

From Kernel Methods to Neural Networks: A Unifying Variational Formulation

Michael Unser

the date of receipt and acceptance should be inserted later

Abstract The minimization of a data-fidelity term and an additive regularization functional gives rise to a powerful framework for supervised learning. In this paper, we present a unifying regularization functional that depends on an operator L and on a generic Radon-domain norm. We establish the existence of a minimizer and give the parametric form of the solution(s) under very mild assumptions. When the norm is Hilbertian, the proposed formulation yields a solution that involves radial-basis functions and is compatible with the classical methods of machine learning. By contrast, for the total-variation norm, the solution takes the form of a two-layer neural network with an activation function that is determined by the regularization operator. In particular, we retrieve the popular ReLU networks by letting the operator be the Laplacian. We also characterize the solution for the intermediate regularization norms $\|\cdot\| = \|\cdot\|_{L_p}$ with $p \in (1, 2]$. Our framework offers guarantees of universal approximation for a broad family of regularization operators or, equivalently, for a wide variety of shallow neural networks, including the cases (such as ReLU) where the activation function is increasing polynomially. It also explains the favorable role of bias and skip connections in neural architectures.

Mathematics Subject Classification (2010) 44A12 · 46N10 · 47A52 · 68T07

Keywords machine learning · convex optimization · regularization · representer theorem · kernel methods · neural networks · Banach space

Communicated by Carola-Bibiane Schönlieb

The research leading to these results has received funding from the European Research Council under Grant ERC-2020-AdG FunLearn-101020573.

Biomedical Imaging Group, École polytechnique fédérale de Lausanne (EPFL), Station 17, CH-1015, Lausanne, Switzerland (michael.unser@epfl.ch).

Address(es) of author(s) should be given

1 Introduction

Regularization theory constitutes a powerful framework for the derivation of algorithms for supervised learning [14,41,42]. Given a series of data points $(\mathbf{x}_m, y_m) \in \mathbb{R}^d \times \mathbb{R}$, $m = 1, \dots, M$, the basic problem (regression) is to find a mapping $f : \mathbb{R}^d \rightarrow \mathbb{R}$ such that $f(\mathbf{x}_m) \approx y_m$, without overfitting. The standard paradigm is to let f be the minimizer of a cost that consists of a data-fidelity term and an additive regularization functional [8]. The minimization proceeds over a prescribed class \mathcal{H} of candidate functions. One usually distinguishes between the parametric approaches (e.g., neural networks), where $\mathcal{H} = \mathcal{H}_\theta$ is a family of functions specified by a finite set of parameters $\theta \in \Theta$ (e.g., the weights of the network), and the nonparametric ones, where the properties of the solution are controlled by the regularization functional. The focus of this paper is on the nonparametric techniques. They rely on *functional optimization*, which means that the minimization proceeds over a space of functions rather than over a set of parameters. The regularization is usually chosen to be an increasing function of the norm associated with a particular Banach space, which results in a well-posed problem [10,9,56].

The functional-optimization point of view is often constructive, in that it suggests or supports explicit learning architectures. For instance, the choice of the Hilbertian regularization $R(f) = \|f\|_{\mathcal{H}}^2$ where \mathcal{H} is a reproducing kernel Hilbert space (RKHS) results in a closed-form solution that is a linear combination of kernels positioned on the data [62,7]. In fact, the RKHS setting yields a generic class of estimators that is compatible with the classical kernel-based methods of machine learning, including support vector machines [62,41,50,49,1,56]. Likewise, adaptive kernel methods are justifiable from the minimization of a generalized total-variation norm, which favors sparse representations [11,12,3]. These latter results actually take their root in spline theory [18,28,60]. Similarly, it has been demonstrated that shallow ReLU networks are solutions of functional-optimization problems with an appropriate regularization. One way to achieve this is to start from an explicit parameterization of an infinite-width network [4] (the reverse engineering/synthesis approach). Another way is to consider a regularization operator that is matched to the neuronal activation with a L_1 -type penalty¹; for instance, a second-order derivative for $d = 1$ [48,36] or, more generally, the Radon-domain counterpart of the Laplace operator whose Green's function is precisely a ReLU ridge [35,37,57]. Similar optimality results can be stated within the framework of reproducing-kernel Banach spaces [6], which is a formal point of view that bridges the synthesis and analysis approach of [4] and [37], respectively. Also relevant to the discussion is a variational formulation that links the ridgelet transform to the training of shallow neural networks with weight-decay regularization [53].

The second important benefit of the functional-optimization approach is that it gives insight on the approximation capabilities (expressivity) of the re-

¹ The precise formulation involves the \mathcal{M} -norm (or total variation), which is the weak form of L_1 associated with space of bounded Radon measures. In our account, we take it as the default norm for the Lebesgue exponent $p = 1$, with a slight abuse of language.

sulting learning architectures. This information is encapsulated in the choice/definition of the native space \mathcal{H} (typically, a Sobolev space), which goes hand-in-hand with the regularization functional. Roughly speaking, the native space \mathcal{H} ought to be “large enough” to allow for the approximation of any continuous function with an arbitrary degree of precision. This universal approximation property is a central theme in the theory of radial-basis functions (RBFs) [31, 63]. In machine learning, the kernel estimators that meet this approximation requirement are called *universal* [32]. When the basis functions are shifted replicates of a single template $h : \mathbb{R}^d \rightarrow \mathbb{R}$, then the condition is equivalent to h being strictly positive definite, which means that its Fourier transform is real-valued symmetric, and (strictly) positive [13]. Similar guarantees of universal approximation exist for (shallow) neural networks under mild conditions on the activation functions [16, 25, 30, 5, 39]. The main difference with the RKHS framework, however, is that the universality results for neural nets usually make the assumption that the input domain is a compact subset of \mathbb{R}^d .

The purpose of this paper is to unify and extend these various approaches by introducing a universal regularization functional. The latter has two components: an admissible differential operator L , and an L_p -type Radon-domain norm. The resulting regularization operator is $L_R = K_{\text{rad}}RL$, where R is the Radon transform and K_{rad} the “filtering” operator of computer tomography [33]. Our main result (Theorem 5) gives the parametric form of the solution of the corresponding functional-optimization problems under minimal hypotheses. For $p = 2$, the outcome of our theorem is compatible with the type of kernel expansions (RBFs) of classical machine learning for which there is a vast literature [52, 24]. For $p = 1$, the solution set is parameterized by a neural network with one hidden layer whose activation function is determined by the regularization operator. In particular, if we take L to be the Laplacian, then one retrieves the popular ReLU activation. Remarkably, the connection with neural networks also works the other way round: Parhi et al. [36, 38] could prove that the training of a shallow ReLU neural network that is sufficiently wide, with weight-decay regularization, converges to the solution of a functional optimization problem that is a special instance of the class considered in this paper.

The foundation for our characterization is an abstract representer theorem for direct-sum Banach spaces [58]. Thus, the primary effort in this paper consists in the development of a functional framework that is adapted to the Radon transform and that fulfills the hypotheses of the abstract theorem. The main contributions can be summarized as follows.

1. Construction and characterization of an extended family of native Banach spaces $\mathcal{X}'_{L_R}(\mathbb{R}^d)$ associated with a generic Radon-domain norm $\|\cdot\|_{\mathcal{X}'}$ and a differential operator L , under the general admissibility conditions stated in Definition 3 (Theorem 6).
2. Proof that (i) the sampling functionals $\delta(\cdot - \mathbf{x}_m) : \mathcal{X}'_{L_R}(\mathbb{R}^d) \rightarrow \mathbb{R}$ are weak*-continuous; and (ii) the adjoint of the regularization operator has a stable

generalized inverse $L_R^{*\dagger}$ (see Theorem 7 and accompanying explanations). These technical points are essential to the argumentation (existence of solution).

3. Extension and unification of a number of earlier optimality results for RBF expansions and neural networks. While the present setup for $p = 2$ and $L = (-\Delta)^\gamma$ is reminiscent of thin-plate splines [17, 29], the resulting solution for a fixed γ does not depend on the dimension d , which makes it easier to deploy. Likewise, our variational formulation with $\mathcal{X}' = \mathcal{M}$ extends the results of Parhi and Nowak [37] by: (i) proving that the neural network parameterization applies to *all the extreme points* of the solution set, and (ii) by covering a much broader class of activation functions, including those with polynomial growth (of degree n_0).
4. General guarantees of universality, subject to the admissibility condition in Definition 3. While the result for $p = 2$ is consistent with the known criteria for kernel estimators [32], its counterpart for neural networks ($\mathcal{X}' = \mathcal{M}$) brings in a new twist: the addition of a polynomial component. The latter, which is not present in the traditional theory [5, 39], is necessary to lift the hypothesis of a compact input domain. The two cases of greatest practical relevance are the sigmoid and the ReLU activations which, in our formulation, require the addition of a bias ($n_0 = 0$) and an affine term ($n_0 = 1$), respectively. Interestingly, the ReLU case yields a neural architecture with a skip connection akin to ResNet [22], which is highly popular in practice.

The paper is organised as follows: We start with mathematical preliminaries in Section 2. In particular, we state our criteria of admissibility for L and show how to represent its polynomial null space. In Section 3, we review the main properties of the Radon transform and specify the dual pair $(\mathcal{X}_{\text{Rad}}, \mathcal{X}'_{\text{Rad}})$ of hyper-spherical Banach spaces that enter the definition of our native spaces. We also provide formulas for the (filtered) Radon transform of RBFs and ridges (the elementary constituents of neural networks). Section 4 is devoted to the description and interpretation of our main result (Theorem 5). In particular, we draw a connection with RKHS in Section 4.2. We discuss the issue of universality in Section 4.3 and show in Section 4.4 how our framework can be extended to handle antisymmetric activations, including sigmoids. We complement our exposition in Section 4.5 with a listing of specific configurations, many of which are intimately connected to splines. The mathematical developments that support our formulation are presented in Section 5. They include the characterization of the kernel of the inverse operator $L_R^{*\dagger}$ —the enabling ingredient of our formulation—and the construction of the predual Banach space $\mathcal{X}_{L_R}(\mathbb{R}^d)$.

2 Mathematical Preliminaries

2.1 Notations

We shall consider multidimensional functions f on \mathbb{R}^d that are indexed by the variable $\mathbf{x} \in \mathbb{R}^d$. To describe their partial derivatives, we use the multi-index $\mathbf{k} = (k_1, \dots, k_d) \in \mathbb{N}^d$ (where \mathbb{N} includes 0) with the notational conventions $\mathbf{k}! = \prod_{i=1}^d k_i!$, $|\mathbf{k}| = k_1 + \dots + k_d$, $\mathbf{x}^{\mathbf{k}} = \prod_{i=1}^d x_i^{k_i}$ for any $\mathbf{x} \in \mathbb{R}^d$, and $\partial^{\mathbf{k}} f(\mathbf{x}) = \frac{\partial^{|\mathbf{k}|} f(x_1, \dots, x_d)}{\partial x_1^{k_1} \dots \partial x_d^{k_d}}$. This allows us to write the multidimensional Taylor expansion around $\mathbf{x} = \mathbf{x}_0$ of an analytical function $f : \mathbb{R}^d \rightarrow \mathbb{R}$ explicitly as

$$f(\mathbf{x}) = \sum_{n=0}^{\infty} \sum_{|\mathbf{k}|=n} \frac{\partial^{\mathbf{k}} f(\mathbf{x}_0) (\mathbf{x} - \mathbf{x}_0)^{\mathbf{k}}}{\mathbf{k}!} \quad (1)$$

where the internal summation is over all multi-indices \mathbf{k} such that $k_1 + \dots + k_d = n$.

Schwartz' space of smooth and rapidly decreasing test functions $\varphi : \mathbb{R}^d \rightarrow \mathbb{R}$ equipped with the usual Fréchet-Schwartz topology is denoted by $\mathcal{S}(\mathbb{R}^d)$. Its continuous dual is the space $\mathcal{S}'(\mathbb{R}^d)$ of tempered distributions. In this setting, the Lebesgue spaces $L_p(\mathbb{R}^d)$ for $p \in [1, \infty)$ can be specified as the completion of $\mathcal{S}(\mathbb{R}^d)$ equipped with the L_p -norm $\|\cdot\|_{L_p}$, denoted as $L_p(\mathbb{R}^d) = \overline{(\mathcal{S}(\mathbb{R}^d), \|\cdot\|_{L_p})}$. For the endpoint $p = \infty$, we have $\overline{(\mathcal{S}(\mathbb{R}^d), \|\cdot\|_{L_\infty})} = C_0(\mathbb{R}^d)$ with $\|\varphi\|_{L_\infty} = \sup_{\mathbf{x} \in \mathbb{R}^d} |\varphi(\mathbf{x})|$, which is the space of continuous functions that vanish at infinity. The continuous dual of $C_0(\mathbb{R}^d)$ is the space $\mathcal{M}(\mathbb{R}^d) = \{f \in \mathcal{S}'(\mathbb{R}^d) : \|f\|_{\mathcal{M}} < \infty\}$ of bounded Radon measures with

$$\|f\|_{\mathcal{M}} = \sup_{\varphi \in \mathcal{S}(\mathbb{R}^d) : \|\varphi\|_{L_\infty} \leq 1} \langle f, \varphi \rangle. \quad (2)$$

The latter is a superset of $L_1(\mathbb{R}^d)$, which is isometrically embedded in it, in the sense that $\|f\|_{L_1} = \|f\|_{\mathcal{M}}$ for all $f \in L_1(\mathbb{R}^d)$.

The Fourier transform of a function $\varphi \in L_1(\mathbb{R}^d)$ is defined as

$$\widehat{\varphi}(\boldsymbol{\omega}) \triangleq \mathcal{F}\{\varphi\}(\boldsymbol{\omega}) = \frac{1}{(2\pi)^d} \int_{\mathbb{R}^d} \varphi(\mathbf{x}) e^{-i\langle \boldsymbol{\omega}, \mathbf{x} \rangle} d\mathbf{x}. \quad (3)$$

Since the Fourier operator \mathcal{F} continuously maps $\mathcal{S}(\mathbb{R}^d)$ into itself, the transform can be extended by duality to the whole space $\mathcal{S}'(\mathbb{R}^d)$ of tempered distribution. Specifically, $\widehat{f} = \mathcal{F}\{f\} \in \mathcal{S}'(\mathbb{R}^d)$ is the (unique) *generalized Fourier transform* of $f \in \mathcal{S}'(\mathbb{R}^d)$ if and only if $\langle \widehat{f}, \varphi \rangle = \langle f, \widehat{\varphi} \rangle$ for all $\varphi \in \mathcal{S}(\mathbb{R}^d)$, where $\widehat{\varphi} = \mathcal{F}\{\varphi\}$ is the "classical" Fourier transform of φ defined by (3).

To control the minimal order $\alpha \geq 0$ of decay (resp., the maximal rate of growth) of functions, we use the dual pair of spaces $L_{1,\alpha}(\mathbb{R}^d)$ and $L_{\infty,-\alpha}(\mathbb{R}^d) = (L_{1,\alpha}(\mathbb{R}^d))'$. These are the Banach spaces associated with the weighted norms

$$\|f\|_{L_{1,\alpha}} \triangleq \int_{\mathbb{R}^d} (1 + \|\mathbf{x}\|)^\alpha |f(\mathbf{x})| d\mathbf{x} \quad (4)$$

$$\|f\|_{L_{\infty,-\alpha}} \triangleq \operatorname{ess\,sup}_{\mathbf{x} \in \mathbb{R}^d} (1 + \|\mathbf{x}\|)^{-\alpha} |f(\mathbf{x})|, \quad (5)$$

respectively. Specifically, the inclusion $f \in L_{\infty,-n_0}(\mathbb{R}^d)$ with $n_0 \in \mathbb{N}$ indicates that f cannot grow faster than a polynomial of degree n_0 , while the condition $f \in L_{1,\alpha}(\mathbb{R}^d)$ implies that $f(\mathbf{x})$ must be locally integrable and must decay (slightly) faster than $1/\|\mathbf{x}\|^{\alpha+d}$ as $\|\mathbf{x}\| \rightarrow \infty$.

2.2 Admissible Regularization Operators

The regularization operators L that are of interest to us are linear, shift-invariant (LSI), and isotropic. For simplicity, we shall first specify the action of L on test functions, with the understanding that the domain of the operator will be extended to some corresponding “native space” that will be identified as we progress through the paper.

Definition 1 The linear operator $L : \mathcal{S}(\mathbb{R}^d) \rightarrow \mathcal{S}'(\mathbb{R}^d)$ is said to be

- *Shift-invariant*: If $L\{\varphi(\cdot - \mathbf{x}_0)\} = L\{\varphi\}(\cdot - \mathbf{x}_0)$ for all $\varphi \in \mathcal{S}(\mathbb{R}^d)$ and $\mathbf{x}_0 \in \mathbb{R}^d$.
- *Isotropic (or rotation-invariant)*: If $L\{\varphi(\mathbf{R}_{\boldsymbol{\theta}}\cdot)\} = L\{\varphi\}(\mathbf{R}_{\boldsymbol{\theta}}\cdot)$ for all $\varphi \in \mathcal{S}(\mathbb{R}^d)$ and any rotation matrix $\mathbf{R}_{\boldsymbol{\theta}}$ on \mathbb{R}^d .
- *Self-adjoint*: If the adjoint operator $L^* : \mathcal{S}'(\mathbb{R}^d) \rightarrow \mathcal{S}'(\mathbb{R}^d)$ has the same Schwartz kernel (impulse response) as L .

Since the Schwartz kernel of a linear operator is unique [55], the property of self-adjointness will be denoted as $L = L^*$, irrespective of the actual domain and range of the operator. It is well-known that a LSI operator can always be expressed as the convolution $L\{\varphi\} = h * \varphi$, where $h = L\{\delta\} \in \mathcal{S}'(\mathbb{R}^d)$ is the impulse response of L . When L is isotropic, h is a purely radial function. Since all isotropic functions are symmetric, this also implies that an isotropic LSI operator is necessarily self-adjoint. All such operators are characterized by a Fourier symbol (a.k.a. frequency response) $\widehat{L} = \mathcal{F}\{h\}$ that is purely radial, with $\widehat{L}(\boldsymbol{\omega}) = \widehat{L}_{\text{rad}}(\|\boldsymbol{\omega}\|)$, under the implicit assumption that the radial profile \widehat{L}_{rad} is identifiable as a measurable function $\mathbb{R} \rightarrow \mathbb{R}$.

Our condition for admissibility is that L be invertible in an appropriate sense.

Definition 2 (Spline-admissible operators with trivial null space) An isotropic LSI operator L has a trivial null space if its radial frequency profile \widehat{L}_{rad} does not vanish over \mathbb{R} . We then say that it is *spline-admissible* if $1/\widehat{L}_{\text{rad}} \in L_1(\mathbb{R})$ and $\rho_{\text{rad}} = \mathcal{F}^{-1}\{1/\widehat{L}_{\text{rad}}\} \in L_1(\mathbb{R})$ where the operator $\mathcal{F}^{-1} : L_1(\mathbb{R}) \rightarrow C_0(\mathbb{R})$ is the classical inverse Fourier transform.

The typical scenario is $\widehat{L}(\boldsymbol{\omega}) = (1 + \|\boldsymbol{\omega}\|^2)^{\alpha/2}$ with $\alpha \geq 1$, which results in a stable inverse operator L^{-1} whose radially symmetric impulse response is the Bessel potential of order α . These operators play a central role in the theory of Sobolev spaces [21].

Distribution theory allows us to go further and to invert operators with nontrivial null spaces, but only if the zeros of the frequency response are located at isolated points. When the operator is isotropic, this reduces the options to the cases where $\widehat{L}(\boldsymbol{\omega})$ has a (multiple) zero at $\boldsymbol{\omega} = \mathbf{0}$. Specifically, we shall say that L is of order γ_0 if $|\widehat{L}(\boldsymbol{\omega})|/\|\boldsymbol{\omega}\|^{\gamma_0} = C_0$ as $\|\boldsymbol{\omega}\| \rightarrow 0$. The second important parameter is the asymptotic growth exponent of $\widehat{L}(\boldsymbol{\omega})$. This is the largest index γ_1 such that $|\widehat{L}(\boldsymbol{\omega})| \geq C_1\|\boldsymbol{\omega}\|^{\gamma_1}$, for all $\|\boldsymbol{\omega}\| > R$. It determines the smoothness of the Green's function of the operator.

