are going to adopt the parameterization of the k -plane transform from [47, section 4]. There, the space of k -planes is parameterized by the Cartesian product of the Stiefel manifold with \mathbb{R}^{d-k} , where k is an integer such that $1 \leq k < d$. Let

$$(2.6) \quad V_{d-k}(\mathbb{R}^d) := \{\mathbf{A} \in \mathbb{R}^{(d-k) \times d} : \mathbf{A}\mathbf{A}^\top = \mathbf{I}_{d-k}\}$$

denote the Stiefel manifold. Then, the k -plane transform of $\varphi \in \mathcal{S}(\mathbb{R}^d)$ is defined as

$$(2.7) \quad \mathcal{R}_k\{\varphi\}(\mathbf{A}, \mathbf{t}) = \int_{\mathbb{R}^d} \varphi(\mathbf{x})\delta(\mathbf{A}\mathbf{x} - \mathbf{t}) \, d\mathbf{x}, \quad (\mathbf{A}, \mathbf{t}) \in (V_{d-k}(\mathbb{R}^d), \mathbb{R}^{d-k}),$$

where $\delta \in \mathcal{S}'(\mathbb{R}^{d-k})$ is the $(d-k)$ -variate Dirac impulse³ and the integral is understood as the action of $\delta(\mathbf{A}(\cdot) - \mathbf{t}) \in \mathcal{S}'(\mathbb{R}^d)$ on $\varphi \in \mathcal{S}(\mathbb{R}^d)$. Further, the dual transform (often called the “backprojection”) of $g \in L^\infty(V_{d-k}(\mathbb{R}^d) \times \mathbb{R}^{d-k})$ is given by

$$(2.8) \quad \mathcal{R}_k^*\{g\}(\mathbf{x}) = \int_{V_{d-k}(\mathbb{R}^d)} g(\mathbf{A}, \mathbf{A}\mathbf{x}) \, d\mathbf{A}, \quad \mathbf{x} \in \mathbb{R}^d,$$

where $d\mathbf{A}$ denotes integration against the Haar measure of $V_{d-k}(\mathbb{R}^d)$. Since we impose that the rows of \mathbf{A} are orthonormal, we have that (\mathbf{A}, \mathbf{t}) and $(\mathbf{U}\mathbf{A}, \mathbf{U}\mathbf{t})$ define the same k -plane, for any orthogonal transformation $\mathbf{U} \in O_{d-k}(\mathbb{R})$ (the orthogonal group in dimension $(d-k)$). The main advantage of the proposed parameterization is that it will allow us to identify the symmetries of k -plane domain as “isotropic” symmetries.

Letting $\Xi_k := V_{d-k}(\mathbb{R}^d) \times \mathbb{R}^{d-k}$ denote the k -plane domain, we define the space of Schwartz functions on Ξ_k , denoted by $\mathcal{S}(\Xi_k)$, as the space of smooth functions that are rapidly decreasing in the $\mathbf{t} \in \mathbb{R}^{d-k}$ variable [20]. More specifically, we have that $\mathcal{S}(\Xi_k) = C^\infty(V_{d-k}(\mathbb{R}^d)) \widehat{\otimes} \mathcal{S}(\mathbb{R}^{d-k})$, where $\widehat{\otimes}$ denotes the topological tensor product, which is the completion of the algebraic tensor product with respect to the projective topology [64, Chapter 43]. We state in Proposition 2.1 a classical result regarding the continuity and invertibility of the k -plane transform.

Proposition 2.1 (see [16, 20, 27, 47, 51, 59, 60]). *The operator \mathcal{R}_k continuously maps $\mathcal{S}(\mathbb{R}^d)$ into $\mathcal{S}(\Xi_k)$. Moreover,*

$$(2.9) \quad \mathcal{R}_k^* K_{d-k} \mathcal{R}_k = c_{d,k} (-\Delta)^{\frac{k}{2}} \mathcal{R}_k^* \mathcal{R}_k = c_{d,k} \mathcal{R}_k^* \mathcal{R}_k (-\Delta)^{\frac{k}{2}} = \text{Id}$$

on $\mathcal{S}(\mathbb{R}^d)$ with

$$(2.10) \quad c_{d,k} = \frac{1}{(2\pi)^k} \frac{|\mathbb{S}^{k-1}|}{|\mathbb{S}^{d-k-1}|} \frac{1}{\prod_{n=k}^{d-1} |\mathbb{S}^{n-1}|},$$

where $|\cdot|$ denotes the surface area. The underlying operators are the d -variate Laplacian operator Δ and the filtering operator⁴ $K_{d-k} = c_{d,k} (-\Delta_{d-k})^{k/2}$, where Δ_{d-k} denotes the $(d-k)$ -variate Laplacian applied to the $\mathbf{t} \in \mathbb{R}^{d-k}$ variable. The filtering operator is equivalently specified by the frequency response $\widehat{K}_{d-k}(\boldsymbol{\omega}) = c_{d,k} \|\boldsymbol{\omega}\|_2^k$, $\boldsymbol{\omega} \in \mathbb{R}^{d-k}$.

The k -plane transform has tight connections with the Fourier transform. This is summarized in the so-called *Fourier slice theorem*.

³The distribution $\delta \in \mathcal{S}'(\mathbb{R}^{d-k})$ is such that $\langle \delta, \phi \rangle = \phi(\mathbf{0})$ for all $\phi \in \mathcal{S}(\mathbb{R}^{d-k})$.

⁴In CT, this filter is referred to as the backprojection filter found in the filtered backprojection algorithm for CT image reconstruction.

Proposition 2.2 ([47, Corollary 7.5]). Given $\varphi \in \mathcal{S}(\mathbb{R}^d)$, we have that

$$(2.11) \quad \widehat{\mathcal{R}_k\{\varphi\}}(\mathbf{A}, \cdot)(\boldsymbol{\omega}) = \widehat{\varphi}(\mathbf{A}^\top \boldsymbol{\omega}), \quad \boldsymbol{\omega} \in \mathbb{R}^{d-k}, \mathbf{A} \in V_{d-k}(\mathbb{R}^d),$$

where the Fourier transform on the left-hand side is the $(d-k)$ -variate transform and the Fourier transform on the right-hand side is the d -variate Fourier transform.

Remark 2.3. The Fourier slice theorem can be extended to apply to members of $\mathcal{S}'(\mathbb{R}^d)$ as long as some additional care is taken regarding in what sense the equality in (2.11) holds [47, Theorem 7.7].

Let $\mathcal{S}_k := \mathcal{R}_k(\mathcal{S}(\mathbb{R}^d))$ denote the range of the k -plane transform. The range \mathcal{S}_k is a strict subspace of $\mathcal{S}(\Xi_k)$ that satisfies certain consistency conditions (see [36, Chapter 4] for a detailed discussion and references on this matter). We have the following additional result regarding the continuity and invertibility of the k -plane transform.

Proposition 2.4 ([47, Corollary 5.3]). The operator $\mathcal{R}_k : \mathcal{S}(\mathbb{R}^d) \rightarrow \mathcal{S}_k$ is a homeomorphism with inverse $\mathcal{R}_k^{-1} = \mathcal{R}_k^* \mathbf{K}_k : \mathcal{S}_k \rightarrow \mathcal{S}(\mathbb{R}^d)$.

Proposition 2.4 motivates the following *distributional extension* of the k -plane transform and related operators.

Definition 2.5 ([47, Definition 6.1]).

1. The distributional k -plane transform

$$(2.12) \quad \mathcal{R}_k : \mathcal{S}'(\mathbb{R}^d) \rightarrow (\mathbf{K}_{d-k} \mathcal{R}_k(\mathcal{S}(\mathbb{R}^d)))'$$

is defined to be the dual map of the homeomorphism $\mathcal{R}_k^* : \mathbf{K}_{d-k} \mathcal{R}_k(\mathcal{S}(\mathbb{R}^d)) \rightarrow \mathcal{S}(\mathbb{R}^d)$.

2. The distributional filtered k -plane transform

$$(2.13) \quad \mathbf{K}_{d-k} \mathcal{R}_k : \mathcal{S}'(\mathbb{R}^d) \rightarrow \mathcal{S}'_k$$

is defined to be the dual map of the homeomorphism $\mathcal{R}_k^* \mathbf{K}_{d-k} : \mathcal{S}_k \rightarrow \mathcal{S}(\mathbb{R}^d)$.

3. The distributional backprojection

$$(2.14) \quad \mathcal{R}_k^* : \mathcal{S}'_k \rightarrow \mathcal{S}'(\mathbb{R}^d)$$

is defined to be the dual map of the homeomorphism $\mathcal{R}_k : \mathcal{S}(\mathbb{R}^d) \rightarrow \mathcal{S}_k$.

4. The extended distributional backprojection

$$(2.15) \quad \mathcal{R}_k^* : \mathcal{S}'(\Xi_k) \rightarrow \mathcal{S}'(\mathbb{R}^d)$$

is defined to be the dual map of the continuous operator $\mathcal{R}_k : \mathcal{S}(\mathbb{R}^d) \rightarrow \mathcal{S}(\Xi_k)$, which is well-defined since \mathcal{S}_k is continuously embedded in $\mathcal{S}(\Xi_k)$.

Based on these definitions, we state in Theorem 2.6 a result on the invertibility of the filtered k -plane transform on $\mathcal{S}'(\mathbb{R}^d)$, which is the dual of Propositions 2.1 and 2.4.

Theorem 2.6. *It holds that $\mathcal{R}_k^* \mathbf{K}_{d-k} \mathcal{R}_k = \text{Id}$ on $\mathcal{S}'(\mathbb{R}^d)$. Moreover, the filtered k -plane transform $\mathbf{K}_{d-k} \mathcal{R}_k : \mathcal{S}'(\mathbb{R}^d) \rightarrow \mathcal{S}'_k$ is a homeomorphism with the inverse given by the backprojection $(\mathbf{K}_{d-k} \mathcal{R}_k)^{-1} = \mathcal{R}_k^* : \mathcal{S}'_k \rightarrow \mathcal{S}'(\mathbb{R}^d)$.*

While this setup provides an attractive formulation to handle the distributional extension of the k -plane transform and its dual, it turns out that the distribution spaces \mathcal{S}'_k and $(\mathbf{K}_{d-k} \mathcal{R}_k (\mathcal{S}(\mathbb{R}^d)))'$ are actually *equivalence classes* of distributions [47]. Luckily, by working with certain Banach subspaces that continuously embed into these distribution spaces, one can identify a concrete member of the equivalence classes via continuous projection operators. To this end, the results of this paper hinge on a nontrivial result of [47] regarding the invertibility of the distributional dual k -plane transform on the space of isotropic Radon measures.

Consider the operator

$$(2.16) \quad \mathbf{P}_{\text{iso}}\{g\}(\mathbf{A}, \mathbf{t}) = \int_{\mathbf{O}_{d-k}(\mathbb{R})} g(\mathbf{U}\mathbf{A}, \mathbf{U}\mathbf{t}) \, d\sigma(\mathbf{U}),$$

where \int denotes the average integral and σ is the Haar measure on $\mathbf{O}_{d-k}(\mathbb{R})$. This is a well-defined operator that maps $C_0(\Xi_k) \rightarrow C_0(\Xi_k)$ and, in particular, is the self-adjoint continuous projector which extracts the isotropic part of a function [47, equation (8.10)]. By the Riesz–Markov–Kakutani representation theorem, we can extend $\mathbf{P}_{\text{iso}} = \mathbf{P}_{\text{iso}}^*$ by duality to act on $\mathcal{M}(\Xi_k) = (C_0(\Xi_k))'$, the Banach space of finite Radon measures on Ξ_k . To this end, define

$$(2.17) \quad \mathcal{M}_{\text{iso}}(\Xi_k) := \mathbf{P}_{\text{iso}}(\mathcal{M}(\Xi_k)),$$

the Banach subspace of isotropic finite Radon measures on Ξ_k . This Banach subspace is complemented in $\mathcal{M}(\Xi_k)$ [47, Theorem 8.2] and so

$$(2.18) \quad \mathcal{M}(\Xi_k) = \mathcal{M}_{\text{iso}}(\Xi_k) \oplus (\mathcal{M}_{\text{iso}}(\Xi_k))^c$$

with $(\mathcal{M}_{\text{iso}}(\Xi_k))^c := (\text{Id} - \mathbf{P}_{\text{iso}})(\mathcal{M}(\Xi_k))$. We also note that

$$(2.19) \quad \mathcal{M}_{\text{iso}}(\Xi_k) = (C_{0,\text{iso}}(\Xi_k))',$$

where

$$(2.20) \quad C_{0,\text{iso}}(\Xi_k) := \mathbf{P}_{\text{iso}}(C_0(\Xi_k)) = \overline{(\mathcal{S}_k, \|\cdot\|_{L^\infty})},$$

where the last equality is from [47, equation (8.12)].