Definition 3 (Spline-admissible operators with nontrivial null space)

An isotropic LSI operator L with radial frequency profile \widehat{L}_{rad} is said to be *spline-admissible* with a polynomial null space of degree n_0 if the following conditions are satisfied.

1. The profile \widehat{L}_{rad} does not vanish over \mathbb{R} , except for a zero of order $\gamma_0 \in (n_0, n_0 + 1]$ at the origin; that is, $|\widehat{L}_{\text{rad}}(\omega)|/|\omega|^{\gamma_0} = C_0$ as $\omega \rightarrow 0$.
2. There exists an order $\gamma_1 > 1$, a constant $C_1 > 0$, and a radius $R_1 > 0$ such that $|\widehat{L}_{\text{rad}}(\omega)| \geq C_1|\omega|^{\gamma_1}$ for all $|\omega| > R_1$ (ellipticity).
3. For all $\varphi \in \mathcal{S}(\mathbb{R}^d)$, $L^*\{\varphi\} \in L_{1,n_0}(\mathbb{R}^d)$.

The connection between Condition 1 and the null space of L will be explained in Section 2.3. Conditions 1 and 2 with $\gamma_1 > 1$ ensure that $\rho_{\text{rad}} = \mathcal{F}^{-1}\{1/\widehat{L}_{\text{rad}}\}$, which is the generalized inverse Fourier transform of the distribution $1/\widehat{L}_{\text{rad}}$, is identifiable as a continuous function $\mathbb{R} \rightarrow \mathbb{R}$. The order γ_1 actually controls the degree of differentiability (Sobolev smoothness) of ρ_{rad} . Condition 3 is a mild technical constraint on the decay of $L^*\{\varphi\}$; this constraint has not appeared to be a practical limitation so far. For instance, if L is an ordinary differential operator (an arbitrary polynomial of the Laplace operator Δ) then $L^*\{\varphi\} \in \mathcal{S}(\mathbb{R}^d)$, which is included in $L_{1,m}(\mathbb{R}^d)$ for any $m \in \mathbb{Z}$. We use this third condition for the handling of fractional operators whose impulse response decays slowly.

An attractive class of admissible operators with $\gamma_0 = \gamma_1 = \alpha$ and $n_0 = \lceil \alpha - 1 \rceil$ are the fractional Laplacians $(-\Delta)^{\frac{\alpha}{2}}$ with $\alpha \in (1, \infty)$ whose frequency response is $\|\boldsymbol{\omega}\|^\alpha$. The inverse of the fractional Laplacian of order α , which corresponds to a frequency-domain multiplication by $\|\boldsymbol{\omega}\|^{-\alpha}$, is denoted by $(-\Delta)^{-\frac{\alpha}{2}}$. Both operators are part of the same family (isotropic LSI and scale-invariant), their distributional impulse response being given by

$$k_{\alpha,d}(\mathbf{x}) = \mathcal{F}^{-1}\left\{\frac{1}{\|\boldsymbol{\omega}\|^\alpha}\right\}(\mathbf{x}) = \begin{cases} A_{\alpha,d} \|\mathbf{x}\|^{\alpha-d}, & \alpha - d, -\alpha \notin 2\mathbb{N} \\ B_{n,d} \|\mathbf{x}\|^{2n} \log(\|\mathbf{x}\|), & \alpha - d = 2n \in 2\mathbb{N} \\ (-\Delta)^n \{\delta\}, & -\alpha/2 = n \in \mathbb{N}, \end{cases} \quad (6)$$

with constants $A_{d,\alpha} = \frac{\Gamma(\frac{d-\alpha}{2})}{2^\alpha \pi^{d/2} \Gamma(\frac{\alpha}{2})}$ and $B_{d,n} = \frac{(-1)^{1+n}}{2^{2n+d-1} \pi^{d/2} \Gamma(n+\frac{d}{2}) n!}$ [19, 47].

The kernel $k_{\alpha,d}$ can also be interpreted as the Green's function of $(-\Delta)^{\frac{\alpha}{2}}$,

with the corresponding radial profile in Definition 3 being $\rho_{\text{rad}}(t) = k_{\alpha,1}(t)$. In view of (6), this means that $(-\Delta)^{\frac{\alpha}{2}}$ is admissible for $\alpha > 1$.

We note that the impulse response of the filtering operator K in Theorem 1 is proportional to $k_{-d+1,d}(\mathbf{x})$, which tells us that it decays asymptotically like $1/\|\mathbf{x}\|^{2d-1}$ when d is even, or is a power of the Laplacian (local operator) otherwise. Functionally, this means that $K(\mathcal{S}(\mathbb{R}^d)) = \mathcal{S}(\mathbb{R}^d)$ for even dimensions, and $K(\mathcal{S}(\mathbb{R}^d)) \subset L_{1,d-1-\epsilon}(\mathbb{R}^d)$ otherwise. Likewise, the impulse response of the fractional Laplacians (of non-even order) decays asymptotically like $1/\|\mathbf{x}\|^{\alpha+d}$, which implies that $(-\Delta)^{\frac{\alpha}{2}}(\mathcal{S}(\mathbb{R}^d)) \subset L_{1,\alpha-\epsilon}(\mathbb{R}^d)$ for arbitrarily small $\epsilon > 0$, so that the third condition in Definition 3 is met.

2.3 Nontrivial Null Space and Related Projectors

Let L be a LSI operator whose frequency response \widehat{L} satisfies the conditions

$$\partial^{\mathbf{k}} \widehat{L}(\mathbf{0}) = 0, \text{ for all } \mathbf{k} \in \mathbb{N}^d \text{ with } |\mathbf{k}| \leq n_0, \quad (7)$$

for some integer $n_0 \geq 0$. This flatness behavior at the origin implies that L has the capacity to annihilate all polynomials of degree n_0 by mapping them to zero (see [61, p. 131]). The explanation lies in the property that the Fourier transform of any polynomial is entirely concentrated at the origin. If, in addition, we impose that $\widehat{L}(\boldsymbol{\omega}) \neq 0$ for all $\boldsymbol{\omega} \in \mathbb{R}^d \setminus \{\mathbf{0}\}$, then we are also making sure that the null space of L is limited to polynomials.

Next, we recall that the directional derivative of a function along the direction $\boldsymbol{\xi} \in \mathbb{S}^{d-1}$ (i.e., $\boldsymbol{\xi} \in \mathbb{R}^d$ with $\|\boldsymbol{\xi}\| = 1$) is given by

$$D_{\boldsymbol{\xi}} f = \boldsymbol{\xi}^T \nabla f = \xi_1 \partial^{e_1} f + \dots + \xi_d \partial^{e_d} f. \quad (8)$$

The operator $D_{\boldsymbol{\xi}}$ is LSI with frequency response $\widehat{D}_{\boldsymbol{\xi}}(\boldsymbol{\omega}) = (i\boldsymbol{\xi}^T \boldsymbol{\omega})$. The n th iterate of $D_{\boldsymbol{\xi}}$ yields the n th derivative along $\boldsymbol{\xi}$ whose explicit expression in terms of partial derivatives is

$$D_{\boldsymbol{\xi}}^n f(\mathbf{x}) = \mathcal{F}^{-1}\{(i\boldsymbol{\xi}^T \boldsymbol{\omega})^n \widehat{f}(\boldsymbol{\omega})\}(\mathbf{x}) = \sum_{|\mathbf{k}|=n} \frac{n!}{\mathbf{k}!} \boldsymbol{\xi}^{\mathbf{k}} \partial^{\mathbf{k}} f(\mathbf{x}), \quad (9)$$

where the right-hand side follows from the application of the multinomial expansion to $(i\boldsymbol{\xi}^T \boldsymbol{\omega})^n = (\xi_1 i\omega_1 + \dots + \xi_d i\omega_d)^n$.

For isotropic operators, the directional derivatives $D_{\boldsymbol{\xi}}^n \widehat{L}(\mathbf{0})$ do not depend on the direction $\boldsymbol{\xi}$ and coincide with the radial derivatives $\widehat{L}_{\text{rad}}^{(n)}(0)$. In view of (9), (7) then has the radial equivalent

$$\widehat{L}_{\text{rad}}^{(n)}(0) = \frac{d^n \widehat{L}_{\text{rad}}(0)}{d\omega^n} = 0, \text{ for } n = 0, 1, \dots, n_0, \quad (10)$$

which is much simpler to test. It follows that an operator whose radial frequency profile is such that $|\widehat{L}_{\text{rad}}(\omega)|/|\omega|^{\gamma_0} = C_0$ as $\omega \rightarrow 0$ will annihilate all polynomials up to degree $n_0 = \lceil \gamma_0 - 1 \rceil$.

Consequently, the null space of a spline-admissible operator L of order γ_0 consists of the polynomials of degree $n_0 = (\gamma_0 - 1)$ when γ_0 is an integer and $n_0 = \lfloor \gamma_0 \rfloor$ otherwise when $\gamma_0 \notin \mathbb{N}$. We shall represent these polynomials by expanding them in the monomial/Taylor basis

$$m_{\mathbf{k}}(\mathbf{x}) = \frac{\mathbf{x}^{\mathbf{k}}}{\mathbf{k}!} \quad (11)$$

with $|\mathbf{k}| \leq n_0$. We also add a topological structure by equipping the space with the ℓ_2 norm of the Taylor coefficients, which results in the description

$$\mathcal{P}_{n_0} = \{p_0 = \sum_{|\mathbf{k}| \leq n_0} b_{\mathbf{k}} m_{\mathbf{k}} : \|p_0\|_{\mathcal{P}} < \infty\} \text{ with } \|p_0\|_{\mathcal{P}} \triangleq \|(b_{\mathbf{k}})_{|\mathbf{k}| \leq n_0}\|_2. \quad (12)$$

To avoid a notational overload, we shall often denote this null space by \mathcal{P} , with the convention that $\mathcal{P} = \mathcal{P}_{n_0} = \{0\}$ when $n_0 = \lceil \gamma_0 - 1 \rceil < 0$ (for the operators L whose null space is trivial). The important point here is that (12) specifies a finite-dimensional Banach subspace of $\mathcal{S}'(\mathbb{R}^d)$. Its continuous dual \mathcal{P}' is finite-dimensional as well, although it is composed of “abstract” elements $p_0^* \in \mathcal{P}'$ that are, in fact, equivalence classes in $\mathcal{S}'(\mathbb{R}^d)$. Yet, it is possible to identify every dual element $p_0^* \in \mathcal{P}'$ as a true function by selecting a particular dual basis $\{m_{\mathbf{k}}^*\}_{|\mathbf{k}| \leq n_0}$ such that $\langle m_{\mathbf{k}}^*, m_{\mathbf{k}'} \rangle = \delta_{\mathbf{k}-\mathbf{k}'}$ (Kroneker delta). Our specific choice is

$$m_{\mathbf{k}}^* = (-1)^{|\mathbf{k}|} \partial^{\mathbf{k}} \kappa_{\text{iso}} \in \mathcal{S}(\mathbb{R}^d) \quad (13)$$

with $\mathbf{k} \in \mathbb{N}^d$, where κ_{iso} is the isotropic function described in Lemma 1.

Lemma 1 (adapted from [57]) *There exists an isotropic window $\kappa_{\text{iso}} \in \mathcal{S}(\mathbb{R}^d)$ such that*

$$\langle m_{\mathbf{k}}, (-1)^{|\mathbf{n}|} \partial^{\mathbf{n}} \kappa_{\text{iso}} \rangle = \delta_{\mathbf{k}-\mathbf{n}} \quad (14)$$

for all $\mathbf{k}, \mathbf{n} \in \mathbb{N}^d$, subject to the spectral constraints $\widehat{\kappa}_{\text{iso}}(\boldsymbol{\omega}) = 1$ for $\|\boldsymbol{\omega}\| < \frac{1}{2}$, $1 \geq \widehat{\kappa}_{\text{iso}}(\boldsymbol{\omega}) \geq 0$ for $\frac{1}{2} < \|\boldsymbol{\omega}\| < 1$, and $\widehat{\kappa}_{\text{iso}}(\boldsymbol{\omega}) = 0$ for $\|\boldsymbol{\omega}\| \geq 1$.

This allows us to describe the dual space explicitly as

$$\mathcal{P}' = \mathcal{P}'_{n_0} = \{p_0^* = \sum_{|\mathbf{k}| \leq n_0} b_{\mathbf{k}}^* m_{\mathbf{k}}^* : \|p_0^*\|_{\mathcal{P}'} < \infty\} \text{ with } \|p_0^*\|_{\mathcal{P}'} \triangleq \|(b_{\mathbf{k}}^*)\|_2 \quad (15)$$

where each elements p_0^* has a unique representation in terms of its coefficients $(b_{\mathbf{k}}^*)_{|\mathbf{k}| \leq n_0}$. We use the dual basis $\{m_{\mathbf{k}}^*\}$ to specify the projection operator $\text{Proj}_{\mathcal{P}} : \mathcal{S}'(\mathbb{R}^d) \rightarrow \mathcal{P}_{n_0}$ as

$$\text{Proj}_{\mathcal{P}}\{f\} = \sum_{|\mathbf{k}| \leq n_0} \langle f, m_{\mathbf{k}}^* \rangle m_{\mathbf{k}}, \quad (16)$$

which is well-defined for any $f \in \mathcal{S}'(\mathbb{R}^d)$ since $m_{\mathbf{k}}^* \in \mathcal{S}(\mathbb{R}^d)$. The “transpose” of this operator is

$$\text{Proj}_{\mathcal{P}'}\{\nu\} = \sum_{|\mathbf{k}| \leq n_0} \langle m_{\mathbf{k}}, \nu \rangle m_{\mathbf{k}}^*, \quad (17)$$

which returns the projection of ν onto $\mathcal{P}'_{n_0} \subseteq \mathcal{S}(\mathbb{R}^d)$ under the implicit assumption that ν has sufficient decay for $\nu \mapsto \langle m_{\mathbf{k}}, \nu \rangle$ to be well-defined—for instance, $\nu \in L_{1,n_0}(\mathbb{R}^d)$. Correspondingly, we also have that $\text{Proj}_{\mathcal{P}'}\{L^*\varphi\} = 0$ for all φ such that $L^*\{\varphi\} \in L_{1,n_0}(\mathbb{R}^d)$ since $\langle m_{\mathbf{k}}, L^*\varphi \rangle = \langle Lm_{\mathbf{k}}, \varphi \rangle = 0$ for $|\mathbf{k}| \leq n_0$. The latter manipulation of the duality product is legitimate in reason of the inclusion $\mathcal{P}_{n_0} \subset L_{\infty,-n_0}(\mathbb{R}^d) = (L_{1,n_0}(\mathbb{R}^d))'$.

Even though the null space of an admissible operator L may be non-trivial, its intersection with $\mathcal{S}(\mathbb{R}^d)$ is always $\{0\}$. This implies that $L^* = L$ is injective on $\mathcal{S}(\mathbb{R}^d)$ with $L^{*-1}L^*\{\varphi\} = \varphi$ for all $\varphi \in \mathcal{S}(\mathbb{R}^d)$ where $L^{*-1} = L^{-1}$ is the LSI operator whose frequency response is $1/|\widehat{L}|$.

3 Radon Transform

The Radon transform extracts the integrals of a function on \mathbb{R}^d over all hyperplanes of dimension $(d-1)$. These hyperplanes are indexed over $\mathbb{R} \times \mathbb{S}^{d-1}$, where $\mathbb{S}^{d-1} = \{\boldsymbol{\xi} \in \mathbb{R}^d : \|\boldsymbol{\xi}\|_2 = 1\}$ is the unit sphere in \mathbb{R}^d . The coordinates of a hyperplane associated with an offset $t \in \mathbb{R}$ and a normal vector $\boldsymbol{\xi} \in \mathbb{S}^{d-1}$ satisfy

$$\boldsymbol{\xi}^\top \mathbf{x} = \xi_1 x_1 + \cdots + \xi_d x_d = t.$$

Here, we first review the classical theory of the Radon transform [27], starting with the case of test functions (Section 3.1), and extending it by duality to tempered distributions (Section 3.2). Then, in Section 3.3, we specify the Radon transform and its inverse on an appropriate class of intermediate Banach spaces \mathcal{Y} with $\mathcal{S}(\mathbb{R}^d) \xrightarrow{\text{d}} \mathcal{Y} \xrightarrow{\text{d}} \mathcal{S}'(\mathbb{R}^d)$ (Theorem 3). Finally, in Section 3.4, we provide the (filtered) Radon transforms of the dictionary elements— isotropic kernels and ridges—that are relevant to our investigation.

3.1 Classical Integral Formulation

The Radon transform of the function $f \in L_1(\mathbb{R}^d)$ is defined as

$$\mathbb{R}\{f\}(t, \boldsymbol{\xi}) = \int_{\mathbb{R}^d} \delta(t - \boldsymbol{\xi}^\top \mathbf{x}) f(\mathbf{x}) d\mathbf{x}, \quad (t, \boldsymbol{\xi}) \in \mathbb{R} \times \mathbb{S}^{d-1}. \quad (18)$$

The adjoint of \mathbb{R} is the backprojection operator \mathbb{R}^* . Its action on $g : \mathbb{R} \times \mathbb{S}^{d-1} \rightarrow \mathbb{R}$ yields the function

$$\mathbb{R}^*\{g\}(\mathbf{x}) = \int_{\mathbb{S}^{d-1}} \underbrace{g(\boldsymbol{\xi}^\top \mathbf{x}, \boldsymbol{\xi})}_t d\boldsymbol{\xi}, \quad \mathbf{x} \in \mathbb{R}^d. \quad (19)$$

Given the d -dimensional Fourier transform \hat{f} of $f \in L_1(\mathbb{R}^d)$, one can calculate $R\{f\}(\cdot, \boldsymbol{\xi}_0)$ for any fixed $\boldsymbol{\xi}_0 \in \mathbb{S}^{d-1}$ through the relation

$$R\{f\}(t, \boldsymbol{\xi}_0) = \frac{1}{2\pi} \int_{\mathbb{R}} \hat{f}(\omega \boldsymbol{\xi}_0) e^{i\omega t} d\omega = \mathcal{F}_{1D}^{-1}\{\hat{f}(\cdot \boldsymbol{\xi}_0)\}(t), \quad (20)$$

In other words, the restriction of $\hat{f} : \mathbb{R}^d \rightarrow \mathbb{C}$ along the ray $\{\boldsymbol{\omega} = \omega \boldsymbol{\xi}_0 : \omega \in \mathbb{R}\}$ coincides with the 1D Fourier transform of $R\{f\}(\cdot, \boldsymbol{\xi}_0)$, a property that is referred to as the *Fourier-slice theorem*.

To describe the functional properties of the Radon transform, one needs the (hyper)spherical (or Radon-domain) counterparts of the spaces described in Section 2.1. There, the Euclidean indexing with $\boldsymbol{x} \in \mathbb{R}^d$ must be replaced by $(t, \boldsymbol{\xi}) \in \mathbb{R} \times \mathbb{S}^{d-1}$. The spherical counterpart of $\mathcal{S}(\mathbb{R}^d)$ is $\mathcal{S}(\mathbb{R} \times \mathbb{S}^{d-1})$. Correspondingly, an element $g \in \mathcal{S}'(\mathbb{R} \times \mathbb{S}^{d-1})$ is a continuous linear functional on $\mathcal{S}(\mathbb{R} \times \mathbb{S}^{d-1})$ whose action on the test function ϕ is represented by the duality product $g : \phi \mapsto \langle g, \phi \rangle_{\text{Rad}}$. When g can be identified as an ordinary function $g : (t, \boldsymbol{\xi}) \mapsto g(t, \boldsymbol{\xi}) \in \mathbb{R}$, one can write that

$$\langle g, \phi \rangle_{\text{Rad}} = \int_{\mathbb{S}^{d-1}} \int_{\mathbb{R}} g(t, \boldsymbol{\xi}) \phi(t, \boldsymbol{\xi}) dt d\boldsymbol{\xi} \quad (21)$$

where $d\boldsymbol{\xi}$ stands for the surface element on the unit sphere \mathbb{S}^{d-1} .

The key property for analysis is that the Radon transform is continuous on $\mathcal{S}(\mathbb{R}^d)$ and invertible [27, 23, 43]. In addition to a backprojection, the inversion involves the so-called “filtering” operator.

Definition 4 The filtering operator $K : \mathcal{S}(\mathbb{R}^d) \rightarrow \mathcal{S}'(\mathbb{R}^d)$ is defined as

$$K\{\varphi\} = \mathcal{F}^{-1}\{\widehat{K}\widehat{\varphi}\} \quad \text{with} \quad \widehat{K}(\boldsymbol{\omega}) = c_d \|\boldsymbol{\omega}\|^{d-1}, \quad (22)$$

where $c_d = \frac{1}{2(2\pi)^{d-1}}$.

The filtering operator is isotropic LSI and, as such, has a Radon-domain counterpart (see Definition 5) denoted by K_{rad} that exclusively acts along the radial variable.

Definition 5 (Radon-domain counterpart of an isotropic LSI operator) Let $L : \mathcal{S}(\mathbb{R}^d) \rightarrow \mathcal{S}'(\mathbb{R}^d)$ be an isotropic LSI operator with radial frequency profile $\widehat{L}_{\text{rad}} : \mathbb{R} \rightarrow \mathbb{R}$. Then, the Radon-domain counterpart of L is $L_{\text{rad}} : \mathcal{S}(\mathbb{R} \times \mathbb{S}^{d-1}) \rightarrow \mathcal{S}'(\mathbb{R} \times \mathbb{S}^{d-1})$, which is defined as

$$L_{\text{rad}}\{\phi(\cdot, \boldsymbol{\xi})\}(t) = \mathcal{F}_{1D}^{-1}\{\widehat{L}_{\text{rad}}\widehat{\phi}(\cdot, \boldsymbol{\xi})\}(t), \quad (23)$$

where $\widehat{\phi}(\omega, \boldsymbol{\xi}) = \int_{\mathbb{R}} \phi(t, \boldsymbol{\xi}) e^{-i\omega t} dt$ is the 1D Fourier transform of $t \mapsto \phi(t, \boldsymbol{\xi})$.

Theorem 1 (Continuity and invertibility of the Radon transform on $\mathcal{S}(\mathbb{R}^d)$) *The Radon operator R continuously maps $\mathcal{S}(\mathbb{R}^d) \rightarrow \mathcal{S}(\mathbb{R} \times \mathbb{S}^{d-1})$. Moreover, $KR^*R = R^*K_{\text{rad}}R = R^*RK = \text{Id}$ on $\mathcal{S}(\mathbb{R}^d)$.*

Based on this result, we can identify the filtering operator as $K = (R^*R)^{-1} = c_d(-\Delta)^{\frac{d-1}{2}}$ (fractional Laplacian). Alternatively, one can perform the filtering in the Radon domain by means of the operator K_{rad} , which implements a 1D convolution along the radial variable. The connection is that the frequency response of K_{rad} coincides with the radial frequency profile of K so that $\widehat{K}(\boldsymbol{\omega}) = \widehat{K}_{\text{rad}}(\|\boldsymbol{\omega}\|)$ with $\widehat{K}_{\text{rad}}(\omega) = c_d|\omega|^{d-1}$.

While the Radon transform $R : \mathcal{S}(\mathbb{R}^d) \rightarrow \mathcal{S}(\mathbb{R} \times \mathbb{S}^{d-1})$ is invertible, it is not surjective, which means that not every hyper-spherical test function $\phi \in \mathcal{S}(\mathbb{R} \times \mathbb{S}^{d-1})$ can be written as $\phi = R\{\varphi\}$ with $\varphi \in \mathcal{S}(\mathbb{R}^d)$. A necessary condition is that ϕ be even, but this is not sufficient [20, 27, 23]. The good news, however, is that the range of R on $\mathcal{S}(\mathbb{R}^d)$ is a closed subspace of $\mathcal{S}(\mathbb{R} \times \mathbb{S}^{d-1})$ [23, p. 60]. Accordingly, one can identify the range space $\mathcal{S}_{\text{Rad}} \triangleq R(\mathcal{S}(\mathbb{R}^d))$ equipped with the Fréchet topology of $\mathcal{S}(\mathbb{R} \times \mathbb{S}^{d-1})$. Since the domain and range spaces are both Fréchet, we then invoke the open-mapping theorem [45, Theorem 2.11] to deduce that the transform $\varphi \mapsto R\{\varphi\}$ is a homeomorphism of $\mathcal{S}(\mathbb{R}^d)$ onto \mathcal{S}_{Rad} .

Corollary 1 *The operator $R : \mathcal{S}(\mathbb{R}^d) \rightarrow \mathcal{S}_{\text{Rad}}$ is a continuous bijection, with a continuous inverse given by $R^{-1} = (R^*K_{\text{rad}}) : \mathcal{S}_{\text{Rad}} \rightarrow \mathcal{S}(\mathbb{R}^d)$.*

3.2 Distributional Extension

To extend the framework to distributions, one proceeds by duality. By invoking the property that $R^*K_{\text{rad}}R = \text{Id}$ on $\mathcal{S}(\mathbb{R}^d)$, we make the manipulation

$$\begin{aligned} \forall \varphi \in \mathcal{S}(\mathbb{R}^d) \quad \langle f, \varphi \rangle &= \langle f, R^*K_{\text{rad}}R\{\varphi\} \rangle \\ &= \langle R\{f\}, K_{\text{rad}}R\{\varphi\} \rangle_{\text{Rad}} = \langle R\{f\}, \phi \rangle_{\text{Rad}}, \end{aligned} \quad (24)$$

with $\phi = K_{\text{rad}}R\{\varphi\} \in K_{\text{rad}}R(\mathcal{S}(\mathbb{R}^d))$ and $\varphi = R^*\{\phi\}$. Relation (24), which is valid in the classical sense for $f \in L_1(\mathbb{R}^d)$, is then used as definition to extend the scope of R for $f \in \mathcal{S}'(\mathbb{R}^d)$.

Definition 6 The distribution $g = R\{f\} \in (K_{\text{rad}}R(\mathcal{S}(\mathbb{R}^d)))'$ is the (formal) Radon transform of $f \in \mathcal{S}'(\mathbb{R}^d)$ if

$$\forall \phi \in K_{\text{rad}}R(\mathcal{S}(\mathbb{R}^d)) : \quad \langle g, \phi \rangle_{\text{Rad}} = \langle f, R^*\{\phi\} \rangle. \quad (25)$$

Likewise, $\tilde{g} = K_{\text{rad}}R\{f\} \in \mathcal{S}'_{\text{Rad}}$ is the (formal) filtered projection of $f \in \mathcal{S}'(\mathbb{R}^d)$ if

$$\forall \phi \in \mathcal{S}_{\text{Rad}} : \quad \langle \tilde{g}, \phi \rangle_{\text{Rad}} = \langle f, R^*K_{\text{rad}}\{\phi\} \rangle. \quad (26)$$

Finally, $f = R^*\{g\} \in \mathcal{S}'(\mathbb{R}^d)$ is the backprojection of $g \in \mathcal{S}'(\mathbb{R} \times \mathbb{S}^{d-1})$ if

$$\forall \varphi \in \mathcal{S}(\mathbb{R}^d) : \quad \langle R^*\{g\}, \varphi \rangle = \langle g, R\{\varphi\} \rangle_{\text{Rad}}. \quad (27)$$

While (27) identifies $\mathbf{R}^*\{g\}$ as a single, unique distribution in $\mathcal{S}'(\mathbb{R}^d)$, this is not so for (26) (resp., (25)), as the members of $\mathcal{S}'_{\text{Rad}}$ (resp., of $(\mathbf{K}_{\text{rad}}\mathbf{R}(\mathcal{S}))'$) are equivalence classes in $\mathcal{S}'(\mathbb{R} \times \mathbb{S}^{d-1})$. To make this explicit, we take advantage of the equivalence $\mathbf{R}^*\{g\} = 0 \Leftrightarrow \langle g, \phi \rangle_{\text{Rad}} = 0$ to identify the null space of the backprojection operator as being

$$\mathcal{N}_{\mathbf{R}^*} = \{g \in \mathcal{S}'(\mathbb{R} \times \mathbb{S}^{d-1}) : \langle g, \phi \rangle_{\text{Rad}} = 0, \forall \phi \in \mathcal{S}_{\text{Rad}}\}, \quad (28)$$

which is a closed subspace of $\mathcal{S}'(\mathbb{R} \times \mathbb{S}^{d-1})$. It is then possible to describe $\mathcal{S}'_{\text{Rad}}$ as the abstract quotient space $\mathcal{S}'(\mathbb{R} \times \mathbb{S}^{d-1})/\mathcal{N}_{\mathbf{R}^*}$. In other words, if we find a hyper-spherical distribution $g_0 \in \mathcal{S}'(\mathbb{R} \times \mathbb{S}^{d-1})$ such that (26) is met for a given $f \in \mathcal{S}'(\mathbb{R}^d)$, then, strictly speaking, $\mathbf{K}_{\text{rad}}\mathbf{R}\{f\} \in \mathcal{S}'_{\text{Rad}}$ is the equivalence class (or coset) given by

$$\mathbf{K}_{\text{rad}}\mathbf{R}\{f\} = [g_0] = \{g_0 + h : h \in \mathcal{N}_{\mathbf{R}^*}\}. \quad (29)$$

Since $[g_0] = [g]$ for any $g \in \mathbf{K}_{\text{rad}}\mathbf{R}\{f\}$, we refer to the members of $\mathbf{K}_{\text{rad}}\mathbf{R}\{f\}$ as “formal” filtered projections of f to remind us of this lack of unicity.

Based on those definitions, one obtains the classical result on the invertibility of the (filtered) Radon transform on $\mathcal{S}'(\mathbb{R}^d)$ [27], which is the dual of Corollary 1.

Theorem 2 (Invertibility of the Radon transform on $\mathcal{S}'(\mathbb{R}^d)$) *It holds that $\mathbf{R}^*\mathbf{K}_{\text{rad}}\mathbf{R} = \text{Id}$ on $\mathcal{S}'(\mathbb{R}^d)$. More precisely, the filtered Radon transform $\mathbf{K}_{\text{rad}}\mathbf{R} : \mathcal{S}'(\mathbb{R}^d) \rightarrow \mathcal{S}'_{\text{Rad}}$ is a continuous bijection, with a continuous inverse given by $(\mathbf{K}_{\text{rad}}\mathbf{R})^{-1} = \mathbf{R}^* : \mathcal{S}'_{\text{Rad}} \rightarrow \mathcal{S}'(\mathbb{R}^d)$.*

To illustrate the fact that (26) does not identify a single distribution, we consider the Dirac ridge $\delta(\boldsymbol{\xi}_0^{\text{T}}\mathbf{x} - t_0) \in \mathcal{S}'(\mathbb{R}^d)$ and refer to the definition (18) of the Radon transform to deduce that, for all $\phi = \mathbf{R}\{\varphi\} \in \mathcal{S}_{\text{Rad}}$ with $\varphi \in \mathcal{S}(\mathbb{R}^d)$,

$$\begin{aligned} \langle \delta(\boldsymbol{\xi}_0^{\text{T}} \cdot - t_0), \mathbf{R}^*\mathbf{K}_{\text{rad}}\{\phi\} \rangle &= \langle \delta(\boldsymbol{\xi}_0^{\text{T}} \cdot - t_0), \overbrace{\mathbf{R}^*\mathbf{K}_{\text{rad}}\mathbf{R}}^{\text{Id}}\{\varphi\} \rangle \\ &= \int_{\mathbb{R}^d} \delta(\boldsymbol{\xi}_0^{\text{T}}\mathbf{x} - t_0)\varphi(\mathbf{x})d\mathbf{x} = \mathbf{R}\{\varphi\}(-t_0, -\boldsymbol{\xi}_0) \\ &= \langle \delta(\cdot + (t_0, \boldsymbol{\xi}_0)), \mathbf{R}\{\varphi\} \rangle_{\text{Rad}} = \langle \delta(\cdot + (t_0, \boldsymbol{\xi}_0)), \phi \rangle_{\text{Rad}}, \end{aligned}$$

which shows that $\delta(\cdot + \mathbf{z}_0)$ with $\mathbf{z}_0 = (t_0, \boldsymbol{\xi}_0)$ is a formal filtered projection of $\delta(\boldsymbol{\xi}_0^{\text{T}}\mathbf{x} - t_0)$. Moreover, since $\delta(\boldsymbol{\xi}_0^{\text{T}}\mathbf{x} - t_0) = \delta(-\boldsymbol{\xi}_0^{\text{T}}\mathbf{x} + t_0)$, the same holds true for $\delta(\cdot - \mathbf{z}_0)$, as well as for $\delta_{\text{Rad}, \mathbf{z}_0} \triangleq \frac{1}{2}(\delta(\cdot - \mathbf{z}_0) + \delta(\cdot + \mathbf{z}_0))$, which has the advantage of being symmetric. While the general solution in $\mathcal{S}'_{\text{Rad}}$ is $\mathbf{K}_{\text{rad}}\mathbf{R}\{\delta(\boldsymbol{\xi}_0^{\text{T}} \cdot - t_0)\} = [\delta(\cdot \pm \mathbf{z}_0)]$, we shall see that there also exists a way to specify a representer that is unique (i.e., $\delta_{\text{Rad}, \mathbf{z}_0}$) by restricting the range of $\mathbf{K}_{\text{rad}}\mathbf{R}$ to a suitable subspace of measures.

The distributional extension of the Radon transform inherits most of the properties of the “classical” operator defined in (18). Of special relevance to

us is the quasi-commutativity of \mathbb{R} with convolution, also known as the *intertwining property*. Specifically, let $h, f \in \mathcal{S}'(\mathbb{R}^d)$ be two distributions whose convolution $h * f$ is well-defined in $\mathcal{S}'(\mathbb{R}^d)$. Then,

$$\mathbb{R}\{h * f\} = \mathbb{R}\{h\} \otimes \mathbb{R}\{f\} \quad (30)$$

where the symbol “ \otimes ” denotes the 1D convolution along the radial variable $t \in \mathbb{R}$ with $(u \otimes g)(t, \boldsymbol{\xi}) = \langle u(\cdot, \boldsymbol{\xi}), g(t - \cdot, \boldsymbol{\xi}) \rangle$. In particular, when $h = \mathbb{L}\{\delta\}$ is the (isotropic) impulse response of an LSI operator whose frequency response $\widehat{L}(\boldsymbol{\omega}) = \widehat{L}_{\text{rad}}(\|\boldsymbol{\omega}\|)$ is purely radial, we get that

$$\mathbb{R}\{h * f\} = \mathbb{R}\mathbb{L}\{f\} = \mathbb{L}_{\text{rad}}\mathbb{R}\{f\}, \quad (31)$$

where \mathbb{L}_{rad} is the corresponding Radon-domain operator of Definition 5. Likewise, by duality, for $g \in \mathcal{S}'(\mathbb{R} \times \mathbb{S}^{d-1})$ we have that

$$\mathbb{L}\mathbb{R}^*\{g\} = \mathbb{R}^*\mathbb{L}_{\text{rad}}\{g\}, \quad (32)$$

under the implicit assumption that $\mathbb{L}\{\mathbb{R}^*g\}$ and $\mathbb{L}_{\text{rad}}\{g\}$ are well-defined distributions. By taking inspiration from Theorem 1, we can then use these relations for $\mathbb{L} = \mathbb{K} = (\mathbb{R}^*\mathbb{R})^{-1}$ to show that $\mathbb{R}^*\mathbb{K}_{\text{rad}}\mathbb{R}\{f\} = \mathbb{R}^*\mathbb{R}\mathbb{K}\{f\} = \mathbb{K}\mathbb{R}^*\mathbb{R}\{f\} = f$ for a broad class of distributions. The first form is valid for all $f \in \mathcal{S}'(\mathbb{R}^d)$ (Theorem 2), but there is a slight restriction with the second form (resp., third form), which requires that $\mathbb{K}\{f\}$ (resp., $\mathbb{K}\{g\}$ with $g = \mathbb{R}^*\mathbb{R}\{f\} \in \mathcal{S}'(\mathbb{R}^d)$) be well-defined in $\mathcal{S}'(\mathbb{R}^d)$. While the latter condition is always met when d is odd, it may fail² in even dimensions with distributions (e.g., polynomials) whose Fourier transform is singular at the origin. The good news for our regularization framework is that these problematic distributions are either excluded from the native space or annihilated by \mathbb{L} , so that it is legitimate to write that $\mathbb{L}_{\mathbb{R}} = \mathbb{K}_{\text{rad}}\mathbb{R}\mathbb{L} = \mathbb{R}\mathbb{K}\mathbb{L}$, where the second form has the advantage that \mathbb{K} and \mathbb{L} can be pooled into a single augmented operator ($\mathbb{K}\mathbb{L}$). An alternative form is $\mathbb{L}_{\mathbb{R}} = \mathbb{Q}_{\text{rad}}\mathbb{R}$, where $\mathbb{Q}_{\text{rad}} = \mathbb{K}_{\text{rad}}\mathbb{L}_{\text{rad}}$ is the radial Radon-domain operator whose frequency response is $\widehat{\mathbb{Q}}_{\text{rad}}(\boldsymbol{\omega}) = c_d|\boldsymbol{\omega}|^{d-1}\widehat{L}_{\text{rad}}(\boldsymbol{\omega})$.

3.3 Radon-Compatible Banach Spaces

Our formulation requires the identification of Radon-domain Banach spaces over which the backprojection operator \mathbb{R}^* is invertible. This is a non-trivial point because the extended operator $\mathbb{R}^* : \mathcal{S}'(\mathbb{R} \times \mathbb{S}^{d-1}) \rightarrow \mathcal{S}'(\mathbb{R}^d)$ in Definition 6 is not injective. In fact, it has the highly non-trivial null space $\ker(\mathbb{R}^*) = \mathcal{S}_{\text{Rad}}^{\perp}$, which is a superset of the odd Radon-domain distributions [20]. Yet, \mathbb{R}^* is invertible on $\mathcal{S}'_{\text{Rad}}$ and surjective on $\mathcal{S}'(\mathbb{R}^d)$ (Theorem 2).

To ensure invertibility, we therefore need to restrict ourselves to Banach spaces that are embedded in $\mathcal{S}'_{\text{Rad}}$. To identify such objects, we consider a

² For $d = 2n$ even, $\widehat{K}(\boldsymbol{\omega}) \propto \|\boldsymbol{\omega}\|^{2n-1}$ which is C^∞ everywhere, except at the origin where it is only C^{2n-2} . This means that \mathbb{K} can properly handle (and annihilate) polynomials only up to degree $(2n - 2)$.

generic Banach space $\mathcal{X} = (\mathcal{X}, \|\cdot\|_{\mathcal{X}})$ such that $\mathcal{S}(\mathbb{R} \times \mathbb{S}^{d-1}) \xhookrightarrow{d} \mathcal{X} \xhookrightarrow{d} \mathcal{S}'(\mathbb{R} \times \mathbb{S}^{d-1})$. This dense-embedding hypothesis has several implications:

1. The space \mathcal{X} is the completion of $\mathcal{S}(\mathbb{R} \times \mathbb{S}^{d-1})$ in the $\|\cdot\|_{\mathcal{X}}$ norm; i.e.,

$$\mathcal{X} = \overline{(\mathcal{S}(\mathbb{R} \times \mathbb{S}^{d-1}), \|\cdot\|_{\mathcal{X}})}. \quad (33)$$

2. The dual space $\mathcal{X}' \hookrightarrow \mathcal{S}'(\mathbb{R} \times \mathbb{S}^{d-1})$ is equipped with the norm

$$\|g\|_{\mathcal{X}'} = \sup_{\phi \in \mathcal{X}: \|\phi\|_{\mathcal{X}} \leq 1} \langle g, \phi \rangle = \sup_{\phi \in \mathcal{S}(\mathbb{R} \times \mathbb{S}^{d-1}): \|\phi\|_{\mathcal{X}} \leq 1} \langle g, \phi \rangle, \quad (34)$$

where the restriction of $\phi \in \mathcal{S}(\mathbb{R} \times \mathbb{S}^{d-1})$ on the right-hand side of (34) is justified by the denseness of $\mathcal{S}(\mathbb{R} \times \mathbb{S}^{d-1})$ in \mathcal{X} .