Proposition 2.7 (see [47, Theorems 8.1 and 8.2 and Corollary 8.3]). *The Banach space $\mathcal{M}_{\text{iso}}(\Xi_k)$ continuously embeds into \mathcal{S}'_k . Further, the distributional backprojection operator \mathcal{R}_k^* is invertible on $\mathcal{M}_{\text{iso}}(\Xi_k)$, so that*

$$(2.21) \quad \mathbf{K}_{d-k} \mathcal{R}_k \mathcal{R}_k^* = \text{Id} \text{ on } \mathcal{M}_{\text{iso}}(\Xi_k).$$

Furthermore, the null space of $\mathcal{R}_k^* : \mathcal{M}(\Xi_k) \rightarrow \mathcal{S}'(\mathbb{R}^d)$ is $(\mathcal{M}_{\text{iso}}(\Xi_k))^c$ and so $\mathcal{R}_k^*(\mathcal{M}(\Xi_k)) = \mathcal{R}_k^*(\mathcal{M}_{\text{iso}}(\Xi_k))$.

2.1.1. The case $k = 0$. When $k = 0$, the k -plane transform of $\varphi \in \mathcal{S}(\mathbb{R}^d)$ becomes

$$(2.22) \quad \mathcal{R}_0\{\varphi\}(\mathbf{A}, \mathbf{t}) = \int_{\mathbb{R}^d} \varphi(\mathbf{x}) \delta(\mathbf{A}\mathbf{x} - \mathbf{t}) \, d\mathbf{x} = \varphi(\mathbf{A}^\top \mathbf{t}).$$

Indeed, when $k = 0$, the matrix $\mathbf{A} \in \mathbb{R}^{d \times d}$ is now an orthogonal matrix (i.e., $\mathbf{A}^\top \mathbf{A} = \mathbf{A}\mathbf{A}^\top = \mathbf{I}_d$) and $\mathbf{t} \in \mathbb{R}^d$. The last equality of the above integral follows from the change of variables $\mathbf{y} = (\mathbf{A}\mathbf{x} - \mathbf{t})$. In this case, it is clear that the range \mathcal{S}_0 is exactly the closed subspace of isotropic functions in $\mathcal{S}(\Xi_0)$, denoted by $\mathcal{S}_{\text{iso}}(\Xi_0)$. The fundamental results discussed previously trivially hold in this limit setting. In particular, it is easy to check that

$$(2.23) \quad c_{d,0} \mathcal{R}_0^* \mathcal{R}_0 = \text{Id on } \mathcal{S}(\mathbb{R}^d),$$

which is the same statement as $\mathcal{R}_0^* \mathbf{K}_d \mathcal{R}_0 = \text{Id}$ since $\mathbf{K}_d = c_{d,0} \text{Id}$ (Proposition 2.1). Therefore, in the remainder of the paper, when working with the k -plane transform for general k such that $0 \leq k < d$, we shall not treat separately the case $k = 0$. We do warn the reader, however, to be aware that the underlying mathematics of the ($k = 0$)-plane transform is much simpler than that of $1 \leq k < d$.

2.2. Polynomial spaces and related projectors. The null space of the regularizer in the learning problem (1.7) is the space of polynomials of degree n_L (see Lemma 3.5), which depends on the operator L . This space is denoted $\mathcal{P}_{n_L}(\mathbb{R}^d)$. To this end, we will be interested in working with a biorthogonal system for this null space. We will use the biorthogonal system from [67, section 2.2].

Remark 2.8. When the null space of L is trivial, we make the identifications $n_L = (-1)$ and $\mathcal{P}_{n_L}(\mathbb{R}^d) = \{0\}$, noting that any sum taken from $n = 0$ to (-1) is understood as 0.

The space $\mathcal{P}_{n_L}(\mathbb{R}^d)$ is spanned by the monomial/Taylor basis

$$(2.24) \quad m_{\mathbf{n}}(\mathbf{x}) = \frac{\mathbf{x}^{\mathbf{n}}}{\mathbf{n}!},$$

where $\mathbf{n} = (n_1, \dots, n_d) \in \mathbb{N}_0^d$ is a multi-index. Accordingly, we have that

$$(2.25) \quad \mathcal{P}_{n_L}(\mathbb{R}^d) = \left\{ \sum_{|\mathbf{n}| \leq n_L} b_{\mathbf{n}} m_{\mathbf{n}} : b_{\mathbf{n}} \in \mathbb{R} \right\} \subset \mathcal{S}'(\mathbb{R}^d).$$

Importantly, $\mathcal{P}_{n_L}(\mathbb{R}^d)$ forms a finite-dimensional Banach subspace of $\mathcal{S}'(\mathbb{R}^d)$. Since all norms are equivalent in finite dimensions, the exact choice does not matter, but, for concreteness, we equip $\mathcal{P}_{n_L}(\mathbb{R}^d)$ with the norm

$$(2.26) \quad \|p\|_{\mathcal{P}_{n_L}} := \|(b_{\mathbf{n}})_{|\mathbf{n}| \leq n_L}\|_2,$$

where the $b_{\mathbf{n}}$ are the coefficients of p in the monomial/Taylor basis. While the dual space $\mathcal{P}'_{n_L}(\mathbb{R}^d)$ is also finite-dimensional, its “abstract” elements are actually equivalence classes in $\mathcal{S}'(\mathbb{R}^d)$. Following the approach from [67, section 2.2], we identify every dual element $p^* \in \mathcal{P}'_{n_L}(\mathbb{R}^d)$ as a function in $\mathcal{S}(\mathbb{R}^d)$ by working with a concrete dual basis $\{m_{\mathbf{n}}^*\}_{|\mathbf{n}| \leq n_L}$ that

satisfies the biorthogonality property $\langle m_{\mathbf{n}}^*, m_{\mathbf{n}'} \rangle = \delta[\mathbf{n} - \mathbf{n}']$, where $\delta[\cdot]$ is the Kronecker impulse which takes the value 1 when its input is 0 and 0 otherwise. Our specific choice is

$$(2.27) \quad m_{\mathbf{n}}^* := (-1)^{|\mathbf{k}|} \partial^{\mathbf{n}} \kappa_{\text{iso}} \in \mathcal{S}(\mathbb{R}^d),$$

where $\kappa_{\text{iso}} \in \mathcal{S}(\mathbb{R}^d)$ is an isotropic function constructed in [67, Lemma 1]. Its frequency response is such that $\widehat{\kappa}_{\text{iso}}(\boldsymbol{\xi}) = \widehat{\kappa}_{\text{rad}}(\|\boldsymbol{\xi}\|_2)$, where the radial frequency profile $\widehat{\kappa}_{\text{rad}} \in \mathcal{S}(\mathbb{R})$ satisfies $0 \leq \widehat{\kappa}_{\text{rad}} \leq 1$ and, for $|\omega| \geq 1$, $\widehat{\kappa}_{\text{rad}}(\omega) = 0$.

The biorthogonal system $\{(m_{\mathbf{n}}^*, m_{\mathbf{n}})\}_{|\mathbf{n}| \leq n_L}$ allows us to define the projection onto $\mathcal{P}_{n_L}(\mathbb{R}^d)$ by the operator

$$(2.28) \quad P_{\mathcal{P}_{n_L}(\mathbb{R}^d)}\{f\} = \sum_{\mathbf{n} \leq n_L} \langle m_{\mathbf{n}}^*, f \rangle m_{\mathbf{n}}.$$

This projector continuously maps $\mathcal{S}'(\mathbb{R}^d) \rightarrow \mathcal{P}_{n_L}(\mathbb{R}^d)$ (because $m_{\mathbf{n}}^* \in \mathcal{S}(\mathbb{R}^d)$). Moreover, since $L_{-n_L}^\infty(\mathbb{R}^d)$ continuously embeds into $\mathcal{S}'(\mathbb{R}^d)$, the restricted operator

$$(2.29) \quad P_{\mathcal{P}_{n_L}(\mathbb{R}^d)} : L_{-n_L}^\infty(\mathbb{R}^d) \rightarrow \mathcal{P}_{n_L}(\mathbb{R}^d)$$

is continuous as well. Last, the finite dimensionality of $\mathcal{P}_{n_L}(\mathbb{R}^d)$ ensures that $\mathcal{P}_{n_L}(\mathbb{R}^d)$ is complemented in $L_{-n_L}^\infty(\mathbb{R}^d)$ [52, Lemma 4.21]. Thus, the complementary projector

$$(2.30) \quad (\text{Id} - P_{\mathcal{P}_{n_L}(\mathbb{R}^d)}) : L_{-n_L}^\infty(\mathbb{R}^d) \rightarrow L_{-n_L}^\infty(\mathbb{R}^d)$$

is guaranteed to exist and be continuous.

3. Main results. We first define the class of operators that are admissible for the learning problem in (1.7). These operators form a subclass of the so-called *spline-admissible* operators [69, Definition 1].

Definition 3.1. A continuous LSI operator $L : \mathcal{S}(\mathbb{R}^d) \rightarrow \mathcal{S}'(\mathbb{R}^d)$ is said to be *k-plane-admissible with a polynomial null space of degree n_L* if

1. its adjoint L^* is a continuous injection that maps $\mathcal{S}(\mathbb{R}^d)$ into $L_{n_L}^1(\mathbb{R}^d)$;
2. it is isotropic in the sense that its frequency response is continuous and satisfies $\widehat{L}(\boldsymbol{\xi}) = \widehat{L}_{\text{rad}}(\|\boldsymbol{\xi}\|_2)$ for some continuous univariate radial frequency profile $\widehat{L}_{\text{rad}} : \mathbb{R} \rightarrow \mathbb{R}$;
3. the radial frequency profile \widehat{L}_{rad} does not vanish over \mathbb{R} , except for a zero of order $\gamma_L \in (n_L, n_L + 1]$ at the origin, so that there exists a constant $C > 0$ satisfying

$$(3.1) \quad \lim_{\omega \rightarrow 0} \frac{\widehat{L}_{\text{rad}}(\omega)}{|\omega|^{\gamma_L}} = C;$$

4. there exist $\gamma'_L > (d - k)$, $C' > 0$, and $R > 0$ such that

$$(3.2) \quad |\widehat{L}_{\text{rad}}(\omega)| \geq C' |\omega|^{\gamma'_L}$$

for all $|\omega| > R$.

Remark 3.2. Item 3 implies that the extension by duality $L : L^\infty_{-n_L}(\mathbb{R}^d) \rightarrow \mathcal{S}'(\mathbb{R}^d)$ of a k -plane-admissible operator annihilates polynomials of degree at most n_L .

Remark 3.3. Item 4 guarantees that the order of the operator is sufficiently high to ensure that the point evaluation functional is well-defined on the native space $\mathcal{M}_L^k(\mathbb{R}^d)$. In particular, this condition is reminiscent of the condition $s > (d - k)$ on the smoothness index of an L^1 -Sobolev space defined on (subsets of) \mathbb{R}^{d-k} to ensure the continuity of its members. Indeed, when $s > (d - k)$, the Sobolev embedding theorem guarantees that the L^1 -Sobolev space of order s embeds into the space of continuous functions. In our setting, the condition $\gamma'_L > (d - k)$ plays the role of the smoothness index. This property is key to establishing the existence of solutions to the learning problem (1.7).

3.1. Native spaces. The primary technical contribution of this paper is the careful treatment of the distributional extension and (pseudo)invertibility of the operator $L_{\mathcal{R}_k} := K_{d-k} \mathcal{R}_k L$, where L is a k -plane-admissible operator (Definition 3.1). This in turn allows us to define our native (Banach) spaces. Indeed, for the definition of the native space in (1.5) to be coherent, the action of $L_{\mathcal{R}_k}$ on $L^\infty_{-n_L}(\mathbb{R}^d)$ must be well-defined. Furthermore, as we shall see in Lemma 3.5, the null space of $L_{\mathcal{R}_k}$ is exactly the space of polynomials of degree at most n_L . We then use a technique from spline theory to “factor out” the null space of $L_{\mathcal{R}_k}$ and identify the subspace of $\mathcal{M}_L^k(\mathbb{R}^d)$ on which $L_{\mathcal{R}_k}$ is invertible [10, 12]. This allows us to identify $\mathcal{M}_L^k(\mathbb{R}^d)$ as the direct sum of two Banach spaces. Thus, it forms a Banach space when equipped with the composite norm.