3. The definition of $\|g\|_{\mathcal{X}'}$ given by the right-hand side of (34) is valid for any distribution $g \in \mathcal{S}'(\mathbb{R} \times \mathbb{S}^{d-1})$ with $\|g\|_{\mathcal{X}'} = \infty$ for $g \notin \mathcal{X}'$. Accordingly, we can specify the topological dual of \mathcal{X} as

$$\mathcal{X}' = \{g \in \mathcal{S}'(\mathbb{R} \times \mathbb{S}^{d-1}) : \|g\|_{\mathcal{X}'} < \infty\}. \quad (35)$$

Likewise, based on the pair $(\mathcal{S}_{\text{Rad}}, \mathcal{S}'_{\text{Rad}})$, we specify the Radon-compatible Banach subspaces

$$\mathcal{X}_{\text{Rad}} = \overline{(\mathcal{S}_{\text{Rad}}, \|\cdot\|_{\mathcal{X}})} \quad (36)$$

$$\mathcal{X}'_{\text{Rad}} = (\mathcal{X}_{\text{Rad}})' = \{g \in \mathcal{S}'_{\text{Rad}} : \|g\|_{\mathcal{X}'_{\text{Rad}}} < \infty\} \quad (37)$$

where the underlying dual norms have a definition that is analogous to (34) with \mathcal{S}_{Rad} and \mathcal{X}_{Rad} substituting for $\mathcal{S}(\mathbb{R} \times \mathbb{S}^{d-1})$ and \mathcal{X} .

Theorem 3 (adapted from [57]) *Let $(\mathcal{X}_{\text{Rad}}, \mathcal{X}'_{\text{Rad}})$ be the dual pair of spaces specified by (36) and (37). Then,*

1. the map $\mathbb{R}^* \mathbb{K}_{\text{rad}} : \mathcal{X}_{\text{Rad}} \rightarrow \mathcal{Y} = \mathbb{R}^* \mathbb{K}_{\text{rad}}(\mathcal{X}_{\text{Rad}})$ is an isometric bijection, with $\mathbb{R} \mathbb{R}^* \mathbb{K}_{\text{rad}} = \text{Id}$ on \mathcal{X}_{Rad} ;
2. the map $\mathbb{R}^* : \mathcal{X}'_{\text{Rad}} \rightarrow \mathcal{Y}' = \mathbb{R}^*(\mathcal{X}'_{\text{Rad}})$ is an isometric bijection, with $\mathbb{K}_{\text{rad}} \mathbb{R} \mathbb{R}^* = \text{Id}$ on $\mathcal{X}'_{\text{Rad}}$.

Moreover, if there exists a complementary Banach space $\mathcal{X}_{\text{Rad}}^c$ such that $\mathcal{X} = \mathcal{X}_{\text{Rad}} \oplus \mathcal{X}_{\text{Rad}}^c$, then $\mathcal{X}' = \mathcal{X}'_{\text{Rad}} \oplus (\mathcal{X}_{\text{Rad}}^c)'$ where $(\mathcal{X}_{\text{Rad}}^c)'$ can be identified as the null space of the backprojection operator $\mathbb{R}^* : \mathcal{X}' \rightarrow \mathcal{Y}' \hookrightarrow \mathcal{S}(\mathbb{R}^d)$.

The prototypical examples where those properties are met are $(\mathcal{X}, \mathcal{X}') = (L_p(\mathbb{R} \times \mathbb{S}^{d-1}), L_q(\mathbb{R} \times \mathbb{S}^{d-1}))$ with $p \in [1, \infty)$ and $q = p/(p-1)$ (conjugate exponent), as well as $(\mathcal{X}, \mathcal{X}') = (C_0(\mathbb{R} \times \mathbb{S}^{d-1}), \mathcal{M}(\mathbb{R} \times \mathbb{S}^{d-1}))$. In fact, those hyper-spherical spaces have the convenient feature of admitting a decomposition in their even and odd components.

Lemma 2 *Let $\mathcal{Z} = \mathbb{R} \times \mathbb{S}^{d-1}$. Then, for $\mathcal{X} = L_p(\mathcal{Z})$ with $p \in [1, \infty)$ and $\mathcal{X} = C_0(\mathcal{Z})$ for $p = \infty$, we have that $\mathcal{X} = \mathcal{X}_{\text{Rad}} \oplus \mathcal{X}_{\text{Rad}}^c$ where*

$$\mathcal{X}_{\text{Rad}} = \mathcal{X}_{\text{even}} = \{g \in \mathcal{X} : g(\mathbf{z}) = g(-\mathbf{z}), \forall \mathbf{z} \in \mathcal{Z}\} \quad (38)$$

$$\mathcal{X}_{\text{Rad}}^c = \mathcal{X}_{\text{odd}} = \{g \in \mathcal{X} : g(\mathbf{z}) = -g(-\mathbf{z}), \forall \mathbf{z} \in \mathcal{Z}\}. \quad (39)$$

Proof To establish this result directly is tricky because the characterization of \mathcal{S}_{Rad} involves some general moment conditions [20, 23, 27]. Instead, we consider the smaller space of even Radon-domain Lizorkin test functions [26] described by

$$\mathcal{S}_{\text{Liz,Rad}} = \{\phi \in \mathcal{S}_{\text{even}}(\mathcal{Z}) : \int_{\mathbb{R}} t^k \phi(t, \boldsymbol{\xi}) dt = 0, \forall \boldsymbol{\xi} \in \mathbb{S}^{d-1}, k \in \mathbb{N}\}, \quad (40)$$

which is such that $\mathcal{S}_{\text{Liz,Rad}} \subset \mathcal{S}_{\text{Rad}} \subset \mathcal{S}_{\text{even}}(\mathcal{Z})$. We then invoke a general result by Samko [46] that implies that $\overline{(\mathcal{S}_{\text{Liz,Rad}}, \|\cdot\|_{L_p})} = L_{p,\text{even}}(\mathcal{Z}) \supset (\mathcal{S}_{\text{Rad}}, \|\cdot\|_{L_p})$ for $p \in [1, \infty)$ and $\overline{(\mathcal{S}_{\text{Liz,Rad}}, \|\cdot\|_{L_\infty})} = C_{0,\text{even}}(\mathcal{Z})$ otherwise [34]. The claim then follows from the observation that $L_p(\mathcal{Z}) = L_{p,\text{even}}(\mathcal{Z}) \oplus L_{p,\text{odd}}(\mathcal{Z})$ with $L_{p,\text{even}}(\mathcal{Z}) = \overline{(\mathcal{S}_{\text{Rad}}, \|\cdot\|_{L_p})}$ (because the completion is unique) and suitable adaptation for $p = \infty$. \square

Correspondingly, we get that $\mathcal{X}'_{\text{Rad}} = \text{P}_{\text{even}}(\mathcal{X}') = \mathcal{X}'_{\text{even}}$ and $(\mathcal{X}'_{\text{Rad}})^c = (\text{Id} - \text{P}_{\text{even}})(\mathcal{X}') = \mathcal{X}'_{\text{odd}}$, with the cases of greatest interest to us being $\mathcal{M}_{\text{Rad}} = \mathcal{M}_{\text{even}}(\mathbb{R} \times \mathbb{S}^{d-1})$ and $L_{2,\text{Rad}} = L_{2,\text{even}}(\mathbb{R} \times \mathbb{S}^{d-1})$.

3.4 Specific Radon Transforms

The Fourier-slice theorem expressed by (20) remains valid for tempered distributions [43] and therefore also yields a characterization of $\text{R}\{f\}$ that is compatible with the Banach framework of Theorem 3. It is especially helpful when the underlying function ρ_{iso} is isotropic with a known radial frequency profile $\widehat{\rho}_{\text{rad}}$ such that $\mathcal{F}\{\rho_{\text{iso}}\}(\boldsymbol{\omega}) = \widehat{\rho}_{\text{rad}}(\|\boldsymbol{\omega}\|)$.

Proposition 1 (Radon transform of isotropic distributions) *Let ρ_{iso} be an isotropic distribution whose radial frequency profile is $\widehat{\rho}_{\text{rad}} : \mathbb{R} \rightarrow \mathbb{R}$. Then,*

$$\text{R}\{\rho_{\text{iso}}(\cdot - \mathbf{x}_0)\}(t, \boldsymbol{\xi}) = \rho_{\text{rad}}(t - \boldsymbol{\xi}^\top \mathbf{x}_0) \quad (41)$$

$$\text{K}_{\text{rad}}\text{R}\{\rho_{\text{iso}}(\cdot - \mathbf{x}_0)\}(t, \boldsymbol{\xi}) = \nu_{\text{rad}}(t - \boldsymbol{\xi}^\top \mathbf{x}_0) \quad (42)$$

with $\rho_{\text{rad}}(t) = \mathcal{F}^{-1}\{\widehat{\rho}_{\text{rad}}(\omega)\}(t)$ and $\nu_{\text{rad}}(t) = \frac{1}{2(2\pi)^{d-1}} \mathcal{F}^{-1}\{|\omega|^{d-1} \widehat{\rho}_{\text{rad}}(\omega)\}(t)$.

The other important building blocks for representing functions are ridges. Specifically, a ridge is a multidimensional function

$$r_{\boldsymbol{\xi}_0} : \mathbb{R}^d \rightarrow \mathbb{R} : \mathbf{x} \mapsto r(\boldsymbol{\xi}_0^\top \mathbf{x}) \quad (43)$$

that is characterized by a profile $r : \mathbb{R} \rightarrow \mathbb{R}$ and a direction $\boldsymbol{\xi}_0 \in \mathbb{S}^{d-1}$. In effect, $r_{\boldsymbol{\xi}_0}$ varies along the axis specified by $\boldsymbol{\xi}_0$ and is constant within any hyperplane perpendicular to $\boldsymbol{\xi}_0$. The connection between ridges and the Radon transform is given by the *ridge identity*

$$\forall \varphi \in \mathcal{S}(\mathbb{R}^d) : \langle r_{\boldsymbol{\xi}_0}, \varphi \rangle = \langle r, \text{R}\{\varphi\}(\cdot, \boldsymbol{\xi}_0) \rangle, \quad (44)$$

where the right-hand side duality product (1D) is well-defined for any $r \in \mathcal{S}'(\mathbb{R})$ because $R\{\varphi\}(\cdot, \xi_0) \in \mathcal{S}(\mathbb{R})$ (by Theorem 1). When the profile $r : \mathbb{R} \rightarrow \mathbb{R}$ is locally integrable, (44) is a simple consequence of Fubini's theorem. For more general distributional profiles $r \in \mathcal{S}'(\mathbb{R})$, we use the ridge identity as definition, which then leads to the following characterization [57].

Theorem 4 (Filtered Radon transform of a ridge) *The filtered Radon transform of the (generalized) ridge r_{ξ_0} with profile $r \in \mathcal{S}'(\mathbb{R}^d)$ and direction $\xi_0 \in \mathbb{S}^{d-1}$ is given by*

$$K_{\text{rad}}R\{r_{\xi_0}\}(t, \xi) = [r(t)\delta(\xi - \xi_0)], \quad (45)$$

where $[r(t)\delta(\xi - \xi_0)] \in \mathcal{S}'_{\text{Rad}}$ is the equivalence class of distributions specified by (29). If $r \in \mathcal{M}(\mathbb{R})$, then the latter has the unique, concrete representer

$$K_{\text{rad}}R\{r_{\xi_0}\}(t, \xi) = \frac{1}{2}(r(t)\delta(\xi - \xi_0) + r(-t)\delta(\xi + \xi_0)) \quad (46)$$

in $\mathcal{M}_{\text{Rad}} = \mathcal{M}_{\text{even}}(\mathbb{R} \times \mathbb{S}^{d-1})$.

An important special case of (46) is the Radon transform of a Dirac ridge: $K_{\text{rad}}R\{\delta(\xi_0^T \cdot - t_0)\} = \delta_{\text{Rad},(t_0, \xi_0)} = \frac{1}{2}(\delta(\cdot - t_0)\delta(\cdot - \xi_0) + \delta(\cdot + t_0)\delta(\cdot + \xi_0))$, which has already been mentioned in Section 3.2 (see also [35, Example 1]).

4 Unifying Variational Formulation

4.1 Representer Theorem for Radon-Domain Regularization

From now on, we shall use the generic symbol \mathcal{X} to designate the hyperspherical Banach space $L_q(\mathbb{R} \times \mathbb{S}^{d-1})$ with $q = (1, \infty)$ or $C_0(\mathbb{R} \times \mathbb{S}^{d-1})$ for $q = \infty$, which fall into the category described by (33) with $\|\cdot\|_{\mathcal{X}} = \|\cdot\|_{L_q}$.

The formulation of Theorem 5 below requires the specification of a *native space*, $\mathcal{X}'_{L_R}(\mathbb{R}^d)$, that is tied to a Radon-domain norm $\|\cdot\|_{\mathcal{X}'}$ and an admissible regularization operator L_R . Informally, our native space is the largest function space over which the proposed regularization functional $f \mapsto \|L_R\{f\}\|_{\mathcal{X}'}$ is well-defined. The precise description of $\mathcal{X}'_{L_R}(\mathbb{R}^d)$, however, is a bit more involved. As first step, we need to restrict the dual pair $(\mathcal{X}, \mathcal{X}')$ to the range of the (filtered) Radon transform. This yields the Banach spaces $(\mathcal{X}_{\text{Rad}}, \mathcal{X}'_{\text{Rad}})$, as defined by (36) and (37), with the pairs of interest being $(C_{0, \text{Rad}}, \mathcal{M}_{0, \text{Rad}})$ and $(L_{q, \text{Rad}}, L_{p, \text{Rad}})$ with $\frac{1}{p} + \frac{1}{q} = 1$ and $p \in [1, \infty)$. Given some spline-admissible operator L (Definition 2 and 3), we then define our regularization operator and its adjoint as

$$\begin{aligned} L_R &\stackrel{\Delta}{=} K_{\text{rad}}RL : \mathcal{X}'_{L_R}(\mathbb{R}^d) \rightarrow \mathcal{X}'_{\text{Rad}}, \\ L_R^* &= L^*R^*K_{\text{rad}} : \mathcal{X}_{\text{Rad}} \rightarrow \mathcal{X}'_{L_R}(\mathbb{R}^d) \end{aligned}$$

where R denotes the Radon transform and K_{rad} is the (self-adjoint) filtering operator such that $K_{\text{rad}}RR^* = \text{Id}$ on $\mathcal{X}'_{\text{Rad}}$ (Theorem 3). In order to establish

isometries, one needs to be able to invert $L : \mathcal{X}'_{L_R}(\mathbb{R}^d) \rightarrow \mathbb{R}^*(\mathcal{X}'_{\text{Rad}})$ as well as L_R , which is feasible if one factors out the null space \mathcal{P} , which is common to both. This motivates us to define the directed inverse operators

$$L_{\mathcal{P}}^{-1} \triangleq (\text{Id} - \text{Proj}_{\mathcal{P}})L^{-1} : \mathbb{R}^*(\mathcal{X}'_{\text{Rad}}) \rightarrow \mathcal{X}'_{L_R}(\mathbb{R}^d) \quad (\text{Right inverse of } L)$$

$$L_{\mathcal{P}}^{-1*} = L^{-1*}(\text{Id} - \text{Proj}_{\mathcal{P}'}) : \mathcal{X}'_{L_R}(\mathbb{R}^d) \rightarrow \mathbb{K}_{\text{rad}}\mathbb{R}(\mathcal{X}'_{\text{Rad}}) \quad (\text{Left inverse of } L^*)$$

$$L_R^\dagger \triangleq L_{\mathcal{P}}^{-1}L^* : \mathcal{X}'_{\text{Rad}} \rightarrow \mathcal{X}'_{L_R}(\mathbb{R}^d) \quad (47)$$

$$L_R^{*\dagger} = RL_{\mathcal{P}}^{-1*} : \mathcal{X}'_{L_R}(\mathbb{R}^d) \rightarrow \mathcal{X}'_{\text{Rad}} \quad (48)$$

where the operators L_R^\dagger and $L_R^{*\dagger}$ are generalized inverses³ of L_R and L_R^* , respectively. We now have all the ingredients to specify our native space as

$$\begin{aligned} \mathcal{X}'_{L_R}(\mathbb{R}^d) &= L_R^\dagger(\mathcal{X}'_{\text{Rad}}) \oplus \mathcal{P} \\ &= \{f \in L_{\infty, -n_0}(\mathbb{R}^d) : \|L_R\{f\}\|_{\mathcal{X}'} + \|\text{Proj}_{\mathcal{P}}\{f\}\|_{\mathcal{P}} < \infty\} \\ &= \{L_R^\dagger\{w\} + p_0 : (w, p_0) \in \mathcal{X}'_{\text{Rad}} \times \mathcal{P}\}, \end{aligned} \quad (49)$$

which is isometrically isomorphic to $\mathcal{X}'_{\text{Rad}} \times \mathcal{P}$, as expressed by (49). The key property there is that $L_R L_R^\dagger = \text{Id}$ on $\mathcal{X}'_{\text{Rad}}$, while $L_R\{p_0\} = 0$ for all $p_0 \in \mathcal{P}$ (Theorem 6). Moreover, $\mathcal{X}'_{L_R}(\mathbb{R}^d)$ is the topological dual of the *predual space*

$$\begin{aligned} \mathcal{X}_{L_R}(\mathbb{R}^d) &= L_R^*(\mathcal{X}'_{\text{Rad}}) \oplus \mathcal{P}' \\ &= \{\nu \in \mathcal{S}'(\mathbb{R}^d) : \|\nu\|_{\mathcal{X}_{L_R}} = \max(\|L_R^{*\dagger}\{\nu\}\|_{\mathcal{X}'}, \|\text{Proj}_{\mathcal{P}'}\{\nu\}\|_{\mathcal{P}'}) < \infty\} \\ &= \{L_R^*\{v\} + p_0^* : (v, p_0^*) \in \mathcal{X}'_{\text{Rad}} \times \mathcal{P}'\}, \end{aligned} \quad (50)$$

which is a Banach space, as shown in Theorem 9. The validity of this dual pairing can be checked formally in the absence of null space components: For any $(f, \nu) \in \mathcal{X}'_{L_R}(\mathbb{R}^d) \times \mathcal{X}_{L_R}(\mathbb{R}^d)$ with $\text{Proj}_{\mathcal{P}}\{f\} = 0$ and $\text{Proj}_{\mathcal{P}'}\{\nu\} = 0$, we have that

$$\langle f, \nu \rangle = \langle L_R^\dagger\{w\}, L_R^*\{v\} \rangle = \langle L_R L_R^\dagger\{w\}, v \rangle_{\text{Rad}} = \langle w, v \rangle_{\text{Rad}}$$

with $(w, v) \in \mathcal{X}'_{\text{Rad}} \times \mathcal{X}'_{\text{Rad}}$. Finally, since $\text{Proj}_{\mathcal{P}}$ continuously maps $\mathcal{X}'_{L_R}(\mathbb{R}^d) \rightarrow \mathcal{P} \xrightarrow{\text{iso.}} \mathcal{X}'_{L_R}(\mathbb{R}^d)$, we can identify $\text{Proj}_{\mathcal{P}'}$ as its adjoint $(\text{Proj}_{\mathcal{P}})^* = \text{Proj}_{\mathcal{P}'} : \mathcal{X}_{L_R}(\mathbb{R}^d) \rightarrow \mathcal{P}' \xrightarrow{\text{iso.}} \mathcal{X}_{L_R}(\mathbb{R}^d)$.

To ensure that the generic regression problem in Theorem 5 is well-defined, we also need a mild hypothesis on the structure of the data points.

Definition 7 (Admissible data points) Let $N_0 = \dim \mathcal{P}$ where $\mathcal{P} = \mathcal{P}_{n_0}$ is the polynomial null space of L_R . Then, the set of data points $\{\mathbf{x}_1, \dots, \mathbf{x}_M\} \subset \mathbb{R}^d$ with $M \geq N_0$ is said to be *\mathcal{P} -admissible* if the sampling matrix $\mathbf{H} = [\mathbf{v}_1 \ \dots \ \mathbf{v}_M]^\top \in \mathbb{R}^{M \times N_0}$ with $\mathbf{v}_m = (\langle \delta(\cdot - \mathbf{x}_m), m_{\mathbf{k}} \rangle)_{|\mathbf{k}| \leq n_0} \in \mathbb{R}^{N_0}$ is of rank N_0 .

³ The precise properties of these inverse operators are stated in Theorem 6 and 9.

The rank condition is precisely the condition under which the classical least-squares polynomial fitting problem

$$\min_{f \in \mathcal{P}} \sum_{m=1}^M |y_m - f(\mathbf{x}_m)|^2 = \min_{\mathbf{b} \in \mathbb{R}^{N_0}} \|\mathbf{y} - \mathbf{H}\mathbf{b}\|^2 \quad (51)$$

is well-posed. Indeed, the formal solution of (51) is $p_0(\mathbf{x}) = \sum_{|\mathbf{k}| \leq n_0} b_{\mathbf{k}} m_{\mathbf{k}}(\mathbf{x})$ with $\mathbf{b} = (b_{\mathbf{k}}) = (\mathbf{H}^T \mathbf{H})^{-1} \mathbf{H}^T \mathbf{y}$ where the rank condition guarantees the invertibility of the normal matrix $(\mathbf{H}^T \mathbf{H}) \in \mathbb{R}^{N_0 \times N_0}$.

We are now ready to formulate our extended representer theorem, which supports a rich panorama of regression models. A case of direct practical relevance is obtained by setting $\mathbf{L} = (-\Delta)^{\frac{\alpha+1}{2}}$ (fractional Laplacian) with $\widehat{L}_{\text{rad}}(\omega) = |\omega|^{\alpha+1}$. Indeed, we shall see that, for $\mathcal{X} = L_2$, this essentially yields the classical thin-plate splines (see Section 4.2), while, for $\mathcal{X}' = \mathcal{M}$, it results in neural networks with activation functions (ρ_{rad}) labelled as ‘‘ridge splines’’ (including ReLU) and ‘‘fractional splines’’ in Table 1 (Section 4.5).

Theorem 5 *Let us consider the following setting.*

- A proper, lower-semicontinuous, coercive, and convex loss functional $E : \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}^+ \cup \{+\infty\}$.
- An isotropic, spline-admissible operator \mathbf{L} with frequency profile \widehat{L}_{rad} and polynomial null space \mathcal{P}_{n_0} of degree n_0 , possibly trivial with the convention that $\mathcal{P}_{-1} = \{0\}$.
- A convex, increasing function $\psi : \mathbb{R}^+ \rightarrow \mathbb{R}^+$.
- A set $\{\mathbf{x}_1, \dots, \mathbf{x}_M\} \subset \mathbb{R}^d$ of \mathcal{P}_{n_0} -admissible data points.

Then, for any fixed $\mathbf{y} = (y_m) \in \mathbb{R}^M$, the generic functional-optimization problem

$$S = \arg \min_{f \in \mathcal{X}'_{\mathbf{L}_R}(\mathbb{R}^d)} \left(\sum_{m=1}^M E(y_m, f(\mathbf{x}_m)) + \psi(\|\mathbf{L}_R\{f\}\|_{\mathcal{X}'}), \right) \quad (52)$$

with $\mathbf{L}_R = \mathbf{K}_{\text{rad}} \mathbf{R}\mathbf{L}$, and ψ, \mathcal{X}' as stated below, always has a solution.

1. When $\mathcal{X} = \mathcal{X}' = L_2(\mathbb{R} \times \mathbb{S}^{d-1})$ and ψ is strictly convex, the solution of (52) is unique and representable by the linear kernel expansion

$$f(\mathbf{x}) = p_0(\mathbf{x}) + \sum_{m=1}^M a_m \rho_{\text{iso}}(\mathbf{x} - \mathbf{x}_m), \quad (53)$$

where $\rho_{\text{iso}} = 2(2\pi)^{d-1} \mathcal{F}^{-1}\{1/(\widehat{L}_{\text{rad}}(\|\boldsymbol{\omega}\|)^2 \|\boldsymbol{\omega}\|^{d-1})\}$ is a radial basis function, $(a_m) \in \mathbb{R}^M$ an adequate set of coefficients, and $p_0 \in \mathcal{P}_{n_0}$ a polynomial that lies in the null space of \mathbf{L}_R .

2. When $\mathcal{X}' = L_p(\mathbb{R} \times \mathbb{S}^{d-1})$ with $p \in (1, 2]$ and ψ is strictly convex, the solution is unique and admits the parametric representation

$$f(\mathbf{x}) = p_0(\mathbf{x}) + L_{\mathbb{R}}^{\dagger} \circ J_q \left\{ \sum_{m=1}^M a_m \nu_{\mathbf{x}_m} \right\}(\mathbf{x}) \quad (54)$$

with basis functions $\nu_{\mathbf{x}_1}, \dots, \nu_{\mathbf{x}_M} \in L_q(\mathbb{R} \times \mathbb{S}^{d-1})$ specified by (70) and parameters $(a_m) \in \mathbb{R}^M$, $p_0 \in \mathcal{P}_{n_0}$, and where J_q is the pointwise nonlinearity given by (92) with $q = p/(p-1)$. (The latter is the duality map $J_{\mathcal{X}'} : \mathcal{X}' \rightarrow \mathcal{X}'$ with $\mathcal{X}' = L_q(\mathbb{R} \times \mathbb{S}^{d-1})$ —see Appendix A for explanations).

3. When $\mathcal{X}' = \mathcal{M}(\mathbb{R} \times \mathbb{S}^{d-1})$ and ψ is strictly increasing, the solution set is the weak*-closed convex hull of its extreme points, which are all of the form

$$f_{\text{ext}}(\mathbf{x}) = p_0(\mathbf{x}) + \sum_{k=1}^{K_0} a_k \rho_{\text{rad}}(\boldsymbol{\xi}_k^{\top} \mathbf{x} - \tau_k) \quad (55)$$

with activation function $\rho_{\text{rad}} = \mathcal{F}^{-1}\{1/\widehat{L}_{\text{rad}}\}$, for some $K_0 \leq M - \dim \mathcal{P}_{n_0}$, $(a_k, \boldsymbol{\xi}_k, \tau_k) \in \mathbb{R} \times \mathbb{S}^{d-1} \times \mathbb{R}$ for $k = 1, \dots, K_0$, and a null-space component $p_0 \in \mathcal{P}_{n_0}$. The optimal regularization cost associated with (55) is $\|L_{\mathbb{R}} f_{\text{ext}}\|_{\mathcal{M}} = \sum_{k=1}^{K_0} |a_k|$ and is shared by all solutions.

Proof Since $L_{\mathbb{R}}^*$ is injective on \mathcal{S}_{Rad} and, by extension, on the completed space \mathcal{X}_{Rad} , the image space $\mathcal{U} = L_{\mathbb{R}}^*(\mathcal{X}_{\text{Rad}})$ is complete as well (see proof of Theorem 9 for the details of the construction of \mathcal{U}). Its continuous dual is given by $\mathcal{U}' = L_{\mathcal{P}}^{-1} \mathbb{R}^*(\mathcal{X}'_{\text{Rad}})$, in reason of the identities $\mathbb{R} L_{\mathcal{P}}^{-1} L_{\mathbb{R}}^* = \mathbb{R} L_{\mathcal{P}}^{-1} L^* \mathbb{R}^* K_{\text{rad}} = \mathbb{R} \mathbb{R}^* K_{\text{rad}} = \text{Id}$ on \mathcal{X}_{Rad} . Likewise, \mathcal{P}' , as identified by (15), is a finite-dimensional Banach space. Its dual is simply $(\mathcal{P}')' = \mathcal{P}$ (the null space of both L and $L_{\mathbb{R}}$), owing to the property that all finite-dimensional spaces are reflexive. Using the notation for direct-sum topologies of [58], we then observe that $\mathcal{X}'_{L_{\mathbb{R}}}(\mathbb{R}^d) = (\mathcal{U} \oplus \mathcal{P}')_{\ell_{\infty}}$, whose formal dual $(\mathcal{U} \oplus \mathcal{P}')'_{\ell_{\infty}} = (\mathcal{U}' \oplus \mathcal{P})_{\ell_1}$ is precisely the native space $\mathcal{X}'_{L_{\mathbb{R}}}(\mathbb{R}^d)$ described by (49). By writing $f = L_{\mathcal{P}}^{-1} \mathbb{R}^* \{w\} + p_0$ and recalling that $L L_{\mathcal{P}}^{-1} = \text{Id}$ (right-inverse property), we then identify the regularization functional as

$$\begin{aligned} \|L_{\mathbb{R}} f\|_{\mathcal{X}'} &= \|K_{\text{rad}} \mathbb{R} \{L L_{\mathcal{P}}^{-1} \mathbb{R}^* w + L p_0\}\|_{\mathcal{X}'} \\ &= \|K_{\text{rad}} \mathbb{R} \mathbb{R}^* \{w\} + K_{\text{rad}} \mathbb{R} \{0\}\|_{\mathcal{X}'} \\ &= \|w\|_{\mathcal{X}'_{\text{Rad}}} = \|\text{Proj}_{\mathcal{U}'} f\|_{\mathcal{U}'} \end{aligned}$$

which, in effect, converts the seminorm over $\mathcal{X}'_{L_{\mathbb{R}}}(\mathbb{R}^d)$ into a norm over \mathcal{U}' by factoring out the null space of $L_{\mathbb{R}}$. The other mathematical ingredient for the optimization problem (52) to be well-posed is the weak* continuity of the sampling functionals $\delta(\cdot - \mathbf{x}_m) : f \mapsto f(\mathbf{x}_m)$ in the underlying topology. This is equivalent to $\delta(\cdot - \mathbf{x}_m) \in \mathcal{X}'_{L_{\mathbb{R}}}(\mathbb{R}^d)$ for any $\mathbf{x}_m \in \mathbb{R}^d$. In the present context, this condition can be reframed as $\nu_{\mathbf{x}_m} = L_{\mathbb{R}}^{*\dagger} \{\delta(\cdot - \mathbf{x}_m)\} \in \mathcal{X}_{\text{Rad}}$, which is a

property that is established in Theorem 7 for the cases $\mathcal{X} = C_0$ and $\mathcal{X} = L_q$ for any $q \geq 2$.

The existence of the solution and the parametric descriptions (53), (54), and (55) then follows from the three cases of the abstract representer theorem for direct sums [58, Theorem 3]. The link with the abstract theorem is made by identifying $\mathcal{N}_{\mathcal{P}} = \mathcal{P}$, $\mathcal{U} \oplus \mathcal{N}_{\mathcal{P}^*} = \mathcal{X}_{L_R}(\mathbb{R}^d)$, $\mathcal{U}' \oplus \mathcal{N}_{\mathcal{P}'} = \mathcal{X}'_{L_R}(\mathbb{R}^d)$ and $\nu_m = \delta(\cdot - \mathbf{x}_m)$ for $m = 1, \dots, M$. As for the required technical assumptions, they directly follow from the weak* continuity of the sampling functionals (i.e., $\nu_m \in \mathcal{U} \oplus \mathcal{N}_{\mathcal{P}^*}$) and the \mathcal{P} -admissibility hypothesis in Definition 7, with the \mathbf{v}_m being the same as in the statement of the abstract theorem. The duality map that is required for the first and second scenarios is $J_{\mathcal{U}} = L_R^\dagger \circ J \circ L_R^{*\dagger} : \mathcal{U} \rightarrow \mathcal{U}'$ (see Appendix A, Proposition 4 with $T = L_R^*$, $T^{-1} = L_R^{*\dagger} = R^* L_{\mathcal{P}'}^{-1}$, and $(T^*)^{-1} = L_R^\dagger = L_{\mathcal{P}}^{-1} R^*$).

To describe the solution set for the non-reflexive case $\mathcal{X}' = \mathcal{M}$, we invoke the third case of [58, Theorem 3], which tells us that the extreme points of S can all be expressed as the sum of a null-space component plus a linear combination of $K_0 \leq M - \dim(\mathcal{P})$ atoms e_k that are selected adaptively within a dictionary consisting of the extreme points of the unit ball in \mathcal{U}' : $B(\mathcal{U}') = \{u \in \mathcal{U}' : \|u\|_{\mathcal{U}'} \leq 1\}$. Because of the linear isomorphism $\mathcal{U}' = L_R^\dagger(\mathcal{M}_{\text{Rad}})$, $\text{ext}B(\mathcal{U}') = L_R^\dagger(\text{ext}B(\mathcal{M}_{\text{Rad}}))$. Next, we use the property that $\mathcal{M}_{\text{Rad}} = \mathcal{M}_{\text{even}}(\mathbb{R} \times \mathbb{S}^{d-1})$ whose extreme points are $\{\pm \delta_{\text{Rad}, \mathbf{z}}\}_{\mathbf{z} \in \mathbb{P}^d}$. Each $\delta_{\text{Rad}, \mathbf{z}_k} \in \text{ext}B(\mathcal{M}_{\text{Rad}})$ then bijectively maps into an extreme point $e_k = L_R^\dagger\{\delta_{\text{Rad}, \mathbf{z}_k}\} \in \text{ext}B(\mathcal{U}')$, and vice versa. Finally, by recalling that $L_R^\dagger = (\text{Id} - \text{Proj}_{\mathcal{P}})L^{-1}R^*$ and by invoking Theorem 4 to show that $L^{-1}R^*\{\delta_{\text{Rad}, (t_k, \boldsymbol{\xi}_k)}\} = L^{-1}R^*K_{\text{rad}}R\{\delta(\boldsymbol{\xi}_k^\top \mathbf{x} - t_k)\} = L^{-1}\{\delta(\boldsymbol{\xi}_k^\top \mathbf{x} - t_k)\} = \rho_{\text{rad}}(\boldsymbol{\xi}_k^\top \mathbf{x} - t_k)$, we find that

$$e_k = \pm \rho_{\text{rad}}(\boldsymbol{\xi}_k^\top \mathbf{x} - t_k) \mp p_{0,k}, \quad (56)$$

where $p_{0,k} = \text{Proj}_{\mathcal{P}}\{\rho_{\text{rad}}(\boldsymbol{\xi}_k^\top \mathbf{x} - \tau_k)\} \in \mathcal{P}$. Since every extreme point of $B(\mathcal{U}')$ is necessarily of the form (56), we can substitute this expression in the generic expansion $f_{\text{ext}} = \tilde{p}_0 + \sum_{k=1}^{K_0} a_k e_k$ which, upon collection of all null-space components, yields (55). \square

The two cases in Theorem 5 that are of direct practical relevance to machine learning are Items 1 and 3. The first scenario yields a learning architecture that is a linear expansion of RBFs, which also has a classical RKHS interpretation, as explained in Section 4.2.

The form of the solution in Item 3 is equivalent to a shallow network with the weights $(\boldsymbol{\xi}_k)$ of the hidden layer being normalized. It actually turns out that this normalization is inconsequential when the activation is a homogeneous function. This happens for instance when the regularization operator $L = (-\Delta)^{\frac{\alpha+1}{2}}$ is a fractional Laplacian which maps into $\rho_{\text{Rad}}(t) \propto |t|^\alpha$. Indeed, for any $(\mathbf{w}_k, b_k) \in \mathbb{R}^d \times \mathbb{R}$, we have that

$$|\mathbf{w}_k^\top \mathbf{x} - b_k|^\alpha = \|\mathbf{w}_k\|^\alpha |\boldsymbol{\xi}_k^\top \mathbf{x} - t_k|^\alpha \quad (57)$$

with $\boldsymbol{\xi}_k = \mathbf{w}_k / \|\mathbf{w}_k\|$ and $t_k = b_k / \|\mathbf{w}_k\|$, which indicates that the normalization (or lack thereof) can be absorbed in the weights (a_k) of the output layer. The case $\alpha = 1$ with $L = \Delta$ (Laplacian) is particularly attractive, as it nicely maps into a ReLU network with one hidden layer and a skip connection to implement the affine component [37].

The form of the solution in Item 2 is more involved, but still useful to get insight into the transition with p from RBFs to neural nets. The equivalence with Item 1 for $p = 2$ is demonstrated in the Section 4.2. As for the behavior as $p \rightarrow 1$, we observe that the effect of the duality map J_q as $q \rightarrow \infty$ is to pull a few maximal values to infinity, while attenuating all other (non-supremum) values towards zero. In effect, this means that $J_q\{\sum_{m=1}^M a_m \nu_{\mathbf{x}_m}\}$ will exhibit peaks that become more and more pronounced, and eventually converge to a sum of impulses as $p \rightarrow 1$, which is consistent with the limit form given by (55).

4.2 Connection with RKHS Methods

The scenario $\mathcal{X} = L_2$ in Theorem 5 is compatible with the kernel models of “classical” machine learning [40, 50]. This is because the underlying native space is a reproducing-kernel Hilbert space whose topological structure is now made explicit.