Lemma 3.4. *Let L be a k -plane-admissible operator in the sense of Definition 3.1. Then, the operator $L_{\mathcal{R}_k} = K_{d-k} \mathcal{R}_k L$ continuously maps $L^\infty_{-n_L}(\mathbb{R}^d) \rightarrow \mathcal{S}'_k$.*

Proof. It suffices to prove that the adjoint operator $L_{\mathcal{R}_k}^* = L^* \mathcal{R}_k^* K_{d-k}$ continuously maps $\mathcal{S}_k \rightarrow L^1_{n_L}(\mathbb{R}^d)$. The result then follows by duality. Since $\mathcal{R}_k^* K_{d-k} : \mathcal{S}_k \rightarrow \mathcal{S}(\mathbb{R}^d)$ is a homeomorphism (Proposition 2.4), the lemma follows from item 1 in Definition 3.1. ■

Lemma 3.5. *Let L be a k -plane-admissible operator in the sense of Definition 3.1. Then, the null space of the operator $L_{\mathcal{R}_k} = K_{d-k} \mathcal{R}_k L : L^\infty_{-n_L}(\mathbb{R}^d) \rightarrow \mathcal{S}'_k$ is the space of polynomials of degree at most n_L on \mathbb{R}^d , denoted by $\mathcal{P}_{n_L}(\mathbb{R}^d)$.*

Proof. Let $\mathcal{N}(L_{\mathcal{R}_k})$ denote the null space of $L_{\mathcal{R}_k}$. More precisely,

$$(3.3) \quad \mathcal{N}(L_{\mathcal{R}_k}) = \{f \in L^\infty_{-n_L}(\mathbb{R}^d) : L_{\mathcal{R}_k} \{f\} = 0\}.$$

From Remark 3.2, L annihilates polynomials of degree at most n_L . This implies that $\mathcal{N}(L_{\mathcal{R}_k}) \supset \mathcal{P}_{n_L}(\mathbb{R}^d)$. However, given $f \in L^\infty_{-n_L}(\mathbb{R}^d) \subset \mathcal{S}'(\mathbb{R}^d)$, from the Fourier slice theorem,

$$(3.4) \quad \widehat{L_{\mathcal{R}_k} \{f\}}(\mathbf{A}, \cdot)(\boldsymbol{\omega}) = c_{d,k} \|\boldsymbol{\omega}\|_2^k \widehat{L}(\mathbf{A}^\top \boldsymbol{\omega}) \widehat{f}(\mathbf{A}^\top \boldsymbol{\omega}) = c_{d,k} \|\boldsymbol{\omega}\|_2^k \widehat{L}_{\text{rad}}(\|\boldsymbol{\omega}\|_2) \widehat{f}(\mathbf{A}^\top \boldsymbol{\omega}).$$

Note that the product in the last equality must be a well-defined tempered distribution (in the $\boldsymbol{\omega}$ variable) since $L_{\mathcal{R}_k} : L^\infty_{-n_L}(\mathbb{R}^d) \rightarrow \mathcal{S}'_k$ is well-defined by Lemma 3.4. Due to the vanishing property in item 3 of Definition 3.1, this quantity is 0 if and only if \widehat{f} is supported only at $\mathbf{0}$, in which case f is a polynomial. Combined with the growth restriction $f \in L^\infty_{-n_L}(\mathbb{R}^d)$, we have that f must be a polynomial of degree at most n_L . Therefore, $\mathcal{N}(L_{\mathcal{R}_k}) \subset \mathcal{P}_{n_L}(\mathbb{R}^d)$. ■

From Lemmas 3.4 and 3.5, we deduce that the native space in (1.5) is well-defined and that the null space of $L_{\mathcal{R}_k}$ is the finite-dimensional Banach space $\mathcal{P}_{n_L}(\mathbb{R}^d)$. The next two technical theorems (Theorems 3.6 and 3.7) establish the Banach structure of the native space.

Theorem 3.6. *Let L be a k -plane-admissible operator in the sense of Definition 3.1. Then, the operator $L_{\mathcal{R}_k} = K_{d-k} \mathcal{R}_k L$ maps $\mathcal{M}_L^k(\mathbb{R}^d) \rightarrow \mathcal{M}_{\text{iso}}(\Xi_k)$. Furthermore, there exists an operator $L_{\mathcal{R}_k}^\dagger$ that continuously maps $\mathcal{M}_{\text{iso}}(\Xi_k) \rightarrow L_{-n_L}^\infty(\mathbb{R}^d) \subset \mathcal{S}'(\mathbb{R}^d)$ and is such that*

$$(3.5) \quad L_{\mathcal{R}_k} L_{\mathcal{R}_k}^\dagger \{\mu\} = \mu \text{ for all } \mu \in \mathcal{M}_{\text{iso}}(\Xi_k),$$

$$(3.6) \quad L_{\mathcal{R}_k}^\dagger L_{\mathcal{R}_k} \{f\} = (\text{Id} - P_{\mathcal{P}_{n_L}(\mathbb{R}^d)})\{f\} \text{ for all } f \in \mathcal{M}_L^k(\mathbb{R}^d).$$

This operator is realized by

$$(3.7) \quad L_{\mathcal{R}_k}^\dagger = (\text{Id} - P_{\mathcal{P}_{n_L}(\mathbb{R}^d)}) L^{-1} \mathcal{R}_k^*,$$

where L^{-1} is the operator specified by the frequency response $\boldsymbol{\xi} \mapsto 1/\widehat{L}(\boldsymbol{\xi})$. Furthermore, $L_{\mathcal{R}_k}^\dagger$ is an integral operator specified by the kernel $(\mathbf{A}, \mathbf{t}) \mapsto g_{\mathbf{A}, \mathbf{t}}(\mathbf{x})$ that takes the form

$$(3.8) \quad g_{\mathbf{A}, \mathbf{t}}(\mathbf{x}) = \rho_L(\mathbf{A}\mathbf{x} - \mathbf{t}) - \sum_{|\mathbf{k}| \leq n_L} \langle m_{\mathbf{k}}^*, \rho_L(\mathbf{A}(\cdot) - \mathbf{t}) \rangle m_{\mathbf{k}}(\mathbf{x}),$$

where $\rho_L = \mathcal{F}_{d-k}^{-1} \{1/\widehat{L}_{\text{rad}}(\|\cdot\|_2)\}$ and $\{(m_{\mathbf{n}}^*, m_{\mathbf{n}})\}_{|\mathbf{n}| \leq n_L}$ are specified in subsection 2.2, and where \mathcal{F}_{d-k}^{-1} denotes the $(d-k)$ -variate inverse Fourier transform. This kernel satisfies the stability/continuity bound

$$(3.9) \quad \sup_{\substack{\mathbf{x} \in \mathbb{R}^d \\ (\mathbf{A}, \mathbf{t}) \in \Xi_k}} |g_{\mathbf{A}, \mathbf{t}}(\mathbf{x})| (1 + \|\mathbf{x}\|_2)^{-n_L} < \infty$$

with

$$(3.10) \quad (\mathbf{A}, \mathbf{t}) \mapsto g_{\mathbf{A}, \mathbf{t}}(\mathbf{x}_0) \in C_{0, \text{iso}}(\Xi_k)$$

for any fixed $\mathbf{x}_0 \in \mathbb{R}^d$. Thus, for $\mu \in \mathcal{M}_{\text{iso}}(\Xi_k)$,

$$(3.11) \quad L_{\mathcal{R}_k}^\dagger \{\mu\}(\mathbf{x}) = \int_{\Xi_k} g_{\mathbf{A}, \mathbf{t}}(\mathbf{x}) d\mu(\mathbf{A}, \mathbf{t}), \quad \mathbf{x} \in \mathbb{R}^d.$$

Theorem 3.7. *Consider the setting of Theorem 3.6. Then, the following hold.*

1. The range space $\mathcal{V} := L_{\mathcal{R}_k}^\dagger(\mathcal{M}_{\text{iso}}(\Xi_k))$ is a Banach space when equipped with the norm

$$(3.12) \quad \|f\|_{\mathcal{V}} := \|L_{\mathcal{R}_k} f\|_{\mathcal{M}}.$$

This Banach space is isometrically isomorphic to $\mathcal{M}_{\text{iso}}(\Xi_k)$.

2. The native space $\mathcal{M}_L^k(\mathbb{R}^d)$ is decomposable as the direct sum of Banach spaces

$$(3.13) \quad \mathcal{M}_L^k(\mathbb{R}^d) = \mathcal{V} \oplus \mathcal{P}_{n_L}(\mathbb{R}^d) = L_{\mathcal{R}_k}^\dagger(\mathcal{M}_{\text{iso}}(\Xi_k)) \oplus \mathcal{P}_{n_L}(\mathbb{R}^d).$$

3. The native space $\mathcal{M}_L^k(\mathbb{R}^d)$ forms a bona fide Banach space when equipped with the norm

$$(3.14) \quad \|f\|_{\mathcal{M}_L^k} := \|\mathbf{L}_{\mathcal{R}_k} f\|_{\mathcal{M}} + \|\mathbf{P}_{\mathcal{P}_{n_L}(\mathbb{R}^d)} f\|_{\mathcal{P}_{n_L}}.$$

Furthermore, $\mathcal{M}_L^k(\mathbb{R}^d)$ is isometrically isomorphic to $\mathcal{M}_{\text{iso}}(\Xi_k) \times \mathcal{P}_{n_L}(\mathbb{R}^d)$ via the map

$$(3.15) \quad f = \mathbf{L}_{\mathcal{R}_k}^\dagger \{u\} + p \mapsto (u, p),$$

where $u = \mathbf{L}_{\mathcal{R}_k} \{f\}$ and $p = \mathbf{P}_{\mathcal{P}_{n_L}(\mathbb{R}^d)} f$.

4. For any $\mathbf{x}_0 \in \mathbb{R}^d$, the shifted Dirac impulse (point evaluation functional) $\delta(\cdot - \mathbf{x}_0) : f \mapsto f(\mathbf{x}_0)$ is weak*-continuous⁵ on $\mathcal{M}_L^k(\mathbb{R}^d)$.

The proofs of Theorems 3.6 and 3.7 appear in Appendices A and B, respectively.

3.2. Optimality of neural architectures with multivariate nonlinearities. Having established the properties of the native space $\mathcal{M}_L^k(\mathbb{R}^d)$ in subsection 3.1, we can now prove our main theorem regarding the function-space optimality of neural architectures with multivariate nonlinearities.

Theorem 3.8. Let $\mathcal{L}(\cdot, \cdot) : \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$ be convex, coercive, and lower-semicontinuous in its second argument and let \mathbf{L} be a k -plane-admissible operator in the sense of Definition 3.1. Then, for any finite data set $\{(\mathbf{x}_m, y_m)\}_{m=1}^M \subset \mathbb{R}^d \times \mathbb{R}$ for which the data-fitting problem is well-posed⁶ over $\mathcal{P}_{n_L}(\mathbb{R}^d)$, the solution set to the data-fitting variational problem

$$(3.16) \quad S := \arg \min_{f \in \mathcal{M}_L^k(\mathbb{R}^d)} \sum_{m=1}^M \mathcal{L}(y_m, f(\mathbf{x}_m)) + \lambda \|\mathbf{K}_{d-k} \mathcal{R}_k \mathbf{L} f\|_{\mathcal{M}}$$

is nonempty, convex, and weak*-compact. If $\mathcal{L}(\cdot, \cdot)$ is strictly convex (or if it imposes the equality $y_m = f(\mathbf{x}_m)$ for $m = 1, \dots, M$), then S is the weak*-closure of the convex hull of its extreme points, which can all be expressed as

$$(3.17) \quad f_{\text{extreme}}(\mathbf{x}) = c(\mathbf{x}) + \sum_{n=1}^N v_n \rho_L(\mathbf{A}_n \mathbf{x} - \mathbf{t}_n),$$

where the number N of atoms satisfies $N \leq (M - \dim \mathcal{P}_{n_L}(\mathbb{R}^d))$, and, for $n = 1, \dots, N$, $v_n \in \mathbb{R} \setminus \{0\}$, $\mathbf{A}_n \in V_{d-k}(\mathbb{R}^d)$, and $\mathbf{t}_n \in \mathbb{R}^{d-k}$. The function $c \in \mathcal{P}_{n_L}(\mathbb{R}^d)$ is a polynomial of degree at most n_L and $\rho_L : \mathbb{R}^{d-k} \rightarrow \mathbb{R}^d$ is a $(d-k)$ -variate nonlinearity given by $\rho_L = \mathcal{F}_{d-k}^{-1} \{1/\widehat{L}_{\text{rad}}(\|\cdot\|_2)\}$, where \mathcal{F}_{d-k}^{-1} is the $(d-k)$ -variate inverse Fourier transform. Finally, the regularization cost, which is common to all solutions, is

$$(3.18) \quad \|\mathbf{K}_{d-k} \mathcal{R}_k \mathbf{L} f_{\text{extreme}}\|_{\mathcal{M}} = \sum_{n=1}^N |v_n| = \|\mathbf{v}\|_1.$$

The proof of the theorem requires the following proposition.

⁵In particular, we prove that $(\mathcal{M}_L^k(\mathbb{R}^d), \|\cdot\|_{\mathcal{M}_L^k})$ can be identified as the dual of some primary Banach space, which allows us to equip $\mathcal{M}_L^k(\mathbb{R}^d)$ with a weak* topology.

⁶The data-fitting problem is well-posed over the null space when the classical least-squares polynomial-fitting problem admits a unique solution with respect to the data $\{(\mathbf{x}_m, y_m)\}_{m=1}^M \subset \mathbb{R}^d \times \mathbb{R}$.

Proposition 3.9 ([47, Lemma 10.2]). *The isotropic shifted Dirac impulse $\delta_{\text{iso}}(\cdot - (\mathbf{A}_0, \mathbf{t}_0)) \in \mathcal{M}_{\text{iso}}(\Xi_k)$ defined by*

$$(3.19) \quad \delta_{\text{iso}}(\cdot - (\mathbf{A}_0, \mathbf{t}_0)) := \text{P}_{\text{iso}}\{\delta(\cdot - (\mathbf{A}_0, \mathbf{t}_0))\},$$

where $\delta(\cdot - (\mathbf{A}_0, \mathbf{t}_0)) = \delta(\cdot - \mathbf{A}_0)\delta(\cdot - \mathbf{t}_0) \in \mathcal{M}(\Xi_k)$ is the “classical” Dirac impulse on Ξ_k , satisfies the following properties:

1. *Sampling:* For any $\psi \in C_{0,\text{iso}}(\Xi_k)$,

$$(3.20) \quad \langle \delta_{\text{iso}}(\cdot - (\mathbf{A}_0, \mathbf{t}_0)), \psi \rangle_k = \psi(\mathbf{A}_0, \mathbf{t}_0).$$

2. *Rotation invariance:* For any $\mathbf{U} \in \text{O}_{d-k}(\mathbb{R})$,

$$(3.21) \quad \delta_{\text{iso}}(\cdot - (\mathbf{A}_0, \mathbf{t}_0)) = \delta_{\text{iso}}(\cdot - (\mathbf{U}\mathbf{A}_0, \mathbf{U}\mathbf{t}_0)).$$

3. *Unit norm:* $\|\delta_{\text{iso}}(\cdot - (\mathbf{A}_0, \mathbf{t}_0))\|_{\mathcal{M}} = 1$.

4. *Linear combination:* For any set $\{(\mathbf{A}_n, \mathbf{t}_n)\}_{n=1}^N \subset \Xi_k$ of distinct points,

$$(3.22) \quad \left\| \sum_{n=1}^N a_n \delta_{\text{iso}}(\cdot - (\mathbf{A}_n, \mathbf{t}_n)) \right\|_{\mathcal{M}} = \sum_{n=1}^N |a_n| = \|\mathbf{a}\|_1.$$

5. *Extreme points of $B_{\mathcal{M}_{\text{iso}}} := \{e \in \mathcal{M}_{\text{iso}}(\Xi_k) : \|e\|_{\mathcal{M}} \leq 1\}$:* If $e \in \text{Ext } B_{\mathcal{M}_{\text{iso}}}$, then $e = \pm \delta_{\text{iso}}(\cdot - (\mathbf{A}_n, \mathbf{t}_n))$ for some $(\mathbf{A}_n, \mathbf{t}_n) \in \Xi_k$.

Proof of Theorem 3.8. The proof relies on the abstract representer theorem in [68] (see also [6, 7, 65]). From the assumptions on the loss function combined with the weak*-continuity of the point evaluation functional on $\mathcal{M}_{\text{L}}^k(\mathbb{R}^d)$ (item 4 in Theorem 3.7), our setting coincides with the hypotheses of [68, Theorem 3]. First, this abstract result ensures that the solution set S is nonempty, convex, and weak*-compact. Second, it ensures that, when the loss function is strictly convex (or if it imposes the equality $y_m = f(\mathbf{x}_m)$ for $m = 1, \dots, M$), S is the weak*-closure of the convex hull of its extreme points, which can all be expressed as

$$(3.23) \quad f_{\text{extreme}}(\mathbf{x}) = c(\mathbf{x}) + \sum_{n=1}^N v_n e_n(\mathbf{x}),$$

where the number N of atoms satisfies $N \leq (M - \dim \mathcal{P}_{n_{\text{L}}}(\mathbb{R}^d))$, $c(\cdot)$ is in the null space of the regularizer (i.e., $c \in \mathcal{P}_{n_{\text{L}}}(\mathbb{R}^d)$), and, for $n = 1, \dots, N$, $v_n \in \mathbb{R} \setminus \{0\}$ and e_n is an extreme point of the unit regularization ball

$$(3.24) \quad B := \left\{ f \in \mathcal{M}_{\text{L}}^k(\mathbb{R}^d) : \|\mathbf{K}_{d-k} \mathcal{R}_k \mathbf{L} f\|_{\mathcal{M}} \leq 1 \right\}.$$

From item 2 in Theorem 3.7, we have the direct-sum decomposition of the native space as

$$(3.25) \quad \mathcal{M}_{\text{L}}^k(\mathbb{R}^d) = \mathcal{V} \oplus \mathcal{P}_{n_{\text{L}}}(\mathbb{R}^d) = \text{L}_{\mathcal{R}_k}^{\dagger}(\mathcal{M}_{\text{iso}}(\Xi_k)) \oplus \mathcal{P}_{n_{\text{L}}}(\mathbb{R}^d).$$

Since $\text{L}_{\mathcal{R}_k}^{\dagger} : \mathcal{M}_{\text{iso}}(\Xi_k) \rightarrow \mathcal{V}$ is an isometric isomorphism (item 1 in Theorem 3.7), the extreme points of B take the form $\text{L}_{\mathcal{R}_k}^{\dagger}(\text{Ext } B_{\mathcal{M}_{\text{iso}}})$ plus a polynomial term in $\mathcal{P}_{n_{\text{L}}}(\mathbb{R}^d)$. From item 5 in Proposition 3.9, it then follows that

$$(3.26) \quad e_n = \text{L}_{\mathcal{R}_k}^{\dagger} \{ \pm \delta_{\text{iso}}(\cdot - (\mathbf{A}_n, \mathbf{t}_n)) \} + p = \pm \rho_{\text{L}}(\mathbf{A}_n(\cdot) - \mathbf{t}_n) + \tilde{p}$$

with $(\mathbf{A}_n, \mathbf{t}_n) \in \Xi_k$ and $p, \tilde{p} \in \mathcal{P}_{n_{\text{L}}}(\mathbb{R}^d)$.

A calculation in the Fourier domain reveals that $L\{\rho_L(\mathbf{A}(\cdot) - \mathbf{t})\} = \delta(\mathbf{A}(\cdot) - \mathbf{t})$ for $(\mathbf{A}, \mathbf{t}) \in \Xi_k$. Next, we invoke the property that $K_{d-k} \mathcal{R}_k \{ \delta(\mathbf{A}(\cdot) - \mathbf{t}) \} = \delta_{\text{iso}}(\cdot - (\mathbf{A}, \mathbf{t}))$ [47, equation (9.16)], which yields that

$$(3.27) \quad K_{d-k} \mathcal{R}_k L\{f_{\text{extreme}}\} = \sum_{n=1}^N v_n \delta_{\text{iso}}(\cdot - (\mathbf{A}_n, \mathbf{t}_n)).$$

Finally, by item 4 in Proposition 3.9, we have that $\|K_{d-k} \mathcal{R}_k L f_{\text{extreme}}\|_{\mathcal{M}} = \sum_{n=1}^N |v_n| = \|\mathbf{v}\|_1$, which proves the theorem. \blacksquare

3.2.1. Discussion. The main takeaway from Theorem 3.8 is that *sparse* neural architectures (sparse in the sense that there are fewer neurons than data) are solutions to variational problems over $\mathcal{M}_L^k(\mathbb{R}^d)$. In particular, the regularity of functions imposed by the Banach structure of $\mathcal{M}_L^k(\mathbb{R}^d)$ explains the variational optimality of the architectures in (3.17). Furthermore, by the isomorphism in item 3 of Theorem 3.7, we see that $\mathcal{M}_L^k(\mathbb{R}^d)$ is a *nonreflexive Banach space* (since $\mathcal{M}_{\text{iso}}(\Xi_k)$ is nonreflexive), which shows that $\mathcal{M}_L^k(\mathbb{R}^d)$ differs in a fundamental way from a Hilbert space. Said differently, Theorem 3.8 provides a function-space framework for neural networks that differs in a fundamental way from the (Hilbertian) framework of the neural tangent kernel [26].

The two extremes of the theorem ($k = 0$ and $k = (d - 1)$) capture well-studied problems. Indeed, when $k = 0$, we can take advantage of the fact that the effect of the k -plane transform essentially disappears. Indeed, when $k = 0$ we have (see subsection 2.1.1) that

$$(3.28) \quad \|K_d \mathcal{R}_0 L f\|_{\mathcal{M}(\Xi_0)} = c_{d,0} \|(\mathbf{A}, \mathbf{t}) \mapsto L\{f\}(\mathbf{A}^\top \mathbf{t})\|_{\mathcal{M}(\Xi_0)} = \|L f\|_{\mathcal{M}(\mathbb{R}^d)}.$$

Therefore, the variational problem in (3.16) reduces to the well-studied variational problem for L-splines [13, 69]. Since L is isotropic from the admissibility assumptions (Definition 3.1), the atoms ρ_L are radial basis functions.

For these problems, the classical theory [13, 69] suggests that the extreme point solutions are built from atoms of the form $\rho_L(\cdot - \boldsymbol{\tau}_n)$, $\boldsymbol{\tau}_n \in \mathbb{R}^d$, where $\rho_L : \mathbb{R}^d \rightarrow \mathbb{R}$ is the (canonical) Green's function of L defined in the Fourier domain by $\widehat{\rho}_L = 1/\widehat{L}$. We can quickly see that Theorem 3.8 recovers this result since the atoms take the form

$$(3.29) \quad \mathbf{x} \mapsto \rho_L(\mathbf{A}_n \mathbf{x} - \mathbf{t}_n) = \rho_L(\mathbf{A}_n^\top \mathbf{A}_n \mathbf{x} - \mathbf{A}_n^\top \mathbf{t}_n) = \rho_L(\mathbf{x} - \mathbf{A}_n^\top \mathbf{t}_n) = \rho_L(\mathbf{x} - \boldsymbol{\tau}_n),$$

where we made the substitution $\boldsymbol{\tau}_n = \mathbf{A}_n^\top \mathbf{t}_n \in \mathbb{R}^d$ in the last equality. Here, we used the fact that $\mathbf{A} \mathbf{A}^\top = \mathbf{A}^\top \mathbf{A} = \mathbf{I}_d$ when $\mathbf{A} \in V_d(\mathbb{R}^d)$ (i.e., the $k = 0$ Stiefel manifold is the space of $(d \times d)$ orthogonal matrices).

At the opposite extreme, when $k = (d - 1)$, the atoms take the form

$$(3.30) \quad \mathbf{x} \mapsto \rho_L(\boldsymbol{\alpha}^\top \mathbf{x} - t),$$

where $\rho_L : \mathbb{R} \rightarrow \mathbb{R}$ is the Green's function of the univariate operator L_{rad} , specified by the frequency response \widehat{L}_{rad} of the univariate radial profile of L. These atoms are classical neurons with univariate nonlinearities. This problem was first studied in [43] with $L_{\text{rad}} = \partial_t^m$, which corresponds to nonlinearities proportional to the truncated power functions $t \mapsto t_+^{m-1}$ (which is the ReLU when $m = 2$), and then generalized to other regularization operators L in [66].

4. Observations and examples of compatible neural architectures. Since Theorem 3.8 guarantees the existence of a solution to (3.16) that takes the form in (3.17), we can always find an admissible solution by solving the neural network training problem

$$(4.1) \quad \begin{aligned} \min_{\boldsymbol{\theta}} \quad & \left(\sum_{m=1}^M \mathcal{L}(y_m, f_{\boldsymbol{\theta}}(\mathbf{x}_m)) + \lambda \sum_{n=1}^N |v_n| \right) \\ \text{s.t.} \quad & \mathbf{A}_n \mathbf{A}_n^T = \mathbf{I}_{d-k}, \quad n = 1, \dots, N, \end{aligned}$$

for some fixed width $N \geq M$ with

$$(4.2) \quad f_{\boldsymbol{\theta}}(\mathbf{x}) = c(\mathbf{x}) + \sum_{n=1}^N v_n \rho_L(\mathbf{A}_n \mathbf{x} - \mathbf{t}_n), \quad \mathbf{x} \in \mathbb{R}^d, \boldsymbol{\theta} = \{v_n, \mathbf{A}_n, \mathbf{t}_n\}_{n=1}^N \cup \{c(\cdot)\}.$$

The assumption that $N \geq M$ ensures that a solution to the variational problem in (3.16) exists in the neural network parameter space (indexed by $\boldsymbol{\theta}$) thanks to the bound in Theorem 3.8. This assumption implies that, as long as the neural network problem is critically parameterized or overparameterized, its solutions will also be solutions to the variational problem in Theorem 3.8. Thus, this result provides insight on the role of overparameterization. We also remark that the constraint on the weight matrices in (4.1) corresponds to orthogonal weight normalization. The latter has become a popular architectural choice as it has been shown to increase the stability and improve the generalization performance of neural networks [2, 24, 25, 33].

The nonlinearity $\rho_L : \mathbb{R}^{d-k} \rightarrow \mathbb{R}^{d-k}$ that appears in (3.17) can be viewed as the Green’s function of the operator $L_{d-k} : \mathcal{S}(\mathbb{R}^{d-k}) \rightarrow \mathcal{S}'(\mathbb{R}^{d-k})$ which shares the radial profile \widehat{L}_{rad} of the k -plane-admissible operator L . That is to say,

$$(4.3) \quad \widehat{L}_{d-k}(\boldsymbol{\omega}) = \widehat{L}_{\text{rad}}(\|\boldsymbol{\omega}\|_2).$$

Furthermore, due to the intertwining properties of the k -plane transform [47, Corollary 7.8], it turns out that the regularization operator in (3.16) has the alternative specification

$$(4.4) \quad K_{d-k} \mathcal{R}_k L = L_{d-k} K_{d-k} \mathcal{R}_k.$$

The framework of Theorem 3.8 encapsulates many neural architectures. The prototypical example of such an operator is the fractional Laplacian $L = (-\Delta)^{\frac{\alpha}{2}}$ and, so, $L_{d-k} = (-\Delta_{d-k})^{\frac{\alpha}{2}}$. The radial profile for this family of operators is

$$(4.5) \quad \widehat{L}_{\text{rad}}(\boldsymbol{\omega}) = |\boldsymbol{\omega}|^{\alpha}.$$

From Definition 3.1, the reader can readily verify that this operator is k -plane-admissible for $\alpha > (d - k)$. This simple operator encapsulates several known results. At one extreme ($k = (d - 1)$), we recover⁷ the classical ReLU neurons by the choice $L = (-\Delta)$ [41, 43]. At the opposite extreme ($k = 0$), from subsection 3.2.1, we see that we recover a sparse variant of the

⁷Technically, $\rho_{(-\Delta)}(t) = |t|/2$ in this case, but since $t \mapsto |t|/2$ differs from the ReLU $t \mapsto t_+$ by a null space component (affine function), the ReLU and absolute value nonlinearity are treated the same in this framework.

classical thin-plate/polyharmonic spline radial basis functions of Duchon [12] by the choice $L = (-\Delta)^{\frac{\alpha}{2}}$, $\alpha > d$.