Proposition 2 (Characterization of RKHS for $\mathcal{X} = L_2$) *Let L be a spline admissible operator with a polynomial null space of degree n_0 and consider the self-adjoint operator $A = (L^*KL)^{-1}$. Then, the native space $\mathcal{H} = L'_{2,LR}(\mathbb{R}^d)$, defined by (49) with $\mathcal{X}' = \mathcal{X} = L_2$, is the reproducing-kernel Hilbert space $\mathcal{H} = \mathcal{U}' \oplus \mathcal{P}$ associated with the composite inner product*

$$\begin{aligned} \langle f_1, f_2 \rangle_{\mathcal{H}} &= \langle L_R\{f_1\}, L_R\{f_2\} \rangle_{\text{Rad}} + \sum_{|\mathbf{k}| \leq n_0} \langle m_{\mathbf{k}}^*, f_1 \rangle \langle m_{\mathbf{k}}^*, f_2 \rangle \\ &= \langle (L^*KL)\{f_1\}, f_2 \rangle + \sum_{|\mathbf{k}| \leq n_0} \langle m_{\mathbf{k}}^*, f_1 \rangle \langle m_{\mathbf{k}}^*, f_2 \rangle \end{aligned} \quad (58)$$

whose leading term can also be written as

$$\begin{aligned} \langle (L^*KL)\{f_1\}, f_2 \rangle &= \langle (L^*KL)\text{Proj}_{\mathcal{U}'}\{f_1\}, \text{Proj}_{\mathcal{U}'}\{f_2\} \rangle \\ &= \langle \text{Proj}_{\mathcal{U}'}\{f_1\}, \text{Proj}_{\mathcal{U}'}\{f_2\} \rangle_{\mathcal{U}'}, \end{aligned} \quad (59)$$

where $\text{Proj}_{\mathcal{U}'} = \text{Id} - \text{Proj}_{\mathcal{P}} : \mathcal{H} \rightarrow \mathcal{U}'$. The topological dual of \mathcal{H} is the Hilbert space $\mathcal{H}' = L_{2,LR}(\mathbb{R}^d) = \mathcal{U} \oplus \mathcal{P}'$ equipped with the inner product

$$\begin{aligned} \langle g_1, g_2 \rangle_{\mathcal{H}'} &= \langle L_R^{\dagger}\{g_1\}, L_R^{\dagger}\{g_2\} \rangle_{\text{Rad}} + \sum_{|\mathbf{k}| \leq n_0} \langle m_{\mathbf{k}}, g_1 \rangle \langle m_{\mathbf{k}}, g_2 \rangle \\ &= \langle L_R^{\dagger}L_R^{\dagger}\{g_1\}, g_2 \rangle + \sum_{|\mathbf{k}| \leq n_0} \langle m_{\mathbf{k}}, g_1 \rangle \langle m_{\mathbf{k}}, g_2 \rangle \\ &= \langle A\text{Proj}_{\mathcal{U}}\{g_1\}, \text{Proj}_{\mathcal{U}}\{g_2\} \rangle + \sum_{|\mathbf{k}| \leq n_0} \langle m_{\mathbf{k}}, g_1 \rangle \langle m_{\mathbf{k}}, g_2 \rangle, \end{aligned} \quad (60)$$

where $\text{Proj}_{\mathcal{U}} = \text{Id} - \text{Proj}_{\mathcal{P}'}$, $\mathcal{H}' \rightarrow \mathcal{U}$. The corresponding linear isometries (Riesz maps) that map $\mathcal{U} \rightarrow \mathcal{U}'$, $\mathcal{P}' \rightarrow \mathcal{P}$, and $\mathcal{H} \rightarrow \mathcal{H}'$ are

$$\begin{aligned} J_{\mathcal{U}} &= A : \mathcal{U} \rightarrow \mathcal{U}' \\ &= (\text{Id} - \text{Proj}_{\mathcal{P}})A(\text{Id} - \text{Proj}_{\mathcal{P}'}) : \mathcal{U} \oplus \mathcal{P}' \rightarrow \mathcal{U}', \\ J_{\mathcal{P}'} &= \text{Proj}_{\mathcal{P}} : \mathcal{P}' \rightarrow \mathcal{P}, \\ &= \text{Proj}_{\mathcal{P}}\text{Proj}_{\mathcal{P}'} : \mathcal{U} \oplus \mathcal{P}' \rightarrow \mathcal{P}, \\ J_{\mathcal{H}} &= ((\text{Id} - \text{Proj}_{\mathcal{P}})A(\text{Id} - \text{Proj}_{\mathcal{P}'}) + \text{Proj}_{\mathcal{P}}\text{Proj}_{\mathcal{P}'}) : \mathcal{H} \rightarrow \mathcal{H}', \end{aligned}$$

where the second, alternative forms of $J_{\mathcal{U}}$ and $J_{\mathcal{P}'}$ that include projectors are the proper extension of those operators to the whole space $\mathcal{H} = \mathcal{U}' \oplus \mathcal{P}$.

Proposition 2 is obtained as a corollary of Theorems 6 and 9 with $\mathcal{X} = \mathcal{X}' = L_2$, with the help of the intertwining property $L_R = K_{\text{rad}}RL = RKL$ (see the discussion at the end of Section 3.2). The technical part is to establish the completeness of \mathcal{U} (resp., $\mathcal{H}' = \mathcal{U} \oplus \mathcal{P}'$), which then implies that of \mathcal{U}' (resp., $\mathcal{H} = \mathcal{H}' = \mathcal{U}' \oplus \mathcal{P}$) by duality. As for the Hilbert-space property, the maps defined by (58) and (60) are obviously bilinear and symmetric. To show that (58) is also positive definite, we invoke Theorem 6, which states that any $f \in \mathcal{H} = \mathcal{U}' \oplus \mathcal{P}$ has a unique decomposition as $f = L_R^\dagger\{w\} + p_0$, with $w = L_R\{f\} \in \mathcal{X}'$ and $p_0 = \text{Proj}_{\mathcal{P}}\{f\} = \sum_{|\mathbf{k}| \leq n_0} b_{\mathbf{k}}m_{\mathbf{k}} \in \mathcal{P}$ with $b_{\mathbf{k}} = \langle m_{\mathbf{k}}^*, f \rangle$. This then yields that

$$\langle f, f \rangle_{\mathcal{H}} = \langle L_R\{f\}, L_R\{f\} \rangle + \sum_{|\mathbf{k}| \leq n_0} |b_{\mathbf{k}}|^2 = \underbrace{\langle w, w \rangle_{L_2(\mathbb{R} \times \mathbb{S}^{d-1})}}_{\|w\|_{\mathcal{X}'}^2} + \underbrace{\|\mathbf{b}\|_2^2}_{\|p_0\|_{\mathcal{P}}^2} \geq 0,$$

which vanishes if and only if $f = 0$. Likewise, one readily verifies that the semi-inner products associated with each individual term in (60) induce the two component norms $\|v\|_{\mathcal{X}}$ and $\|p_0^*\|_{\mathcal{P}'}$ for $\mathcal{X} = L_2$ that appear in the definition (50) of the predual space.

The denomination RKHS applies to any Hilbert space \mathcal{H} of functions on \mathbb{R}^d such that $\delta(\cdot - \mathbf{x}_0) \in \mathcal{H}'$ for any $\mathbf{x}_0 \in \mathbb{R}^d$. In our case, this property is equivalent to $L_R^{*\dagger}\{\delta(\cdot - \mathbf{x}_0)\} \in L_{2, \text{Rad}}$, which follows from Theorem 7.

To get further insight, we now show that the RBF solution (53) is a particular case of the L_p solution (54) with $p = 2$. Since $L_2 = (L_2)'$ is its own

dual, $\mathbf{J} = \text{Id}$, which allows us to manipulate (54) as

$$\begin{aligned} f &= p_0 + \mathbf{L}_{\mathcal{P}}^{-1} \mathbf{R}^* \mathbf{R} \mathbf{L}_{\mathcal{P}}^{-1*} \left\{ \sum_{m=1}^M a_m \delta(\cdot - \mathbf{x}_m) \right\} \\ &= p_0 + \mathbf{L}_{\mathcal{P}}^{-1} \mathbf{K}^{-1} \mathbf{L}_{\mathcal{P}}^{-1*} \left\{ \sum_{m=1}^M a_m \delta(\cdot - \mathbf{x}_m) \right\} \\ &= p_0 + (\text{Id} - \text{Proj}_{\mathcal{P}}) \mathbf{A} (\text{Id} - \text{Proj}_{\mathcal{P}'}) \left\{ \sum_{m=1}^M a_m \delta(\cdot - \mathbf{x}_m) \right\} \end{aligned} \quad (61)$$

$$= p_0 + \mathbf{A} \left\{ \sum_{m=1}^M a_m \delta(\cdot - \mathbf{x}_m) \right\} = p_0 + \sum_{m=1}^M a_m \rho_{\text{iso}}(\mathbf{x} - \mathbf{x}_m), \quad (62)$$

where $\rho_{\text{iso}} = (\mathbf{L}^* \mathbf{K} \mathbf{L})^{-1} \{\delta\} = \mathbf{A} \{\delta\} : \mathbb{R}^d \rightarrow \mathbb{R}$ is the Green's function of $(\mathbf{L}^* \mathbf{K} \mathbf{L})$ and $p_0 \in \mathcal{P}$. The nonobvious simplification from (61) to (62) results from two crucial observations: (1) the ‘‘orthogonality-to-the-null-space’’ condition $\sum_{m=1}^M a_m \delta(\cdot - \mathbf{x}_m) \in \mathcal{U}$ is necessary for optimality; and (2) the availability of the identity

$$\forall u \in \mathcal{U} : (\text{Id} - \text{Proj}_{\mathcal{P}}) \mathbf{A} (\text{Id} - \text{Proj}_{\mathcal{P}'}) \{u\} = \mathbf{A} \{u\} = u^* \in \mathcal{U}',$$

which is tightly linked to the specification of the underlying spaces in Proposition 2. Likewise, we find that the quadratic regularization cost associated with the linear model (62) is $\mathbf{a}^T \mathbf{G} \mathbf{a}$, where $\mathbf{G} \in \mathbb{R}^{M \times M}$ is a symmetric, conditionally positive-definite matrix (see [31]) whose entries are calculated as follows:

$$\begin{aligned} [\mathbf{G}]_{m,n} &= \langle \mathbf{L}_R \rho_{\text{iso}}(\cdot - \mathbf{x}_m), \mathbf{L}_R \rho_{\text{iso}}(\cdot - \mathbf{x}_n) \rangle_{\text{Rad}} \\ &= \langle \text{RKL} \{ \rho_{\text{iso}}(\cdot - \mathbf{x}_m) \}, \text{RKL} \{ \rho_{\text{iso}}(\cdot - \mathbf{x}_n) \} \rangle_{\text{Rad}} \\ &= \langle \text{KL} \{ \rho_{\text{iso}}(\cdot - \mathbf{x}_m) \}, \mathbf{R}^* \text{RKL} \{ \rho_{\text{iso}}(\cdot - \mathbf{x}_n) \} \rangle \\ &= \langle \rho_{\text{iso}}(\cdot - \mathbf{x}_m), \mathbf{L}^* \mathbf{K} \underbrace{\mathbf{R}^* \mathbf{R}}_{\text{Id}} \mathbf{L} \{ \rho_{\text{iso}}(\cdot - \mathbf{x}_n) \} \rangle \\ &= \langle \rho_{\text{iso}}(\cdot - \mathbf{x}_m), \delta(\cdot - \mathbf{x}_n) \rangle = \rho_{\text{iso}}(\mathbf{x}_n - \mathbf{x}_m) = \rho_{\text{iso}}(\mathbf{x}_m - \mathbf{x}_n). \end{aligned}$$

As variant of the L_2 result in Theorem 5, we may also consider the modified regularization operator $\tilde{\mathbf{L}} = \mathbf{K}^{-\frac{1}{2}} \mathbf{L}$ whose frequency response is $\widehat{\tilde{\mathbf{L}}}(\boldsymbol{\omega}) = \sqrt{2}(2\pi)^{(d-1)/2} \widehat{\mathbf{L}}(\boldsymbol{\omega}) / \|\boldsymbol{\omega}\|^{(d-1)/2}$. For this particular setting, $\tilde{\mathbf{L}}_R = \mathbf{R} \mathbf{K}^{\frac{1}{2}} \mathbf{L}$, which translates into

$$\|\tilde{\mathbf{L}}_R \{f\}\|_{L_2(\mathbb{R} \times \mathbb{S}^{d-1})}^2 = \|\mathbf{R} \mathbf{K}^{\frac{1}{2}} \mathbf{L} \{f\}\|_{L_2(\mathbb{R} \times \mathbb{S}^{d-1})}^2 = \|\mathbf{L} \{f\}\|_{L_2(\mathbb{R}^d)}^2,$$

owing to the property that $\mathbf{R} \mathbf{K}^{\frac{1}{2}}$ is an L_2 isometry [27]. The proposed Radon-domain regularization therefore reduces to the standard energy functional associated with (semi-)reproducing-kernel Hilbert spaces. The resulting basis function is $\tilde{\rho}_{\text{iso}}(\mathbf{x}) = \mathcal{F}^{-1} \{1/|\widehat{\tilde{\mathbf{L}}}|^2\}(\mathbf{x})$, which is the same as the one encountered in the classical formulation that does not involve the Radon transform.

While this may suggest that the two formulations are equivalent, there is an important difference that concerns the dimension of the null space.

As a matter of illustration, we now compare two schemes that utilize the same “linear” radial basis $\rho_{\text{iso}}(\mathbf{x}) \propto \|\mathbf{x}\|$. Based on (6), we deduce that this corresponds to the choice $|\widehat{L}(\boldsymbol{\omega})|^2 = \|\boldsymbol{\omega}\|^{d+1}$ in the classical formulation, which induces the regularization norm $\|(-\Delta)^{(d+1)/2}\{f\}\|_{L_2}$ with a polynomial null space of degree $n_0 = \lceil (d+1)/2 \rceil > d/2$ [17, 63]. In our proposed Radon-domain variant, the appropriate regularization norm is $\|\text{RK}(-\Delta)^{1/2}\{f\}\|_{L_2}$ with a null space of polynomial degree $n_0 = 0$. This solution is attractive because it does not depend on the dimensionality of the data.

4.3 Universal Approximation Properties

The universal-approximation properties of the supervised-learning scheme specified by (52) are supported by Theorem 6 below, which summarizes the properties of the native space \mathcal{X}'_{LR} in relation to the regularization operator L_{R} and its generalized inverse L_{R}^\dagger . This result is a direct corollary (dual counterpart) of Theorem 9, whose proof is given in Section 5.2.

Theorem 6 (Properties of the native space \mathcal{X}'_{LR}) *Let L be an isotropic LSI operator that is spline-admissible with a polynomial null space \mathcal{P} (possibly trivial) of degree n_0 . Then, the operators $L_{\text{R}} = K_{\text{rad}}RL : \mathcal{X}'_{\text{LR}} \rightarrow \mathcal{X}'_{\text{Rad}}$ and $L_{\text{R}}^\dagger = (\text{Id} - \text{Proj}_{\mathcal{P}})L^{-1}R^* : \mathcal{X}'_{\text{Rad}} \rightarrow L_{\infty, -n_0}(\mathbb{R}^d)$ (the adjoint of $L_{\text{R}}^{*\dagger}$ in Theorem 9) are continuous and have the properties*

$$\forall w \in \mathcal{X}'_{\text{Rad}} : L_{\text{R}}L_{\text{R}}^\dagger\{w\} = w \quad (63)$$

$$\forall p_0 \in \mathcal{P} : L_{\text{R}}\{p_0\} = 0 \quad (64)$$

$$\forall f \in \mathcal{X}'_{\text{LR}}(\mathbb{R}^d) : L_{\text{R}}^\dagger L_{\text{R}}\{f\} = (\text{Id} - \text{Proj}_{\mathcal{P}})\{f\} = \text{Proj}_{\mathcal{U}'}\{f\}, \quad (65)$$

where $\mathcal{X}'_{\text{LR}}(\mathbb{R}^d) \triangleq L_{\text{R}}^\dagger(\mathcal{X}'_{\text{Rad}}) \oplus \mathcal{P}$ is equipped with the composite norm $\|f\|_{\mathcal{X}'_{\text{LR}}} = \|L_{\text{R}}\{f\}\|_{\mathcal{X}'_{\text{Rad}}} + \|\text{Proj}_{\mathcal{P}}\{f\}\|_{\mathcal{P}}$. The space \mathcal{X}'_{LR} is complete and isomorphic to $\mathcal{X}'_{\text{Rad}} \times \mathcal{P}$ with $f = L_{\text{R}}^\dagger\{w\} + p_0 \mapsto (w, p_0) = (L_{\text{R}}\{f\}, \text{Proj}_{\mathcal{P}}\{f\})$. Moreover, $\mathcal{S}(\mathbb{R}^d) \hookrightarrow \mathcal{X}'_{\text{LR}}(\mathbb{R}^d) \hookrightarrow L_{\infty, -n_0}(\mathbb{R}^d) \xrightarrow{d} \mathcal{S}'(\mathbb{R}^d)$.

To explain how Theorem 6 relates to universality, let us first consider the case $\mathcal{X}' = L_2$ for which we have just shown that $\mathcal{X}'_{\text{LR}} = \mathcal{H}$ is a RKHS whose (semi-)reproducing kernel is $\rho_{\text{iso}} = (L^*KL)^{-1}\{\delta\}$. From the general properties of RKHS [2, 63], we know that \mathcal{H} (as a set) can be specified as $\mathcal{H} = \overline{\text{span}\{\rho_{\text{iso}}(\cdot - \mathbf{y})\}_{\mathbf{y} \in \mathbb{R}^d}} + \mathcal{P}$, which tells us that the class of RBF estimators of the form given by (53) is dense in \mathcal{H} . This means that such estimators can yield an approximation of any $f \in \mathcal{H}$ to an arbitrary degree of precision. In particular, this applies to any $f \in \mathcal{S}(\mathbb{R}^d)$, due to the inclusion $\mathcal{S}(\mathbb{R}^d) \subset \mathcal{X}'_{\text{LR}}(\mathbb{R}^d)$, as guaranteed by Theorem 6. Now, the key to universal approximation is that $\mathcal{S}(\mathbb{R}^d)$ is dense in most of the classical function spaces

[51], in particular, $C_0(\mathbb{R}^d)$. We then readily deduce that (53), for M sufficiently large and a suitable choice of the \mathbf{x}_m , has the ability to reproduce any continuous function $f : \mathbb{R}^d \rightarrow \mathbb{R}$. This deduction, of course, is consistent with the classical theory of kernel estimators: Our admissibility conditions for \hat{L}_{rad} in Definition 2 (resp., in Definition 3) ensure that the function $\rho_{\text{iso}} : \mathbb{R}^d \rightarrow \mathbb{R}$ is strictly positive-definite (resp., strictly conditionally positive-definite), which is the standard criterion for universality [31, 32, 63].

The same kind of density argument can also be made for $\mathcal{X}' = \mathcal{M}$. The relevant basis elements there are the atomic Radon-compatible Dirac measures $\delta_{\text{Rad},(t_k, \boldsymbol{\xi}_k)} \in \mathcal{M}_{\text{Rad}}$ with $(t_k, \boldsymbol{\xi}_k) \in \mathbb{R} \times \mathbb{S}^{d-1}$. These get mapped into $e_k = L_{\mathbb{R}}^{\dagger} \{ \delta_{\text{Rad},(t_k, \boldsymbol{\xi}_k)} \} \in \mathcal{U}'$, which are essentially ridges (up to some polynomial) characterized by $e_k = \rho_{\text{rad}}(\boldsymbol{\xi}_k^{\text{T}} \cdot -t_k) - p_{0,k}$ with $p_{0,k} = \text{Proj}_{\mathcal{P}} \{ \rho_{\text{rad}}(\boldsymbol{\xi}_k^{\text{T}} \cdot -t_k) \}$. Thus, by setting $w \approx \sum_k w_k \delta_{\text{Rad},(t_k, \boldsymbol{\xi}_k)}$, we can interpret the generative formula $f = L_{\mathbb{R}}^{\dagger} \{ w \} + p_0$ in Theorem 6 as a linear superposition of ridges plus a global polynomial trend, which is compatible with the form of the estimator in (55). We then invoke the property that $\mathcal{S}(\mathbb{R}^d) \subset \mathcal{M}_{L_{\mathbb{R}}}(\mathbb{R}^d)$, which implies that any continuous function can be approximated as closely as desired by a member of $\mathcal{M}_{L_{\mathbb{R}}}(\mathbb{R}^d)$. As it turns out, the latter is representable by a superposition of ridges plus a polynomial of degree n_0 . We emphasize that the presence of the polynomial term—the guarantor of stability for Theorem 8—is essential to counterbalance the growth of the individual atoms at infinity. This is an important aspect where our analysis and conclusions deviate from those of [54].

4.4 Regularization Operators for Antisymmetric Activations

While the framework that has been discussed so far is very powerful, it has one shortcoming: it yields canonical activation functions ρ_{rad} that are necessarily symmetric. In some cases such as $L_{\mathbb{R}} = K_{\text{rad}}\mathbb{R}(-\Delta)$, these can be converted into one-sided functions such as the ReLU by doctoring the polynomial term. That said, the original scheme that is described by Theorem 5 does not allow for sigmoids, which are frequently used for classification [8]. This is the reason why we now introduce a variant that systematically produces anti-symmetric activations, including sigmoids for $n_0 = 0$.