In the intermediate regime $1 \leq k \leq (d-2)$, we can choose $L_{d-k} = (-\Delta_{d-k})^{\frac{1+(d-k)}{2}}$ so that $\rho_L(\mathbf{t}) \propto \|\mathbf{t}\|_2$, $\mathbf{t} \in \mathbb{R}^{d-k}$, is the norm activation function that has been used for neural architectures in [23]. These observations follow from the fact that the Green's function of the fractional Laplacian $(-\Delta_n)^{\frac{\alpha}{2}}$ (which acts on n -variables) for $\alpha > n$ takes the form

$$(4.6) \quad k_{\alpha,n}(\mathbf{t}) = \mathcal{F}^{-1} \left\{ \frac{1}{\|\cdot\|_2} \right\}(\mathbf{t}) = \begin{cases} A_{\alpha,n} \|\mathbf{t}\|_2^{\alpha-n}, & \alpha - n \notin 2\mathbb{N}, \\ B_{m,n} \|\mathbf{t}\|_2^{2m} \log \|\mathbf{t}\|_2, & \alpha - n = 2m, m \in \mathbb{N}, \\ \Delta_n^{-m} \{\delta\}, & -\alpha/2 = m, m \in \mathbb{N}, \end{cases}$$

with $\mathbf{t} \in \mathbb{R}^n$, $A_{\alpha,n} = \frac{\Gamma((n-\alpha)/2)}{2^\alpha \pi^{n/2} \Gamma(\alpha/2)}$, and $B_{m,n} = \frac{(-1)^{1+m}}{2^{2m+n-1} \pi^{n/2} \Gamma(m+n/2) m!}$ [17, 53]. In general, there exist many nonlinearities that are compatible with the presented framework. All that needs to be verified is the admissibility conditions (Definition 3.1) of the underlying regularization operator.

5. Connections to RKBS methods and variation spaces. After [43], a recent line of research has been trying to understand neural networks through the lens of RKBSs [5, 62]. These works consider Banach spaces defined on, say, \mathbb{R}^d whose members are defined as integral combinations of atoms from some continuously indexed dictionary \mathcal{D} . The elements of the dictionary are assumed to be continuously indexed by $\xi \in \Xi$, where Ξ is assumed to be some locally compact Hausdorff space. That is, $\mathcal{D} = \{\varphi_\xi\}_{\xi \in \Xi}$, with the additional hypothesis that $\xi \mapsto \varphi_\xi(\mathbf{x}) \in C_0(\Xi)$ for any $\mathbf{x} \in \mathbb{R}^d$.

It turns out that the space

$$(5.1) \quad \mathcal{B}(\mathbb{R}^d) := \left\{ f : \mathbb{R}^d \rightarrow \mathbb{R} \text{ is measurable} : \text{there exists } \mu \in \mathcal{M}(\Xi) \text{ s.t. } f = \int_{\Xi} \varphi_\xi \, d\mu(\xi) \right\}$$

forms a Banach space when equipped with the norm

$$(5.2) \quad \|f\|_{\mathcal{B}} := \inf_{\mu \in \mathcal{M}(\Xi)} \|\mu\|_{\mathcal{M}} \quad \text{s.t.} \quad f = \int_{\Xi} \varphi_\xi \, d\mu(\xi).$$

The assumptions on $\mathcal{D} = \{\varphi_\xi\}_{\xi \in \Xi}$ ensure that the point evaluation is continuous on $\mathcal{B}(\mathbb{R}^d)$ (i.e., $\delta(\cdot - \mathbf{x}_0) \in \mathcal{B}'(\mathbb{R}^d)$). Such Banach spaces are referred to as *RKBSs* [34, 72]. An RKBS formed from integral combinations of atoms from some continuously indexed dictionary is an *integral RKBS* (I-RKBS) [62]. With this formalism, [5, 62] study many properties of $\mathcal{B}(\mathbb{R}^d)$ as well as data-fitting problems over these spaces with associated representer theorems. We remark that, thanks to the assumption $\xi \mapsto \varphi_\xi(\mathbf{x}) \in C_0(\Xi)$ for any $\mathbf{x} \in \mathbb{R}^d$, these authors implicitly ensure that the point evaluation is actually weak*-continuous on $\mathcal{B}(\mathbb{R}^d)$, which is stronger than standard continuity. This property is critical in proving the existence of solutions to data-fitting problems over these spaces.

Our native spaces $\mathcal{M}_L^k(\mathbb{R}^d)$ are compatible with the I-RKBS framework. We first note that item 4 in Theorem 3.7 ensures that the point evaluation is weak*-continuous and, hence, continuous on $\mathcal{M}_L^k(\mathbb{R}^d)$. Thus, $\mathcal{M}_L^k(\mathbb{R}^d)$ is an RKBS. Next, we have the direct-sum decomposition

$$(5.3) \quad \mathcal{M}_L^k(\mathbb{R}^d) = L_{\mathcal{R}_k}^\dagger(\mathcal{M}_{\text{iso}}(\Xi_k)) \oplus \mathcal{P}_{n_L}(\mathbb{R}^d) = L_{\mathcal{R}_k}^\dagger(\mathcal{M}(\Xi_k)) \oplus \mathcal{P}_{n_L}(\mathbb{R}^d),$$

where the first equality is from item 2 in Theorem 3.7 and the second inequality follows since the null space of \mathcal{R}_k^* is $(\mathcal{M}_{\text{iso}}(\Xi_k))^c$ (recall that $L_{\mathcal{R}_k}^\dagger = (\text{Id} - P_{\mathcal{P}_{n_L}(\mathbb{R}^d)})L^{-1}\mathcal{R}_k^*$ and see Proposition 2.7). This immediately implies that $\mathcal{M}_L^k(\mathbb{R}^d)$ is the direct sum of an I-RKBS with $\mathcal{P}_{n_L}(\mathbb{R}^d)$, where the dictionary is composed of the kernels $g_{\mathbf{A},\mathbf{t}}$ from (3.8), continuously indexed by $(\mathbf{A}, \mathbf{t}) \in \Xi_k$. This correspondence allows us to directly apply any I-RKBS developments to $\mathcal{M}_L^k(\mathbb{R}^d)$.

The study of variation spaces to understand neural networks is a classical endeavor [31, 37]. These spaces have received renewed interest [4, 11, 56, 58, 57] as a means toward the understanding of the reason why neural networks seem to “break” the curse of dimensionality through the lens of nonlinear approximation theory. It turns out that the variation space for a dictionary \mathcal{D} exactly coincides with the I-RKBS as long as the members of \mathcal{D} are sufficiently regular (see [57, Lemma 3]). Indeed, in that case, the variation space for \mathcal{D} is the Banach space $(\mathcal{B}(\mathbb{R}^d), \|\cdot\|_{\mathcal{B}})$ defined in (5.1). Thus, $\mathcal{M}_L^k(\mathbb{R}^d)$ can also be viewed as a variation space. The investigation of the implications of these tight connections to I-RKBS and variation spaces toward the understanding of neural architectures with multivariate nonlinearities is a direction for future work.

Appendix A. Proof of Theorem 3.6.

Proof. We first note that $L_{\mathcal{R}_k}$ maps $\mathcal{M}_L^k(\mathbb{R}^d) \rightarrow \mathcal{M}_{\text{iso}}(\Xi_k)$ by design due to the isotropic symmetry of the k -plane domain. The remainder of the proof is divided into three parts.

Part (I): Existence/continuity of $L_{\mathcal{R}_k}^\dagger$ and the right-inverse property (3.5). We first note that $L^* : \mathcal{S}(\mathbb{R}^d) \rightarrow L_{n_L}^1(\mathbb{R}^d)$ is a continuous injection (item 1 in Definition 3.1). Therefore, there exists an inverse operator $L^{*-1} : L^*(\mathcal{S}(\mathbb{R}^d)) \rightarrow \mathcal{S}(\mathbb{R}^d)$ such that $L^{*-1}L^* = \text{Id}$ on $\mathcal{S}(\mathbb{R}^d)$. In particular, L^{*-1} is the LSI operator specified by the frequency response $\boldsymbol{\xi} \mapsto 1/\widehat{L}(\boldsymbol{\xi})$ (since L^* is necessarily self-adjoint from item 2 in Definition 3.1).

Next, we define the operator

$$(A.1) \quad L_{\mathcal{R}_k}^{\dagger*} := \mathcal{R}_k L^{-1*}(\text{Id} - P_{\mathcal{P}_{n_L}(\mathbb{R}^d)}^*),$$

where

$$(A.2) \quad P_{\mathcal{P}_{n_L}(\mathbb{R}^d)}^*\{f\} = \sum_{|\mathbf{n}| \leq n_L} \langle m_{\mathbf{n}}, f \rangle m_{\mathbf{n}}^*,$$

which is the projection of f onto the dual space $(\mathcal{P}_{n_L}(\mathbb{R}^d))' \subset \mathcal{S}(\mathbb{R}^d)$. By recalling from (2.24) that $m_{\mathbf{n}}(\mathbf{x}) = \mathbf{x}^{\mathbf{n}}/\mathbf{n}!$, we see that (A.2) is a well-defined operator as long as f has sufficient decay (e.g., $f \in L^*(\mathcal{S}(\mathbb{R}^d)) \subset L_{n_L}^1(\mathbb{R}^d)$). This reveals that

$$(A.3) \quad L_{\mathcal{R}_k}^{\dagger*} : L^*(\mathcal{S}(\mathbb{R}^d)) \rightarrow \mathcal{S}_k.$$

To check the continuity of this operator, we characterize the boundedness of the (Schwartz) kernel of $L_{\mathcal{R}_k}^{\dagger*}$. This kernel can be formally identified with $((\mathbf{A}, \mathbf{t}), \mathbf{x}) \mapsto h_{\mathbf{x}}(\mathbf{A}, \mathbf{t}) := L_{\mathcal{R}_k}^{\dagger*} \{\delta(\cdot - \mathbf{x})\}(\mathbf{A}, \mathbf{t})$. By the Fourier slice theorem,

$$\begin{aligned}
\widehat{h_{\mathbf{x}}(\mathbf{A}, \cdot)}(\boldsymbol{\omega}) &= \frac{\mathcal{F}\left\{\delta(\cdot - \mathbf{x}) - \sum_{|\mathbf{n}| \leq n_L} \langle m_{\mathbf{n}}, \delta(\cdot - \mathbf{x}) \rangle m_{\mathbf{n}}^* \right\}(\mathbf{A}^\top \boldsymbol{\omega})}{\widehat{L}_{\text{rad}}(\|\boldsymbol{\omega}\|_2)} \\
&= \frac{e^{-i\boldsymbol{\omega}^\top \mathbf{A}\mathbf{x}} - \sum_{|\mathbf{n}| \leq n_L} \frac{\mathbf{x}^{\mathbf{n}}}{n!} \widehat{m}_{\mathbf{n}}^*(\mathbf{A}^\top \boldsymbol{\omega})}{\widehat{L}_{\text{rad}}(\|\boldsymbol{\omega}\|_2)} \\
\text{(A.4)} \quad &= \frac{e^{-i\boldsymbol{\omega}^\top \mathbf{A}\mathbf{x}} - \sum_{|\mathbf{n}| \leq n_L} \frac{\mathbf{x}^{\mathbf{n}}}{n!} (-i\mathbf{A}^\top \boldsymbol{\omega})^{\mathbf{n}} \widehat{\kappa}_{\text{rad}}(\|\boldsymbol{\omega}\|_2)}{\widehat{L}_{\text{rad}}(\|\boldsymbol{\omega}\|_2)} \\
&= \frac{e^{-i\boldsymbol{\omega}^\top \mathbf{A}\mathbf{x}} - \sum_{n=0}^{n_L} \frac{(-i\boldsymbol{\omega}^\top \mathbf{A}\mathbf{x})^n}{n!} \widehat{\kappa}_{\text{rad}}(\|\boldsymbol{\omega}\|_2)}{\widehat{L}_{\text{rad}}(\|\boldsymbol{\omega}\|_2)}, \\
&= \frac{e^{-i\boldsymbol{\omega}^\top \mathbf{A}\mathbf{x}} - \widehat{\kappa}_{\text{rad}}(\|\boldsymbol{\omega}\|_2) \sum_{n=0}^{n_L} \frac{(-i\boldsymbol{\omega}^\top \mathbf{A}\mathbf{x})^n}{n!}}{\widehat{L}_{\text{rad}}(\|\boldsymbol{\omega}\|_2)},
\end{aligned}$$

where the penultimate line holds by the multinomial expansion. Note that the quantity in (A.4) is well-defined despite the pole of multiplicity γ_L at $\boldsymbol{\omega} = \mathbf{0}$. Since $\gamma_L \in (n_L, n_L + 1]$, the form of the numerator ensures a proper pole-zero cancellation with the denominator. Indeed, by Taylor's theorem, when $t \in \mathbb{R}$ is in a neighborhood of 0, we have that

$$\text{(A.5)} \quad e^t - \sum_{n=1}^{n_L} \frac{t^n}{n!} = O(t^{n_L+1}).$$

By the identification of the numerator⁸ in (A.4) with the above display combined with the property of item 3 in Definition 3.1, we have that (A.4) is well-defined. Next, since $\widehat{\kappa}_{\text{rad}}(\|\boldsymbol{\omega}\|_2) \leq 1$ for $\|\boldsymbol{\omega}\|_2 < 1$ and $\widehat{\kappa}_{\text{rad}}(\|\boldsymbol{\omega}\|_2) = 0$ for $\|\boldsymbol{\omega}\|_2 \geq 1$ (subsection 2.2), on one hand we have for $\|\boldsymbol{\omega}\|_2 < 1$ that $|\widehat{h_{\mathbf{x}}(\mathbf{A}, \cdot)}(\boldsymbol{\omega})|$ is bounded by a constant which depends on \mathbf{A} and \mathbf{x} . To see the dependence on \mathbf{A} and \mathbf{x} , we note that

$$\begin{aligned}
&\sum_{n=0}^{n_L} \frac{1}{n!} |i\boldsymbol{\omega}^\top \mathbf{A}\mathbf{x}|^n \\
&\leq \sum_{n=0}^{n_L} \frac{1}{n!} \|\boldsymbol{\omega}\|_2^n \|\mathbf{A}\mathbf{x}\|_2^n \\
&\leq \sum_{n=0}^{n_L} \frac{(n_L - n)!}{n_L!} \frac{n_L!}{n!(n_L - n)!} \|\mathbf{A}\mathbf{x}\|_2^n \\
&\leq \sum_{n=0}^{n_L} \frac{n_L!}{n!(n_L - n)!} \|\mathbf{A}\mathbf{x}\|_2^n \\
\text{(A.6)} \quad &= (1 + \|\mathbf{A}\mathbf{x}\|_2)^{n_L}.
\end{aligned}$$

Thus, there exists a universal constant $C_0 > 0$ such that

$$\text{(A.7)} \quad |\widehat{h_{\mathbf{x}}(\mathbf{A}, \cdot)}(\boldsymbol{\omega})| \leq C_0(1 + \|\mathbf{A}\mathbf{x}\|_2)^{n_L}, \quad \|\boldsymbol{\omega}\|_2 < 1.$$

⁸This identification is valid by the substitution $t = -i\boldsymbol{\omega}^\top \mathbf{A}\mathbf{x}$ and noting that $\widehat{\kappa}_{\text{rad}}(\|\boldsymbol{\omega}\|_2) = 1$ for $\|\boldsymbol{\omega}\|_2 < R_0$, for some $R_0 \leq 1/2$ [67, p. 6].

On the other hand, when $\|\boldsymbol{\omega}\|_2 \geq 1$, we have that

$$(A.8) \quad |\widehat{h_{\mathbf{x}}(\mathbf{A}, \cdot)}(\boldsymbol{\omega})| \leq \frac{1}{\widehat{L}_{\text{rad}}(\|\boldsymbol{\omega}\|_2)} \leq \frac{C_0(1 + \|\mathbf{A}\mathbf{x}\|_2)^{n_L}}{\widehat{L}_{\text{rad}}(\|\boldsymbol{\omega}\|_2)}.$$

From item 4 in Definition 3.1, we have that

$$(A.9) \quad \frac{1}{\widehat{L}_{\text{rad}}(\|\boldsymbol{\omega}\|_2)} \leq \frac{\|\boldsymbol{\omega}\|_2^{-\gamma'_L}}{C'},$$

where $\gamma'_L > (d-k)$. This property ensures that $\widehat{h_{\mathbf{x}}(\mathbf{A}, \cdot)} \in L^1(\mathbb{R}^{d-k})$. Combining (A.7)–(A.9), implies that

$$(A.10) \quad (1 + \|\mathbf{A}\mathbf{x}\|_2)^{-n_L} \|\widehat{h_{\mathbf{x}}(\mathbf{A}, \cdot)}\|_{L^1} < \infty.$$

Since this bound is uniform in $\mathbf{x} \in \mathbb{R}^d$ and $\mathbf{A} \in V_{d-k}(\mathbb{R}^d)$ and the inverse Fourier transform $\mathcal{F}_{d-k}^{-1} : L^1(\mathbb{R}^{d-k}) \rightarrow C_0(\mathbb{R}^{d-k})$ is a bounded operator (Riemann–Lebesgue lemma), we have that

$$(A.11) \quad \sup_{\substack{\mathbf{x} \in \mathbb{R}^d \\ (\mathbf{A}, \mathbf{t}) \in \Xi_k}} |h_{\mathbf{x}}(\mathbf{A}, \mathbf{t})|(1 + \|\mathbf{A}\mathbf{x}\|_2)^{-n_L} < \infty$$

with the property that $h_{\mathbf{x}} \in C_{0,\text{iso}}(\Xi_k)$, due to the isotropic symmetry of the k -plane transform. Finally, if we write

$$(A.12) \quad \mathbf{A}\mathbf{x} = \begin{bmatrix} \boldsymbol{\alpha}_1^\top \mathbf{x} \\ \vdots \\ \boldsymbol{\alpha}_{d-k}^\top \mathbf{x} \end{bmatrix},$$

where $\boldsymbol{\alpha}_n$ is the n th row of \mathbf{A} , then we have that

$$(A.13) \quad \|\mathbf{A}\mathbf{x}\|_2 = \sqrt{(\boldsymbol{\alpha}_1^\top \mathbf{x})^2 + \cdots + (\boldsymbol{\alpha}_{d-k}^\top \mathbf{x})^2} \leq |\boldsymbol{\alpha}_1^\top \mathbf{x}| + \cdots + |\boldsymbol{\alpha}_{d-k}^\top \mathbf{x}| \leq (d-k)\|\mathbf{x}\|_2.$$

This results in

$$(A.14) \quad \sup_{\substack{\mathbf{x} \in \mathbb{R}^d \\ (\mathbf{A}, \mathbf{t}) \in \Xi_k}} |h_{\mathbf{x}}(\mathbf{A}, \mathbf{t})|(1 + \|\mathbf{x}\|_2)^{-n_L} < \infty.$$

This bound implies that the operator is actually well-defined on the larger space $L^1_{n_L}(\mathbb{R}^d) \supset L^*(\mathcal{S}(\mathbb{R}^d))$. In fact, since $h_{\mathbf{x}} \in C_{0,\text{iso}}(\Xi_k)$ we have for any $f \in L^1_{n_L}(\mathbb{R}^d)$ that

$$(A.15) \quad L_{\mathcal{R}_k}^{\dagger*} \{f\} = \int_{\mathbb{R}^d} f(\mathbf{x}) h_{\mathbf{x}}(\cdot) d\mathbf{x} \in C_{0,\text{iso}}(\Xi_k).$$

This, combined with (A.14), ensures that the operator

$$(A.16) \quad L_{\mathcal{R}_k}^{\dagger*} : (L^1_{n_L}(\mathbb{R}^d), \|\cdot\|_{L^1_{n_L}}) \rightarrow (C_{0,\text{iso}}(\Xi_k), \|\cdot\|_{L^\infty})$$

is continuous and, subsequently, that its adjoint

$$(A.17) \quad (L_{\mathcal{R}_k}^\dagger)^* = L_{\mathcal{R}_k}^\dagger = (\text{Id} - P_{\mathcal{P}_{n_L}(\mathbb{R}^d)}) L^{-1} \mathcal{R}_k^* : (\mathcal{M}_{\text{iso}}(\Xi_k), \|\cdot\|_{\mathcal{M}}) \rightarrow (L_{-n_L}^\infty(\mathbb{R}^d), \|\cdot\|_{L_{-n_L}^\infty})$$

is also continuous.

We now prove the right-inverse property (3.5). Recall that $L^{*-1}L^* = \text{Id}$ on $\mathcal{S}(\mathbb{R}^d)$. Thus, by duality, we have the identity

$$(A.18) \quad LL^{-1} = \text{Id} \text{ on } \mathcal{S}'(\mathbb{R}^d).$$

Since $\mathcal{M}_{\text{iso}}(\Xi_k)$ continuously embeds into \mathcal{S}'_k (Proposition 2.7), given $u \in \mathcal{M}_{\text{iso}}(\Xi_k)$, we have that

$$(A.19) \quad \begin{aligned} L_{\mathcal{R}_k} L_{\mathcal{R}_k}^\dagger \{u\} &= K_{d-k} \mathcal{R}_k L (\text{Id} - P_{\mathcal{P}_{n_L}(\mathbb{R}^d)}) L^{-1} \mathcal{R}_k^* \{u\} \\ &= K_{d-k} \mathcal{R}_k LL^{-1} \mathcal{R}_k^* \{u\} - K_{d-k} \mathcal{R}_k L \underbrace{\left\{ P_{\mathcal{P}_{n_L}(\mathbb{R}^d)} \left\{ L^{-1} \mathcal{R}_k^* \{u\} \right\} \right\}}_{=0} \\ &= K_{d-k} \mathcal{R}_k \mathcal{R}_k^* \{u\} \\ &= u, \end{aligned}$$

where the last equality follows from Proposition 2.7.

Part (II): The pseudo-left-inverse property (3.6). Observe that, for any $f \in L_{-n_L}^\infty(\mathbb{R}^d)$,

$$(A.20) \quad L^{-1}L\{f\} = f + p$$

for some $p \in \mathcal{P}_{n_L}(\mathbb{R}^d)$. Therefore, given $f \in \mathcal{M}_L^k(\mathbb{R}^d) \subset L_{-n_L}^\infty(\mathbb{R}^d)$, we have that

$$(A.21) \quad \begin{aligned} L_{\mathcal{R}_k}^\dagger L_{\mathcal{R}_k} \{f\} &= (\text{Id} - P_{\mathcal{P}_{n_L}(\mathbb{R}^d)}) L^{-1} \mathcal{R}_k^* K_{d-k} \mathcal{R}_k L \{f\} \\ &= (\text{Id} - P_{\mathcal{P}_{n_L}(\mathbb{R}^d)}) L^{-1} L \{f\} \\ &= (\text{Id} - P_{\mathcal{P}_{n_L}(\mathbb{R}^d)}) \{f + p\} \\ &= f + p - \underbrace{P_{\mathcal{P}_{n_L}(\mathbb{R}^d)} \{f\} - P_{\mathcal{P}_{n_L}(\mathbb{R}^d)} \{p\}}_{=p} \\ &= (\text{Id} - P_{\mathcal{P}_{n_L}(\mathbb{R}^d)}) \{f\}. \end{aligned}$$

Part (III): The form (3.8), stability (3.9), and continuity (3.10) of the kernel. From (A.17) we immediately see that the kernel takes the form

$$(A.22) \quad g_{\mathbf{A}, \mathbf{t}}(\mathbf{x}) = \rho_L(\mathbf{A}\mathbf{x} - \mathbf{t}) - \sum_{|\mathbf{n}| \leq n_L} \langle m_{\mathbf{n}}^*, \rho_L(\mathbf{A}(\cdot) - \mathbf{t}) \rangle m_{\mathbf{n}}(\mathbf{x}),$$

where $\langle m_{\mathbf{n}}^*, \rho_L(\mathbf{A}(\cdot) - \mathbf{t}) \rangle$ is well-defined since $m_{\mathbf{n}}^* \in \mathcal{S}(\mathbb{R}^d)$. Next, we note that the kernel of $L_{\mathcal{R}_k}^\dagger$ is the “transpose” of the kernel of $L_{\mathcal{R}_k}^\dagger$. Consequently, we have the equality $g_{\mathbf{A}, \mathbf{t}}(\mathbf{x}) = h_{\mathbf{x}}(\mathbf{A}, \mathbf{t})$. Therefore, (A.14) is equivalent to the stability bound

$$(A.23) \quad \sup_{\substack{\mathbf{x} \in \mathbb{R}^d \\ (\mathbf{A}, \mathbf{t}) \in \Xi_k}} |g_{\mathbf{A}, \mathbf{t}}(\mathbf{x})| (1 + \|\mathbf{x}\|_2)^{-n_L} < \infty.$$

Finally, in Part (I) of the proof we showed that $h_{\mathbf{x}} \in C_{0, \text{iso}}(\Xi_k)$, which proves (3.10). ■

Appendix B. Proof of Theorem 3.7.

Proof.

1. From (3.5) in Theorem 3.6, we readily deduce that

$$(B.1) \quad \begin{aligned} L_{\mathcal{R}_k}^\dagger &: \mathcal{M}_{\text{iso}}(\Xi_k) \rightarrow \mathcal{V}, \\ L_{\mathcal{R}_k} &: \mathcal{V} \rightarrow \mathcal{M}_{\text{iso}}(\Xi_k) \end{aligned}$$

are continuous bijections. Therefore, if we equip \mathcal{V} with the norm in (3.12), \mathcal{V} is isometrically isomorphic to $\mathcal{M}_{\text{iso}}(\Xi_k)$.