The idea is to substitute the (symmetric) filtering operator K_{rad} by its anti-symmetric counterpart \tilde{K}_{rad} , which includes an additional Hilbert transform. Specifically, \tilde{K}_{rad} is the hyper-spherical radial filter whose frequency response is $\hat{K}_{\text{rad}}(\omega) = -i \text{sign}(\omega) c_p |\omega|^{-d-1}$ and whose adjoint is $\tilde{K}_{\text{rad}}^* = -\tilde{K}_{\text{rad}}$. Since the action of the (radial) Hilbert transform $\tilde{H}_{\text{rad}} : \phi(\cdot, \boldsymbol{\xi}) \mapsto \phi \otimes 1/(\pi t)$ is well-defined on $\mathcal{S}(\mathbb{R} \times \mathbb{S}^{d-1})$ with $\tilde{H}_{\text{rad}}^* = -\tilde{H}_{\text{rad}} = \tilde{H}_{\text{rad}}^{-1}$, we have the identity $\mathbb{R}^* K_{\text{rad}} \mathbb{R} = \mathbb{R}^* K_{\text{rad}} \tilde{H}_{\text{rad}}^* \tilde{H}_{\text{rad}} \mathbb{R} = \mathbb{R}^* \tilde{K}_{\text{rad}}^* \tilde{\mathbb{R}} = \text{Id}$ with $\tilde{\mathbb{R}} \triangleq \tilde{H}_{\text{rad}} \mathbb{R}$. Accordingly, we can essentially replicate the whole mechanism of construction of spaces in Section 3.3 by substituting \mathcal{S}_{Rad} by $\tilde{\mathcal{S}}_{\text{Rad}} = \tilde{\mathbb{R}}(\mathcal{S}(\mathbb{R} \times \mathbb{S}^{d-1})) = \tilde{H}_{\text{rad}}(\mathcal{S}_{\text{Rad}})$, which is a space of odd functions that are smooth (C^∞) and included in $L_p(\mathbb{R} \times \mathbb{S}^{d-1})$

for all $p \geq 1$. While the members of $\tilde{\mathcal{S}}_{\text{Rad}}$ do not necessarily decay rapidly, the mapping $\mathbb{R}^* \tilde{\mathbb{K}}_{\text{rad}}^* : \tilde{\mathcal{S}}_{\text{Rad}} \rightarrow \mathcal{S}(\mathbb{R}^d)$ is still guaranteed to be an isomorphism (see [43, Theorem 3.3.1, p. 83] where similar arguments are used). Under the assumption that the hyper-spherical norm $\|\cdot\|_{\mathcal{X}}$ is continuous on $\tilde{\mathcal{S}}_{\text{Rad}}$, we can readily adapt the proof of [57, Theorem 8] to establish the following.

Proposition 3 (Odd Radon-compatible Banach spaces) *Consider the Banach space $\tilde{\mathcal{X}}_{\text{Rad}} = (\tilde{\mathcal{S}}_{\text{Rad}}, \|\cdot\|_{\mathcal{X}})$ of odd hyper-spherical functions. Then, the following holds.*

1. The map $\mathbb{R}^* \tilde{\mathbb{K}}_{\text{rad}}^* : \tilde{\mathcal{X}}_{\text{Rad}} \rightarrow \tilde{\mathcal{Y}} = \mathbb{R}^* \tilde{\mathbb{K}}_{\text{rad}}^*(\tilde{\mathcal{X}}_{\text{Rad}})$ is an isometric bijection, with $\tilde{\mathbb{R}} \mathbb{R}^* \tilde{\mathbb{K}}_{\text{rad}}^* = \text{Id}$ on $\tilde{\mathcal{X}}_{\text{Rad}}$.
2. The map $\tilde{\mathbb{R}}^* : \tilde{\mathcal{X}}'_{\text{Rad}} \rightarrow \tilde{\mathcal{Y}}' = \tilde{\mathbb{R}}^*(\tilde{\mathcal{X}}'_{\text{Rad}})$ is an isometric bijection, with $\tilde{\mathbb{K}}_{\text{rad}} \tilde{\mathbb{R}}^* = \text{Id}$ on $\tilde{\mathcal{X}}'_{\text{Rad}}$.

The underlying definition of the ‘‘oddified’’ backprojection operator $\tilde{\mathbb{R}}^* : \tilde{\mathcal{X}}'_{\text{Rad}} \rightarrow \tilde{\mathcal{Y}}'$ for $g \in \tilde{\mathcal{X}}'_{\text{Rad}}$ is

$$\langle \tilde{\mathbb{R}}^* \{g\}, \varphi \rangle = \langle g, \tilde{\mathbb{R}} \{\varphi\} \rangle_{\text{Rad}} = \langle \tilde{\mathbb{H}}_{\text{rad}}^* \{g\}, \mathbb{R} \{\varphi\} \rangle_{\text{Rad}} \quad (66)$$

for all $\varphi \in \tilde{\mathcal{Y}}$ or, equivalently, $\varphi \in \mathcal{S}(\mathbb{R}^d)$ since $\mathcal{S}(\mathbb{R}^d)$ is dense in \mathcal{Y} by construction. Likewise, by using the property that the Hilbert transform is a homeomorphism from $\mathcal{S}_{\text{Liz,even}}$ onto $\mathcal{S}_{\text{Liz,odd}}$ [47] with the underlying Lizorkin spaces being included in \mathcal{S}_{Rad} and $\tilde{\mathcal{S}}_{\text{Rad}} \subset L_{p,\text{odd}}$, respectively, we can adapt the proof of Lemma 2 to show that $\tilde{L}_{q,\text{Rad}} = L_{q,\text{odd}}$ for $q \in (1, \infty)$, $\tilde{C}_{0,\text{Rad}} = C_{0,\text{odd}}$ for $p = \infty$, and $\tilde{\mathcal{M}}_{\text{Rad}} = (\tilde{C}_{0,\text{Rad}})' = \mathcal{M}_{\text{odd}}$.

The bottom line is that the whole argumentation, including the critical Fourier-based proofs of Section 5, applies in this modified setting as well. Accordingly, all theorems that mention the regularization operator $L_{\text{R}} = \mathbb{K}_{\text{rad}} \mathbb{R} L$ and the radial profile $\rho_{\text{rad}} = \mathcal{F}^{-1}\{1/\tilde{L}_{\text{rad}}\}$ are also valid for the odd setting where these quantities are substituted by

$$\tilde{L}_{\text{R}} \triangleq \tilde{\mathbb{K}}_{\text{rad}} \mathbb{R} L = \tilde{\mathbb{K}}_{\text{rad}} L_{\text{rad}} \mathbb{R}, \quad (67)$$

$$\tilde{\rho}_{\text{rad}}(t) \triangleq \mathcal{F}^{-1} \left\{ \frac{i \text{sign}(\omega)}{\tilde{L}_{\text{rad}}(\omega)} \right\} (t), \quad (68)$$

where (68) directly follows from (66) and Theorem 4. The conditions for admissibility in Definitions 2 and 3, which are all Fourier-based, remain the same, while the adjusted activation $\tilde{\rho}_{\text{rad}}$ (the 1D Hilbert transform of ρ_{rad}) is real-valued and anti-symmetric because the original Fourier profile $\tilde{L}_{\text{rad}} : \mathbb{R} \rightarrow \mathbb{R}$ is symmetric.

We note that our admissibility condition with $\gamma_0 = 1$ for odd variant ($n_0 = 0$) is compatible with the condition used by Barron to prove the universality of neural networks with sigmoidal activations [5]. It is also worth mentioning that the substitution of L_{R} by \tilde{L}_{R} has no incidence on the form of the RBF in (53) because of the unitary nature of the Hilbert transform.

$\rho_{\text{rad}}(t)$	$\tilde{\rho}_{\text{rad}}(t)$	\widehat{L}_{rad} or $\widehat{\tilde{L}}_{\text{rad}}$	n_0 (null space)
Exponential			
$e^{- t }$	N/A	$1 + \omega ^2$	-1 (trivial)
Classical sigmoids			
N/A	$\frac{\tanh(\frac{t}{2})}{2} = \frac{1}{2} + \frac{1}{1+e^{-t}}$	$\frac{\sinh(\pi\omega)}{\pi}$	0 (bias)
N/A	$\frac{\arctan(t)}{\pi}$	$\omega e^{ \omega }$	0
Ridge splines (of degree $n \in \mathbb{N}$)			
$\frac{1}{2} t $ (or ReLU)	$t \log t $	$ \omega ^2$	1 (affine)
$\propto t^{2n} \log t $	$\frac{\text{sign}(t) t ^{2n}}{(2n)!}$	$ \omega ^{2n+1}$	$2n \geq 2$ (even)
$\frac{1}{2} \frac{ t ^{2n+1}}{(2n+1)!}$	$\propto t^{2n+1} \log t $	$ \omega ^{2n+2}$	$2n + 1 \geq 1$ (odd)
Fractional splines (degree $\alpha \in \mathbb{R}^+ \setminus \mathbb{N}$)			
$\frac{ t ^\alpha \sin(\frac{\alpha\pi}{2})}{\pi\Gamma(\alpha)}$	$\frac{\text{sign}(t) t ^\alpha \cos(\frac{\alpha\pi}{2})}{\pi\Gamma(\alpha)}$	$ \omega ^{\alpha+1}$	$\lceil \alpha \rceil$

Table 1 Examples of admissible symmetric and anti-symmetric activation functions with their corresponding regularization operator. The anti-symmetric activation is given by (68) and requires the use of the modified filter \tilde{K}_{rad} in the statement of Theorem 5.

4.5 Specific Configurations

The proposed framework encompasses a wide variety of kernels and activation functions, with minimal restrictions. For instance, one can start with any strictly positive-definite function $\rho_{\text{rad},0}$ whose Fourier transform is strictly positive, and construct some higher-order variants by (fractional) integration. The variants are such that $\rho_{\text{rad},\gamma_0}(t) = \mathcal{F}^{-1}\left\{\frac{\widehat{\rho}_{\text{rad},0}(\omega)}{|\omega|^{\gamma_0}}\right\}(t)$ with suitable $\gamma_0 > 0$ and are guaranteed to meet the requirements in Definition 3. The simplest scenario is to set $\rho_{\text{rad},0} = \delta$, which maps into a Laplacian-type regularization.

In Table 1, we provide examples of admissible operators together with their corresponding symmetric and anti-symmetric activations. It is noteworthy that the two most popular sigmoids (tanh and arctan) are part of the framework. We can determine the explicit frequency response of their regularization filter and show that it is first-order-admissible with a null space that consists of the polynomials of degree $n_0 = 0$ (the constants). The symmetric spline activations of odd degree $m - 1$ and the anti-symmetric ones of even degree are also known: they coincide with the ridge splines of Parhi and Novak, which are tied to the Radon-domain operator $L_{\text{rad}} = \frac{\partial^m}{\partial t^m}$ [37]. With the present formulation, we can seamlessly extend this family to fractional orders, in direct analogy with [59], by setting $L = (-\Delta)^{\frac{\alpha+1}{2}}$.

5 Supporting Mathematical Results

To answer the fundamental issue of the existence of a solution in Theorem 5, we need to 1) prove that the “predual” space $\mathcal{X}_L(\mathbb{R}^d)$ is a proper Banach space (Theorem 9); and 2) establish the weak* continuity of the sampling functionals $\delta(\cdot - \mathbf{x}_k)$. As we shall see, both aspects largely rest upon the functional characterization of the Schwartz kernel of the pseudoinverse operator $L_R^{*\dagger}$ provided in Theorem 7.

5.1 Kernel and Stability of Generalized Inverse Operators

Let $\nu : f \mapsto \langle \nu, f \rangle$ be a linear functional that is acting on some Banach space \mathcal{X}' . We recall that ν is weak*-continuous if and only if $\nu \in \mathcal{X}$, where \mathcal{X} is the predual of \mathcal{X}' [44]. The condition that \mathcal{X}' is reflexive (i.e., $\mathcal{X}'' = \mathcal{X}'$) is equivalent to the continuity of ν on \mathcal{X} , which is the standard condition for analysis. However, when \mathcal{X}' is not reflexive (e.g., $\mathcal{X}' = \mathcal{M} = (C_0)'$), the constraint of weak* continuity is stronger than continuity, contrary to what could be suggested by the qualifier “weak.” In that case, the predual space $\mathcal{X} \subseteq \mathcal{X}''$, which is continuously embedded in \mathcal{X}'' , turns out to be smaller than \mathcal{X}'' . Therefore, in order to establish the weak* continuity of the Dirac functionals $\delta(\cdot - \mathbf{x}_0)$ for the scenarios in Theorem 5, we need to show that $\delta(\cdot - \mathbf{x}_0) \in \mathcal{X}_{L_R}$, which can be reduced to proving that $L_R^{*\dagger}\{\delta(\cdot - \mathbf{x}_0)\} \in \mathcal{X}_{\text{Rad}}$.

Theorem 7 (Properties of the impulse response of $L_R^{*\dagger} = \text{RL}_{\mathcal{P}}^{-1*}$) *Let L be an isotropic operator such that $\widehat{L}(\boldsymbol{\omega}) = \widehat{L}_{\text{rad}}(\|\boldsymbol{\omega}\|)$ where $\widehat{L}_{\text{rad}} : \mathbb{R} \rightarrow \mathbb{R}$ is a continuous function and $\rho_{\text{rad}}(t) = \mathcal{F}^{-1}\{1/\widehat{L}_{\text{rad}}\}(t)$. We consider two cases:*

1. *Trivial null space: If L satisfies the admissibility conditions in Definition 2, then $L_R^{*\dagger} = \text{RL}^{-1*}$ and*

$$\nu_{\mathbf{x}_0}(t, \boldsymbol{\xi}) = \text{RL}^{-1*}\{\delta(\cdot - \mathbf{x}_0)\}(t, \boldsymbol{\xi}) = \rho_{\text{rad}}(t - \boldsymbol{\xi}^\top \mathbf{x}_0) \quad (69)$$

with $\mathbf{x}_0 \in \mathbb{R}^d$ and $(t, \boldsymbol{\xi}) \in \mathbb{R} \times \mathbb{S}^{d-1}$. Moreover, $\nu_{\mathbf{x}_0} \in \mathcal{X}_{\text{Rad}}$ for $\mathcal{X} = C_0$ as well as $\mathcal{X} = L_q$ with $q \in [1, \infty]$.

2. *Nontrivial null space: If L satisfies the admissibility conditions in Definition 3 with a polynomial null space of degree n_0 , then*

$$\begin{aligned} \nu_{\mathbf{x}_0}(t, \boldsymbol{\xi}) &= L_R^{*\dagger}\{\delta(\cdot - \mathbf{x}_0)\}(t, \boldsymbol{\xi}) \\ &= \rho_{\text{rad}}(t - \boldsymbol{\xi}^\top \mathbf{x}_0) - \sum_{n=0}^{n_0} \frac{(-\boldsymbol{\xi}^\top \mathbf{x}_0)^n}{n!} (\kappa_{\text{rad}} * \partial^n \rho_{\text{rad}})(t) \end{aligned} \quad (70)$$

with $\nu_{\mathbf{x}_0} \in \mathcal{X}_{\text{Rad}}$ for the same spaces as in Item 1, but with $q \in [2, \infty]$. Moreover,

$$\sup_{(\mathbf{x}_0, \boldsymbol{\xi}) \in \mathbb{R}^d \times \mathbb{S}^{d-1}} (1 + |\boldsymbol{\xi}^\top \mathbf{x}_0|)^{-n_0} \|\nu_{\mathbf{x}_0}(\cdot, \boldsymbol{\xi})\|_{L_q(\mathbb{R})} < \infty \quad (71)$$

for any $q \in [2, \infty]$.

Proof To show that $\nu_{\mathbf{x}_0} \in \mathcal{X}_{\text{Rad}}$, it is sufficient to prove that $\nu_{\mathbf{x}_0} \in \mathcal{X}(\mathbb{R}^d \times \mathbb{S}^{d-1})$ since $\nu_{\mathbf{x}_0}$ is in the range of the Radon transform by construction.

When the null space of L is trivial, L has a stable convolution inverse so that it suffices to show that $\nu_{\mathbf{x}_0} \in L_q(\mathbb{R} \times \mathbb{S}^{d-1}) \cap C_0(\mathbb{R} \times \mathbb{S}^{d-1})$. To that end, we formally identify the isotropic distribution $\rho_{\text{iso}} = L^{-1*}\{\delta\} = L^{-1}\{\delta\} = \mathcal{F}^{-1}\{1/\widehat{L}_{\text{rad}}(\|\boldsymbol{\omega}\|)\}$ and apply Proposition 1, which yields

$$\nu_{\mathbf{x}_0}(t, \boldsymbol{\xi}) = \mathbb{R}\{\rho_{\text{iso}}(\cdot - \mathbf{x}_0)\}(t, \boldsymbol{\xi}) = \rho_{\text{rad}}(t - \boldsymbol{\xi}^\top \mathbf{x}_0).$$

Due to our assumptions, this function is such that $\|\nu_{\mathbf{x}_0}(\cdot, \boldsymbol{\xi})\|_{L_1} = \|\rho_{\text{rad}}\|_{L_1} < \infty$ for any fixed $\boldsymbol{\xi} \in \mathbb{S}^{d-1}$. Moreover, since $1/\widehat{L}_{\text{rad}} \in L_1(\mathbb{R})$, ρ_{rad} is bounded, continuous, and vanishing at infinity (by the Riemann-Lebesgue lemma), which gives $\rho_{\text{rad}}(\cdot - \boldsymbol{\xi}^\top \mathbf{x}_0) \in C_0(\mathbb{R})$ for any $\mathbf{x}_0 \in \mathbb{R}^d$ and $\boldsymbol{\xi} \in \mathbb{S}^{d-1}$. Since $\rho_{\text{rad}} \in L_\infty(\mathbb{R}) \cap L_1(\mathbb{R})$, we readily deduce that $\rho_{\text{rad}}(\cdot - \boldsymbol{\xi}^\top \mathbf{x}_0) \in L_q(\mathbb{R})$ for all intermediate $q \geq 1$, which then also yields $\rho_{\text{rad}}(t - \boldsymbol{\xi}^\top \mathbf{x}_0) \in L_q(\mathbb{R} \times \mathbb{S}^{d-1})$ because the surface of the unit sphere \mathbb{S}^{d-1} is bounded. Finally, since $\rho_{\text{rad}} : \mathbb{R} \rightarrow \mathbb{R}$ is continuous and vanishing at infinity, $\nu_{\mathbf{x}_0}(t, \boldsymbol{\xi})$ is jointly continuous in $(t, \boldsymbol{\xi})$ and vanishing at $t \rightarrow \pm\infty$ for all $\boldsymbol{\xi} \in \mathbb{S}^{d-1}$, which implies that $\nu_{\mathbf{x}_0} \in C_0(\mathbb{R} \times \mathbb{S}^{d-1})$.