2. To check that the sum is direct, we must verify that $\mathcal{V} \cap \mathcal{P}_{n_L}(\mathbb{R}^d) = \{0\}$. By construction, $\mathcal{V} \subset \mathcal{M}_L^k(\mathbb{R}^d)$. Therefore, by (3.6) in Theorem 3.6,

$$(B.2) \quad L_{\mathcal{R}_k}^\dagger L_{\mathcal{R}_k} = \text{Id} - P_{\mathcal{P}_{n_L}(\mathbb{R}^d)} \text{ on } \mathcal{V}.$$

Yet, from item 1,

$$(B.3) \quad L_{\mathcal{R}_k}^\dagger L_{\mathcal{R}_k} = \text{Id} \text{ on } \mathcal{V}.$$

Thus, $(\text{Id} - P_{\mathcal{P}_{n_L}(\mathbb{R}^d)}) (\mathcal{M}_L^k(\mathbb{R}^d)) = \mathcal{V}$. Consequently, \mathcal{V} and $\mathcal{P}_{n_L}(\mathbb{R}^d)$ are complementary Banach subspaces of $\mathcal{M}_L^k(\mathbb{R}^d)$ and so $\mathcal{V} \cap \mathcal{P}_{n_L}(\mathbb{R}^d) = \{0\}$.

3. Since \mathcal{V} and $\mathcal{P}_{n_L}(\mathbb{R}^d)$ are complementary Banach subspaces of $\mathcal{M}_L^k(\mathbb{R}^d)$, we can decompose any $f \in \mathcal{M}_L^k(\mathbb{R}^d)$ as

$$(B.4) \quad \begin{aligned} f &= (\text{Id} - P_{\mathcal{P}_{n_L}(\mathbb{R}^d)})\{f\} + P_{\mathcal{P}_{n_L}(\mathbb{R}^d)}\{f\} \\ &= L_{\mathcal{R}_k}^\dagger L_{\mathcal{R}_k}\{f\} + P_{\mathcal{P}_{n_L}(\mathbb{R}^d)}\{f\} \\ &= L_{\mathcal{R}_k}^\dagger\{u\} + p, \end{aligned}$$

where the second line follows from (3.6) in Theorem 3.6. Therefore, we can equip $\mathcal{M}_L^k(\mathbb{R}^d)$ with the composite norm

$$(B.5) \quad \begin{aligned} \|f\|_{\mathcal{M}_L^k} &:= \|L_{\mathcal{R}_k}^\dagger\{u\}\|_{\mathcal{V}} + \|p\|_{\mathcal{P}_{n_L}} \\ &= \|L_{\mathcal{R}_k}^\dagger L_{\mathcal{R}_k} f\|_{\mathcal{V}} + \|P_{\mathcal{P}_{n_L}(\mathbb{R}^d)} f\|_{\mathcal{P}_{n_L}} \\ &= \underbrace{\|L_{\mathcal{R}_k} L_{\mathcal{R}_k}^\dagger L_{\mathcal{R}_k} f\|_{\mathcal{M}}}_{=\text{Id}} + \|P_{\mathcal{P}_{n_L}(\mathbb{R}^d)} f\|_{\mathcal{P}_{n_L}} \\ &= \|L_{\mathcal{R}_k} f\|_{\mathcal{M}} + \|P_{\mathcal{P}_{n_L}(\mathbb{R}^d)} f\|_{\mathcal{P}_{n_L}}. \end{aligned}$$

This norm is an isometric isomorphism with $\mathcal{V} \times \mathcal{P}_{n_L}(\mathbb{R}^d)$ by design and thus an isometric isomorphism with $\mathcal{M}_{\text{iso}}(\Xi_k) \times \mathcal{P}_{n_L}(\mathbb{R}^d)$ by item 1.

4. We first note that, in order to equip a Banach space with a weak* topology, it must be identifiable as the dual of some primary Banach space.

Next, notice that

$$(B.6) \quad \langle m_n, f \rangle = \langle m_n, L^* \varphi \rangle = \langle L m_n, \varphi \rangle = \langle 0, \varphi \rangle = 0$$

for all $f = L^* \varphi \in L^*(\mathcal{S}(\mathbb{R}^d))$ with $\varphi \in \mathcal{S}(\mathbb{R}^d)$ and $|n| \leq n_L$, where we took advantage of the null space property of L (Remark 3.2). Therefore, $P_{\mathcal{P}_{n_L}(\mathbb{R}^d)}^* \{f\} = 0$ for all $f \in L^*(\mathcal{S}(\mathbb{R}^d))$.




Figure 1. Relationships between the function spaces.

This shows that the operator $L_{\mathcal{R}_k}^*$ (which, the reader can check, maps $\mathcal{S}_k \rightarrow L^*(\mathcal{S}(\mathbb{R}^d))$) is such that

$$\begin{aligned}
 L_{\mathcal{R}_k}^{\dagger*} L_{\mathcal{R}_k}^* \{\psi\} &= \mathcal{R}_k L^{-1*} (\text{Id} - P_{\mathcal{P}_{n_L}(\mathbb{R}^d)}^*) L^* \underbrace{\mathcal{R}_k^* K_{d-k} \{\psi\}}_{\in \mathcal{S}(\mathbb{R}^d)} \\
 &= \mathcal{R}_k L^{-1*} L^* \mathcal{R}_k^* K_{d-k} \{\psi\} \\
 &= \mathcal{R}_k \mathcal{R}_k^* K_{d-k} \{\psi\} \\
 (B.7) \qquad \qquad \qquad &= \psi
 \end{aligned}$$

for all $\psi \in \mathcal{S}_k$. Thus, $L_{\mathcal{R}_k}^{\dagger*}$ is a left-inverse of $L_{\mathcal{R}_k}^*$. This implies that the normed space $(\mathcal{S}_k, \|\cdot\|_{L^\infty})$ is (isometrically) isomorphic to the normed space $(L^*(\mathcal{S}(\mathbb{R}^d)), \|\cdot\|_{\mathcal{U}})$ with $\|u\|_{\mathcal{U}} := \|L_{\mathcal{R}_k}^{\dagger*} \{u\}\|_{L^\infty}$. Recall that $C_{0,\text{iso}}(\Xi_k) = \overline{(\mathcal{S}_k, \|\cdot\|_{L^\infty})}$ from (2.20) and define the Banach space $\mathcal{U} := \overline{(L^*(\mathcal{S}(\mathbb{R}^d)), \|\cdot\|_{\mathcal{U}})}$. We can now invoke the bounded linear transformation theorem [49, Theorem I.7, p. 9] on both $L_{\mathcal{R}_k}^{\dagger*}$ and $L_{\mathcal{R}_k}^*$ to find that these operators have the continuous extensions

$$\begin{aligned}
 (B.8) \qquad \qquad \qquad L_{\mathcal{R}_k}^{\dagger*} : \mathcal{U} &\rightarrow C_{0,\text{iso}}(\Xi_k), \\
 L_{\mathcal{R}_k}^* : C_{0,\text{iso}}(\Xi_k) &\rightarrow \mathcal{U},
 \end{aligned}$$

which establishes that \mathcal{U} and $C_{0,\text{iso}}(\Xi_k)$ are (isometrically) isomorphic Banach spaces. From item 3, we know that $\mathcal{M}_L^k(\mathbb{R}^d)$ is isometrically isomorphic to $\mathcal{M}_{\text{iso}}(\Xi_k) \times \mathcal{P}_{n_L}(\mathbb{R}^d)$. Since $\mathcal{M}_{\text{iso}}(\Xi_k) = (C_{0,\text{iso}}(\Xi_k))'$ (see (2.19)) and $\mathcal{P}_{n_L}(\mathbb{R}^d) = (\mathcal{P}_{n_L}(\mathbb{R}^d))''$ (since $\mathcal{P}_{n_L}(\mathbb{R}^d)$ is finite-dimensional and hence reflexive), we see that there is a predual of $\mathcal{M}_L^k(\mathbb{R}^d)$ that is isometrically isomorphic to $C_{0,\text{iso}}(\Xi_k) \times (\mathcal{P}_{n_L}(\mathbb{R}^d))'$.

From (B.8) and item 1, we have the diagram in Figure 1. Therein, we see that $\mathcal{V} = \mathcal{U}'$ and so $\mathcal{X} = \mathcal{U} \oplus (\mathcal{P}_{n_L}(\mathbb{R}^d))'$ is such that $\mathcal{X}' = \mathcal{M}_L^k(\mathbb{R}^d)$. To complete the proof, we need to establish that $\delta(\cdot - \mathbf{x}_0) \in \mathcal{U} \oplus (\mathcal{P}_{n_L}(\mathbb{R}^d))'$ [49, Theorem IV.20, p. 114]. Clearly, $\delta(\cdot - \mathbf{x}_0) \in (\mathcal{P}_{n_L}(\mathbb{R}^d))'$. Therefore, we only need to check that $\delta(\cdot - \mathbf{x}_0) \in \mathcal{U}$. This is equivalent to $L_{\mathcal{R}_k}^{\dagger*} \{\delta(\cdot - \mathbf{x}_0)\} \in C_{0,\text{iso}}(\Xi_k)$. Since $L_{\mathcal{R}_k}^{\dagger*} \{\delta(\cdot - \mathbf{x}_0)\}(\mathbf{A}, \mathbf{t}) = g_{\mathbf{A},\mathbf{t}}(\mathbf{x}_0)$, the result follows from (3.10) in Theorem 3.6. ■

REFERENCES

- [1] E. ABBE, E. B. ADSERA, AND T. MISIAKIEWICZ, *The merged-staircase property: A necessary and nearly sufficient condition for SGD learning of sparse functions on two-layer neural networks*, in Proceedings of the Conference on Learning Theory, 2022, pp. 4782–4887.
- [2] C. ANIL, J. LUCAS, AND R. GROSSE, *Sorting out Lipschitz function approximation*, in Proceedings of the International Conference on Machine Learning, 2019, pp. 291–301.
- [3] S. AZIZNEJAD AND M. UNSER, *Multikernel regression with sparsity constraint*, SIAM J. Math. Data Sci., 3 (2021), pp. 201–224, <https://doi.org/10.1137/20M1318882>.
- [4] F. BACH, *Breaking the curse of dimensionality with convex neural networks*, J. Mach. Learn. Res., 18 (2017), pp. 629–681.
- [5] F. BARTOLUCCI, E. DE VITO, L. ROSASCO, AND S. VIGOGNA, *Understanding neural networks with reproducing kernel Banach spaces*, Appl. Comput. Harmon. Anal., 62 (2023), pp. 194–236, <https://doi.org/10.1016/j.acha.2022.08.006>.
- [6] C. BOYER, A. CHAMBOLLE, Y. DE CASTRO, V. DUVAL, F. DE GOURNAY, AND P. WEISS, *On representer theorems and convex regularization*, SIAM J. Optim., 29 (2019), pp. 1260–1281, <https://doi.org/10.1137/18M1200750>.
- [7] K. BREDIES AND M. CARIONI, *Sparsity of solutions for variational inverse problems with finite-dimensional data*, Calc. Var. Partial Differential Equations, 59 (2020), <https://doi.org/10.1007/s00526-019-1658-1>.
- [8] A. COHEN, I. DAUBECHIES, R. DEVORE, G. KERKYACHARIAN, AND D. PICARD, *Capturing ridge functions in high dimensions from point queries*, Constr. Approx., 35 (2012), pp. 225–243, <https://doi.org/10.1007/s00365-011-9147-6>.
- [9] A. S. DALALYAN, A. JUDITSKY, AND V. SPOKOINY, *A new algorithm for estimating the effective dimension-reduction subspace*, J. Mach. Learn. Res., 9 (2008), pp. 1647–1678.
- [10] C. DE BOOR AND R. E. LYNCH, *On splines and their minimum properties*, J. Math. Mech., 15 (1966), pp. 953–969.
- [11] R. DEVORE, R. D. NOWAK, R. PARHI, AND J. W. SIEGEL, *Weighted variation spaces and approximation by shallow ReLU networks*, Appl. Comput. Harmon. Anal., 74 (2025), 101713, <https://doi.org/10.1016/j.acha.2024.101713>.
- [12] J. DUCHON, *Splines minimizing rotation-invariant semi-norms in Sobolev spaces*, in Constructive Theory of Functions of Several Variables, Springer, Berlin, 1977, pp. 85–100.
- [13] S. D. FISHER AND J. W. JEROME, *Spline solutions to L^1 extremal problems in one and several variables*, J. Approx. Theory, 13 (1975), pp. 73–83, [https://doi.org/10.1016/0021-9045\(75\)90016-7](https://doi.org/10.1016/0021-9045(75)90016-7).
- [14] G. B. FOLLAND, *Real Analysis: Modern Techniques and Their Applications*, 2nd ed., Pure Appl. Math. (Hoboken), Hoboken, New Jersey, 1999.
- [15] K. FUKUMIZU, F. R. BACH, AND M. I. JORDAN, *Dimensionality reduction for supervised learning with reproducing kernel Hilbert spaces*, J. Mach. Learn. Res., 5 (2004), pp. 73–99.
- [16] I. M. GEL'FAND, M. I. GRAEV, AND N. Y. VILENKIN, *Generalized Functions, Vol. 5: Integral Geometry and Representation Theory*, Academic Press, New York, 1966.
- [17] I. M. GEL'FAND AND G. E. SHILOV, *Generalized Functions, Vol. I: Properties and Operations*, Academic Press, New York, 1964.
- [18] B. GHORBANI, S. MEI, T. MISIAKIEWICZ, AND A. MONTANARI, *When do neural networks outperform kernel methods?*, Adv. Neural Inf. Process. Syst., 33 (2020), pp. 14820–14830.
- [19] X. GLOROT, A. BORDES, AND Y. BENGIO, *Deep sparse rectifier neural networks*, in Proceedings of the 14th International Conference on Artificial Intelligence and Statistics, 2011, pp. 315–323.
- [20] F. B. GONZALEZ, *On the range of the Radon d -plane transform and its dual*, Trans. Amer. Math. Soc., 327 (1991), pp. 601–619, <https://doi.org/10.2307/2001816>.
- [21] I. GOODFELLOW, D. WARDE-FARLEY, M. MIRZA, A. COURVILLE, AND Y. BENGIO, *Maxout networks*, in Proceedings of the International Conference on Machine Learning, 2013, pp. 1319–1327.
- [22] P. GROHS, S. KEIPER, G. KUTYNIOK, AND M. SCHÄFER, *α -molecules*, Appl. Comput. Harmon. Anal., 41 (2016), pp. 297–336, <https://doi.org/10.1016/j.acha.2015.10.009>.
- [23] C. GULCEHRE, K. CHO, R. PASCANU, AND Y. BENGIO, *Learned-norm pooling for deep feedforward and recurrent neural networks*, in Machine Learning and Knowledge Discovery in Databases: European Conference, Proceedings, Part I, Springer, New York, 2014, pp. 530–546.

- [24] L. HUANG, X. LIU, B. LANG, A. YU, Y. WANG, AND B. LI, *Orthogonal weight normalization: Solution to optimization over multiple dependent Stiefel manifolds in deep neural networks*, in Proceedings of the 32nd AAAI Conference on Artificial Intelligence, 2018.
- [25] L. HUANG, J. QIN, Y. ZHOU, F. ZHU, L. LIU, AND L. SHAO, *Normalization techniques in training DNNs: Methodology, analysis and application*, IEEE Trans. Pattern Anal. Mach. Intell., 45 (2023), pp. 10173–10196, <https://doi.org/10.1109/TPAMI.2023.3250241>.
- [26] A. JACOT, F. GABRIEL, AND C. HONGLER, *Neural tangent kernel: Convergence and generalization in neural networks*, Adv. Neural Inf. Process. Syst., 31 (2018).
- [27] F. KEINERT, *Inversion of k -plane transforms and applications in computer tomography*, SIAM Rev., 31 (1989), pp. 273–298, <https://doi.org/10.1137/1031051>.
- [28] S. KEIPER, *Approximation of generalized ridge functions in high dimensions*, J. Approx. Theory, 245 (2019), pp. 101–129, <https://doi.org/10.1016/j.jat.2019.04.006>.
- [29] G. S. KIMELDORF AND G. WAHBA, *Some results on Tchebycheffian spline functions*, J. Math. Anal. Appl., 33 (1971), pp. 82–95, [https://doi.org/10.1016/0022-247X\(71\)90184-3](https://doi.org/10.1016/0022-247X(71)90184-3).
- [30] V. KŮRKOVÁ, P. C. KAINEN, AND V. KREINOVICH, *Estimates of the number of hidden units and variation with respect to half-spaces*, Neural Networks, 10 (1997), pp. 1061–1068, [https://doi.org/10.1016/S0893-6080\(97\)00028-2](https://doi.org/10.1016/S0893-6080(97)00028-2).
- [31] V. KŮRKOVÁ AND M. SANGUINETI, *Bounds on rates of variable-basis and neural-network approximation*, IEEE Trans. Inform. Theory, 47 (2001), pp. 2659–2665, <https://doi.org/10.1109/18.945285>.
- [32] K.-C. LI, *Sliced inverse regression for dimension reduction*, J. Amer. Statist. Assoc., 86 (1991), pp. 316–327, <https://doi.org/10.1080/01621459.1991.10475035>.
- [33] Q. LI, S. HAQUE, C. ANIL, J. LUCAS, R. B. GROSSE, AND J.-H. JACOBSEN, *Preventing gradient attenuation in Lipschitz constrained convolutional networks*, Adv. Neural Inf. Process. Syst., 32 (2019).
- [34] R. R. LIN, H. Z. ZHANG, AND J. ZHANG, *On reproducing kernel Banach spaces: Generic definitions and unified framework of constructions*, Acta Math. Sin. (Engl. Ser.), 38 (2022), pp. 1459–1483, <https://doi.org/10.1007/s10114-022-1397-7>.
- [35] H. LIU AND W. LIAO, *Learning functions varying along a central subspace*, SIAM J. Math. Data Sci., 6 (2024), pp. 343–371, <https://doi.org/10.1137/23M1557751>.
- [36] A. MARKOE, *Analytic Tomography*, Encyclopedia Math. Appl. 106, Cambridge University Press, Cambridge, 2006, <https://doi.org/10.1017/CBO9780511530012>.
- [37] H. N. MHASKAR, *On the tractability of multivariate integration and approximation by neural networks*, J. Complexity, 20 (2004), pp. 561–590, <https://doi.org/10.1016/j.jco.2003.11.004>.
- [38] H. N. MHASKAR, *Approximation by Non-Symmetric Networks for Cross-Domain Learning*, preprint, [arXiv:2305.03890](https://arxiv.org/abs/2305.03890), 2023.
- [39] H. N. MHASKAR AND C. A. MICCHELLI, *Approximation by superposition of sigmoidal and radial basis functions*, Adv. Appl. Math., 13 (1992), pp. 350–373, [https://doi.org/10.1016/0196-8858\(92\)90016-P](https://doi.org/10.1016/0196-8858(92)90016-P).
- [40] H. N. MHASKAR AND C. A. MICCHELLI, *Degree of approximation by neural and translation networks with a single hidden layer*, Adv. Appl. Math., 16 (1995), pp. 151–183, <https://doi.org/10.1006/aama.1995.1008>.
- [41] G. ONGIE, R. WILETT, D. SOUDRY, AND N. SREBRO, *A function space view of bounded norm infinite width ReLU nets: The multivariate case*, in Proceedings of the International Conference on Learning Representations, 2020, pp. 1–24.
- [42] R. PARHI AND R. D. NOWAK, *The role of neural network activation functions*, IEEE Signal Process. Lett., 27 (2020), pp. 1779–1783, <https://doi.org/10.1109/LSP.2020.3027517>.
- [43] R. PARHI AND R. D. NOWAK, *Banach space representer theorems for neural networks and ridge splines*, J. Mach. Learn. Res., 22 (2021).
- [44] R. PARHI AND R. D. NOWAK, *What kinds of functions do deep neural networks learn? Insights from variational spline theory*, SIAM J. Math. Data Sci., 4 (2022), pp. 464–489, <https://doi.org/10.1137/21M1418642>.
- [45] R. PARHI AND R. D. NOWAK, *Deep learning meets sparse regularization: A signal processing perspective*, IEEE Signal Process. Mag., 40 (2023), pp. 63–74, <https://doi.org/10.1109/MSP.2023.3286988>.
- [46] R. PARHI AND R. D. NOWAK, *Near-minimax optimal estimation with shallow ReLU neural networks*, IEEE Trans. Inform. Theory, 69 (2023), pp. 1125–1140, <https://doi.org/10.1109/TIT.2022.3208653>.
- [47] R. PARHI AND M. UNSER, *Distributional extension and invertibility of the k -plane transform and its dual*, SIAM J. Math. Anal., 56 (2024), pp. 4662–4686, <https://doi.org/10.1137/23M1556721>.

- [48] S. PARKINSON, G. ONGIE, AND R. WILLETT, *ReLU Neural Networks with Linear Layers Are Biased Towards Single- and Multi-Index Models*, preprint, [arXiv:2305.15598](https://arxiv.org/abs/2305.15598), 2023.
- [49] M. REED AND B. SIMON, *Methods of Modern Mathematical Physics I: Functional Analysis*, Academic Press, New York, 1972.
- [50] S. ROSSET, G. SWIRSZCZ, N. SREBRO, AND J. ZHU, ℓ_1 regularization in infinite dimensional feature spaces, in Proceedings of the 20th Annual Conference on Learning Theory, Springer, 2007, pp. 544–558.
- [51] B. RUBIN, *Inversion of k -plane transforms via continuous wavelet transforms*, J. Math. Anal. Appl., 220 (1998), pp. 187–203, <https://doi.org/10.1006/jmaa.1997.5852>.
- [52] W. RUDIN, *Functional Analysis*, 2nd ed., Internat. Ser. Pure Appl. Math., McGraw-Hill, New York, 1991.
- [53] S. G. SAMKO, A. A. KILBAS, AND O. I. MARICHEV, *Fractional integrals and derivatives*, in Theory and Applications, S. M. Nikol'skiĭ, ed., Gordon and Breach, New York, 1993, pp. 143–157.
- [54] P. SAVARESE, I. EVRON, D. SOUDRY, AND N. SREBRO, *How do infinite width bounded norm networks look in function space?* in Proceedings of the 32nd Annual Conference on Learning Theory, 2019, pp. 2667–2690.
- [55] B. SCHÖLKOPF AND A. J. SMOLA, *Learning with Kernels: Support Vector Machines, Regularization, Optimization, and Beyond*, Adapt. Comput. Mach. Learn., MIT Press, Cambridge, MA, 2002.
- [56] J. SHENOUDA, R. PARHI, K. LEE, AND R. D. NOWAK, *Variation spaces for multi-output neural networks: Insights on multi-task learning and network compression*, J. Mach. Learn. Res., 25 (2024), pp. 1–40.
- [57] J. W. SIEGEL AND J. XU, *Characterization of the variation spaces corresponding to shallow neural networks*, Constr. Approx., 57 (2023), pp. 1109–1132, <https://doi.org/10.1007/s00365-023-09626-4>.
- [58] J. W. SIEGEL AND J. XU, *Sharp bounds on the approximation rates, metric entropy, and n -widths of shallow neural networks*, Found. Comput. Math., 24 (2024), pp. 481–537, <https://doi.org/10.1007/s10208-022-09595-3>.
- [59] K. T. SMITH, D. C. SOLMON, AND S. L. WAGNER, *Practical and mathematical aspects of the problem of reconstructing objects from radiographs*, Bull. Amer. Math. Soc., 83 (1977), pp. 1227–1270, <https://doi.org/10.1090/S0002-9904-1977-14406-6>.
- [60] D. C. SOLMON, *The X-ray transform*, J. Math. Anal. Appl., 56 (1976), pp. 61–83, [https://doi.org/10.1016/0022-247X\(76\)90008-1](https://doi.org/10.1016/0022-247X(76)90008-1).
- [61] S. SONODA, I. ISHIKAWA, AND M. IKEDA, *A unified Fourier slice method to derive ridgelet transform for a variety of depth-2 neural networks*, J. Stat. Plan. Inference., 233 (2024), 106184, <https://doi.org/10.1016/j.jspi.2024.106184>.
- [62] L. SPEK, T. J. HEERINGA, AND C. BRUNE, *Duality for Neural Networks Through Reproducing Kernel Banach Spaces*, preprint, [arXiv:2211.05020](https://arxiv.org/abs/2211.05020), 2022.
- [63] I. STEINWART, *Sparseness of support vector machines*, J. Mach. Learn. Res., 4 (2003), pp. 1071–1105.
- [64] F. TRÈVES, *Topological Vector Spaces, Distributions and Kernels*, Academic Press, New York-London, 1967.
- [65] M. UNSER, *A unifying representer theorem for inverse problems and machine learning*, Found. Comput. Math., 21 (2021), pp. 941–960, <https://doi.org/10.1007/s10208-020-09472-x>.
- [66] M. UNSER, *From kernel methods to neural networks: A unifying variational formulation*, Found. Comput. Math., 24 (2024), pp. 1779–1818, <https://doi.org/10.1007/s10208-023-09624-9>.
- [67] M. UNSER, *Ridges, neural networks, and the Radon transform*, J. Mach. Learn. Res., 24 (2023).
- [68] M. UNSER AND S. AZIZNEJAD, *Convex optimization in sums of Banach spaces*, Appl. Comput. Harmon. Anal., 56 (2022), pp. 1–25, <https://doi.org/10.1016/j.acha.2021.07.002>.
- [69] M. UNSER, J. FAGEOT, AND J. P. WARD, *Splines are universal solutions of linear inverse problems with generalized TV regularization*, SIAM Rev., 59 (2017), pp. 769–793, <https://doi.org/10.1137/16M1061199>.
- [70] G. WAHBA, *Spline Models for Observational Data*, SIAM, Philadelphia, 1990.
- [71] H. WENDLAND, *Scattered Data Approximation*, Cambridge Monogr. Appl. Comput. Math. 17, Cambridge University Press, Cambridge, 2005.
- [72] H. ZHANG, Y. XU, AND J. ZHANG, *Reproducing kernel Banach spaces for machine learning*, J. Mach. Learn. Res., 10 (2009), pp. 3520–3527.