For the more difficult case of a nontrivial null space, we invoke the Fourier-slice theorem to evaluate the 1D Fourier transform of $\nu_{\mathbf{x}_0}(\cdot, \boldsymbol{\xi})$ with $\boldsymbol{\xi}$ fixed as

$$\begin{aligned} \hat{\nu}_{\mathbf{x}_0}(\omega, \boldsymbol{\xi}) &= \frac{\mathcal{F}\{\delta(\cdot - \mathbf{x}_0) - \sum_{|\mathbf{k}| \leq n_0} \langle m_{\mathbf{k}}, \delta(\cdot - \mathbf{x}_0) \rangle m_{\mathbf{k}}^*\}(\omega \boldsymbol{\xi})}{\widehat{L}_{\text{rad}}(\omega)} \\ &= \frac{e^{-i\omega \boldsymbol{\xi}^\top \mathbf{x}_0} - \sum_{|\mathbf{k}| \leq n_0} \frac{\mathbf{x}_0^{\mathbf{k}}}{\mathbf{k}!} \widehat{m}_{\mathbf{k}}^*(\omega \boldsymbol{\xi})}{\widehat{L}_{\text{rad}}(\omega)} \\ &= \frac{e^{-i\omega \boldsymbol{\xi}^\top \mathbf{x}_0} - \sum_{|\mathbf{k}| \leq n_0} \frac{\mathbf{x}_0^{\mathbf{k}}}{\mathbf{k}!} (-i\omega \boldsymbol{\xi})^{\mathbf{k}} \widehat{\kappa}_{\text{rad}}(\omega)}{\widehat{L}_{\text{rad}}(\omega)} \\ &= \frac{e^{-i\omega \boldsymbol{\xi}^\top \mathbf{x}_0} - \sum_{n=0}^{n_0} \frac{(-\boldsymbol{\xi}^\top \mathbf{x}_0)^n}{n!} (i\omega)^n \widehat{\kappa}_{\text{rad}}(\omega)}{\widehat{L}_{\text{rad}}(\omega)}, \end{aligned} \quad (72)$$

where the simplification of the summation over \mathbf{k} results from the multinomial expansion $(-i\omega \boldsymbol{\xi}^\top \mathbf{x}_0)^n = (y_1 + \dots + y_d)^n = \sum_{|\mathbf{k}|=n} \frac{n!}{\mathbf{k}!} \mathbf{y}^{\mathbf{k}}$ with $\mathbf{y} = (-i\omega \xi_i x_{0,i})_{i=1}^d$. The delicate point here is that (72) is potentially singular because it has a pole of multiplicity γ_0 at $\omega = 0$. Fortunately, the condition $\gamma_0 \leq n_0 + 1$ ensures that there is a proper pole-zero cancellation: By recalling that $\widehat{\kappa}_{\text{rad}}(\omega) = 1$ for $\omega \in \Omega_0 = [-R_0, R_0]$ with $R_0 = \frac{1}{2}$ and setting $t_0 = \boldsymbol{\xi}^\top \mathbf{x}_0$, we identify the numerator as the remainder of the Maclaurin series of $e^{-i\omega t_0}$, which is bounded by

$$\begin{aligned} \left| e^{-i\omega t_0} - \sum_{n=0}^N \frac{(-it_0)^n}{n!} \omega^n \right| &\leq \sup_{\omega \in \mathbb{R}} |(-it_0)^{N+1} e^{-i\omega t_0}| \frac{|\omega|^{N+1}}{(N+1)!} \\ &\leq \frac{|t_0|^{N+1}}{(N+1)!} |\omega|^{N+1}. \end{aligned} \quad (73)$$




Fig. 1 Functions $\omega \mapsto |r_N(\omega)|$ for $N = 1, \dots, 10$. The dilated plot on the right includes the upper bound specified by (77) as an overlay.

This then yields the estimate

$$|\hat{\nu}_{\mathbf{x}_0}(\omega, \boldsymbol{\xi})| \leq \frac{|\boldsymbol{\xi}^\top \mathbf{x}_0|^{n_0+1}}{C_0(n_0+1)!} |\omega|^\epsilon \quad \text{as } \omega \rightarrow 0$$

$$\text{with } \epsilon = n_0 + 1 - \gamma_0 = \begin{cases} 0, & \gamma_0 \in \mathbb{N} \\ 1 - (\gamma_0 - \lfloor \gamma_0 \rfloor) \in (0, 1), & \text{otherwise.} \end{cases} \quad (74)$$

Since the denominator \hat{L}_{rad} in (72) is continuous and non-vanishing away from the origin, $\hat{\nu}_{\mathbf{x}_0}(\cdot, \boldsymbol{\xi})$ is bounded on Ω_0 , and, by extension, on any compact subset of \mathbb{R} . Moreover, since $|e^{-i\omega \boldsymbol{\xi}^\top \mathbf{x}_0}| = 1$ and $\hat{\kappa}_{\text{rad}}$ is rapidly decreasing, there exists a constant C such that $|\hat{\nu}_{\mathbf{x}_0}(\omega, \boldsymbol{\xi})| \leq C|\omega|^{-\gamma_1}$ for $|\omega| > R_1$. This leads to several conclusions. First, if $\gamma_1 > 1$, then $\hat{\nu}_{\mathbf{x}_0}(\cdot, \boldsymbol{\xi}) \in L_1(\mathbb{R})$ so that $\nu_{\mathbf{x}_0} \in C_{0,\text{Rad}}(\mathbb{R} \times \mathbb{S}^{d-1})$ by the same argument as in the nonsingular case. Second, if $\gamma_1 > \alpha + \frac{1}{2}$, then $\nu_{\mathbf{x}_0}(\cdot, \boldsymbol{\xi}) \in W_2^\alpha(\mathbb{R})$, where $W_2^\alpha(\mathbb{R}) = \{f : \int_{\mathbb{R}} (1 + |\omega|^2)^\alpha |\hat{f}(\omega)|^2 d\omega < \infty\}$ is the Sobolev space of functions with finite-energy derivatives up to order α . Therefore, $\nu_{\mathbf{x}_0} \in L_{q,\text{Rad}}(\mathbb{R} \times \mathbb{S}^{d-1})$ for all $q \in [2, \infty]$ provided that $\gamma_1 > 1$. The explicit Radon-domain formula (70) with $\kappa_{\text{rad}}(t) = \mathcal{F}^{-1}\{\hat{\kappa}_{\text{rad}}(\omega)\}(t)$ is obtained by taking the 1D inverse transform of (72).

To refine our characterization of $\hat{\nu}_{\mathbf{x}_0}$, we introduce the function

$$r_N(\omega) = \frac{e^{-i\omega} - \sum_{n=0}^N \frac{(-i\omega)^n}{n!}}{\frac{(i\omega)^N}{N!}} \quad (75)$$

whose modulus is plotted in Figure 1. Some of the remarkable properties of r_N are

$$r_N(\omega) = \frac{-i\omega}{N+1} \quad \text{as } \omega \rightarrow 0 \quad (76)$$

$$\forall \omega \in \mathbb{R} : |r_N(\omega)| \leq \min(|\omega|/2, 1.27) \quad (77)$$

$$\lim_{\omega \rightarrow \infty} |r_N(\omega)| = 1, \quad (78)$$

with the global bound (77) being overlaid on the graph to demonstrate its sharpness. This function will allow us to control the behavior of

$$\begin{aligned}\hat{\nu}_{\mathbf{x}_0}(\omega, \boldsymbol{\xi}) &= \frac{e^{-it_0\omega} - \sum_{n \leq n_0} \frac{(-it_0\omega)^n}{n!} \kappa_{\text{rad}}(\omega)}{\widehat{L}_{\text{rad}}(\omega)} \\ &= \frac{\widehat{\kappa}_{\text{rad}}(\omega) r_{n_0}(t_0\omega) \frac{(-it_0\omega)^{n_0}}{n_0!} + (1 - \widehat{\kappa}_{\text{rad}}(\omega)) e^{-it_0\omega}}{\widehat{L}_{\text{rad}}(\omega)}\end{aligned}$$

by splitting the frequency domain in three regions. First, for $\omega \in \Omega_0 = [0, R_0]$, where $\widehat{\kappa}_{\text{rad}}(\omega) = 1$, we find that

$$|\hat{\nu}_{\mathbf{x}_0}(\omega, \boldsymbol{\xi})| = \left| \frac{r_{n_0}(t_0\omega) \frac{(-it_0\omega)^{n_0}}{n_0!}}{\widehat{L}_{\text{rad}}(\omega)} \right| \leq |t_0|^{n_0} \min(|\omega|, 2) \frac{2 \frac{|\omega|^{n_0}}{n_0!}}{|\widehat{L}_{\text{rad}}(\omega)|}. \quad (79)$$

For the transition region $\omega \in \Omega_{01} = [R_0, \tilde{R}_1]$ with $\tilde{R}_1 = \max(2R_0, R_1)$, where $0 \leq \widehat{\kappa}_{\text{rad}}(\omega) \leq 1$, we bound the numerator by its maximum, which yields

$$|\hat{\nu}_{\mathbf{x}_0}(\omega, \boldsymbol{\xi})| \leq \frac{2 \frac{|t_0 \tilde{R}_1|^{n_0}}{n_0!} + 1}{|\widehat{L}_{\text{rad}}(\omega)|}. \quad (80)$$

Finally, for $\omega \in \Omega_1 = [\tilde{R}_1, \infty]$, where $\widehat{\kappa}_{\text{rad}}(\omega) = 0$, we get the expected tail behavior

$$|\hat{\nu}_{\mathbf{x}_0}(\omega, \boldsymbol{\xi})| \leq \frac{1}{|\widehat{L}_{\text{rad}}(\omega)|} \leq \frac{1}{C_1 |\omega|^{\gamma_1}}. \quad (81)$$

Based on those bounds with $t_0 = 1$, we define the auxiliary function

$$u(\omega) = \begin{cases} \min(|\omega|, 2) \frac{2 \frac{|\omega|^{n_0}}{n_0!}}{|\widehat{L}_{\text{rad}}(\omega)|}, & |\omega| < R_0 \\ \frac{2 \frac{\tilde{R}_1^{n_0}}{n_0!} + 1}{|\widehat{L}_{\text{rad}}(\omega)|}, & R_0 \leq |\omega| \leq \tilde{R}_1 \\ \frac{1}{|\widehat{L}_{\text{rad}}(\omega)|}, & \tilde{R}_1 < |\omega| \end{cases} \quad (82)$$

which, by construction, is such that $\|u\|_{L_p} < \infty$ for any $p \geq 1$. We can now use (79), (80), and (81) to bound the norm of $\hat{\nu}_{\mathbf{x}_0}(\cdot, \boldsymbol{\xi})$ by distinguishing between two cases. For $t_0 \leq 1$, we have $\|\hat{\nu}_{\mathbf{x}_0}(\cdot, \boldsymbol{\xi})\|_{L_p} \leq \|u\|_{L_p}$, while, for $t_0 \geq 1$, we get $\|\hat{\nu}_{\mathbf{x}_0}(\cdot, \boldsymbol{\xi})\|_{L_p} \leq |t_0|^{n_0} \|u\|_{L_p}$. By combining these two inequalities, we obtain the universal norm estimate

$$\|\hat{\nu}_{\mathbf{x}_0}(\cdot, \boldsymbol{\xi})\|_{L_p(\mathbb{R})} \leq (1 + |\boldsymbol{\xi}^\top \mathbf{x}_0|)^{n_0} \|u\|_{L_p} < +\infty, \quad (83)$$

which holds for all $(\mathbf{x}_0, \boldsymbol{\xi}) \in \mathbb{R}^d \times \mathbb{S}^{d-1}$ and $p \geq 1$. Finally, we invoke the boundedness of the (inverse) Fourier transform $\mathcal{F}^{-1} : L_p(\mathbb{R}) \rightarrow L_q(\mathbb{R})$ for $p \in [1, 2]$ to obtain the generic bound

$$\sup_{(\mathbf{x}_0, \boldsymbol{\xi}) \in \mathbb{R}^d \times \mathbb{S}^{d-1}} (1 + |\boldsymbol{\xi}^\top \mathbf{x}_0|)^{-n_0} \|\nu_{\mathbf{x}_0}(\cdot, \boldsymbol{\xi})\|_{L_q(\mathbb{R})} < \infty, \quad (84)$$

which is the desired result. \square

As complement to the proof of Theorem 7, we make the following remarks.

1. The function $\omega \mapsto \hat{\nu}_{\mathbf{x}_0}(\omega, \boldsymbol{\xi})$ in (72) is ϵ -Hölder continuous around the origin for $\gamma_0 \notin 2\mathbb{N}$ and as smooth as \hat{L} when $\gamma_0 = 2n$. For instance, the case where $L = (-\Delta)^n$ is a non-fractional iterate of the Laplacian corresponds to $\hat{L}_{\text{rad}}(\omega) = \omega^{2n}$ and $\hat{\nu}_{\mathbf{x}_0}(\cdot, \boldsymbol{\xi}) \in C^\infty(\mathbb{R})$, which then translates into $t \mapsto \nu_{\mathbf{x}_0}(t, \boldsymbol{\xi})$ being rapidly decreasing, but with a limited order of differentiability controlled by $\gamma_0 = \gamma_1 = 2n$.

2. The proof can be readily adapted to characterize the partial derivatives of $\nu_{\mathbf{x}_0}$ by replacing $\delta(\cdot - \mathbf{x}_0)$ by $\partial^{\mathbf{n}}\delta(\cdot - \mathbf{x}_0)$ with $|\mathbf{n}| < n_0$. These distributions are such that $\langle \partial^{\mathbf{n}}\delta(\cdot - \mathbf{x}_0), m_{\mathbf{k}} \rangle = (-1)^{|\mathbf{n}|} \partial^{\mathbf{n}} m_{\mathbf{k}}(\mathbf{x}_0) = (-1)^{|\mathbf{n}|} m_{\mathbf{k}-\mathbf{n}}(\mathbf{x}_0)$, with the convention that $m_{\mathbf{k}-\mathbf{n}} = 0$ if $k_m < n_m$ for any $m \in \{1, \dots, d\}$.

3. The leading term in (70) is $\rho_{\text{rad}}(t - \tau_0)$ with $\tau_0 = \boldsymbol{\xi}^\top \mathbf{x}_0$, which, in the non-trivial scenario, typically grows as $O(|t|^{\gamma_0-1})$. Our analysis shows that the second correction term in (70), which is unbounded at infinity as well, essentially neutralizes this growth. It is tempting to call this a miraculous cancellation.

The bottom line is that, for any $(\mathbf{x}_0, \boldsymbol{\xi}_0) \in \mathbb{R}^d \times \mathbb{S}^{d-1}$, the function $t \mapsto \nu_{\mathbf{x}_0}(t, \boldsymbol{\xi}_0)$ is continuous, bounded with a maximum that is proportional to $|\boldsymbol{\xi}_0^\top \mathbf{x}_0|^{n_0}$ (see (84) with $p = \infty$), and vanishing at infinity. This is a remarkable property that also guarantees the boundedness of $L_{\mathbb{R}}^{*\dagger} : L_{1, n_0}(\mathbb{R}^d) \rightarrow \mathcal{X}_{\text{Rad}}$, which is not obvious a priori. The enabling ingredient there is (71), which ensures that the corresponding bounding constant in Theorem 8 is finite. Indeed, since $\|\mathbf{x}\| \geq |\boldsymbol{\xi}^\top \mathbf{x}|$ with equality if and only if $\boldsymbol{\xi}$ and \mathbf{x} are colinear, we have that

$$\begin{aligned} \|L_{\mathbb{R}}^{*\dagger}\| &\leq \sup_{\mathbf{x} \in \mathbb{R}^d, \boldsymbol{\xi} \in \mathbb{S}^{d-1}} (1 + \|\mathbf{x}\|)^{-n_0} \|\nu_{\mathbf{x}}(\cdot, \boldsymbol{\xi})\|_{\mathcal{X}(\mathbb{R})} \\ &\leq \sup_{\mathbf{x} \in \mathbb{R}^d, \boldsymbol{\xi} \in \mathbb{S}^{d-1}} (1 + |\boldsymbol{\xi}^\top \mathbf{x}|)^{-n_0} \|\nu_{\mathbf{x}}(\cdot, \boldsymbol{\xi})\|_{\mathcal{X}(\mathbb{R})} < \infty. \end{aligned}$$

Theorem 8 (Stability of Cartesian-to-Radon-domain mappings) *Let $h_{\mathbf{x}}(t, \boldsymbol{\xi}) = \mathbb{T}\{\delta(\cdot - \mathbf{x})\}(t, \boldsymbol{\xi})$ denote the generalized impulse response of the operator $\mathbb{T} : L_{1, \alpha}(\mathbb{R}^d) \rightarrow \mathcal{X}_{\text{Rad}}(\mathbb{R} \times \mathbb{S}^{d-1})$ and let*

$$C_{p, \alpha} = \sup_{\mathbf{x} \in \mathbb{R}^d, \boldsymbol{\xi} \in \mathbb{S}^{d-1}} (1 + \|\mathbf{x}\|)^{-\alpha} \|h_{\mathbf{x}}(\cdot, \boldsymbol{\xi})\|_{L_p(\mathbb{R})}. \quad (85)$$

1. $\mathcal{X} = C_0$: If $C_{\infty, \alpha} < \infty$, then $\mathbb{T} : L_{1, \alpha}(\mathbb{R}^d) \rightarrow C_{0, \text{Rad}}$ is bounded with $\|\mathbb{T}\| \leq C_{\infty, \alpha}$.
2. $\mathcal{X} = L_p$ with $p \in (1, \infty)$: If $C_{p, \alpha} < \infty$, then $\mathbb{T} : L_{1, \alpha}(\mathbb{R}^d) \rightarrow L_{p, \text{Rad}}$ is bounded with $\|\mathbb{T}\| \leq \frac{2\pi^{d/2}}{\Gamma(d/2)} C_{p, \alpha}$.

The same holds true for the adjoint $\mathbb{T}^* : \mathcal{X}'_{\text{Rad}} \rightarrow L_{\infty, -\alpha}(\mathbb{R}^d)$ with $\|\mathbb{T}^*\| = \|\mathbb{T}\|$.

Proof The function $((t, \boldsymbol{\xi}), \mathbf{x}) \mapsto h_{\mathbf{x}}(t, \boldsymbol{\xi})$ is the Schwartz kernel of the operator T , so that

$$\begin{aligned} |g(t, \boldsymbol{\xi})| &= |T\{f\}(t, \boldsymbol{\xi})| = \left| \int_{\mathbb{R}^d} h_{\mathbf{x}}(t, \boldsymbol{\xi}) f(\mathbf{x}) d\mathbf{x} \right| \\ &\leq \sup_{\mathbf{x} \in \mathbb{R}^d, \boldsymbol{\xi} \in \mathbb{S}^{d-1}} \left((1 + \|\mathbf{x}\|)^{-\alpha} \|h_{\mathbf{x}}(\cdot, \boldsymbol{\xi})\|_{L^\infty} \right) \int_{\mathbb{R}^d} (1 + \|\mathbf{x}\|)^\alpha |f(\mathbf{x})| d\mathbf{x}. \end{aligned}$$

Consequently,

$$\|g\|_{L^\infty} = \sup_{(t, \boldsymbol{\xi}) \in \mathbb{R} \times \mathbb{S}^{d-1}} |g(t, \boldsymbol{\xi})| = C_{\infty, \alpha} \|f\|_{L_{1, \alpha}},$$

which is the desired bound for $\mathcal{X} = C_0$.

To handle the reflexive case $\mathcal{X} = L_p$, we consider the adjoint operator T^* whose Schwartz kernel $(\mathbf{x}, (t, \boldsymbol{\xi})) \mapsto h_{\mathbf{x}}(t, \boldsymbol{\xi})$ is obtained by transposition. We now show that $T^* : L_q(\mathbb{R} \times \mathbb{S}^{d-1}) \rightarrow L_{\infty, -\alpha}(\mathbb{R}^d)$ with $q = \frac{p}{p-1}$ is bounded which, by duality, implies that the same holds true for $(T^*)^* = T : L_{1, \alpha}(\mathbb{R}^d) \xrightarrow{c} L_p(\mathbb{R} \times \mathbb{S}^{d-1})$ since $L_{1, \alpha}(\mathbb{R}^d)$ is isometrically embedded in its bidual $(L_{1, \alpha}(\mathbb{R}^d))'' = (L_{\infty, -\alpha}(\mathbb{R}^d))'$. The action of the adjoint operator is described as

$$f(\mathbf{x}) = T^*\{g\}(\mathbf{x}) = \int_{\mathbb{R}} \int_{\mathbb{S}^{d-1}} h_{\mathbf{x}}(t, \boldsymbol{\xi}) g(t, \boldsymbol{\xi}) d\boldsymbol{\xi} dt$$

which, with the help of Hölder's inequality, yields the bound

$$\begin{aligned} |(1 + \|\mathbf{x}\|)^{-\alpha} f(\mathbf{x})| &\leq (1 + \|\mathbf{x}\|)^{-\alpha} \|h_{\mathbf{x}}\|_{L_p(\mathbb{R} \times \mathbb{S}^{d-1})} \|g\|_{L_q(\mathbb{R} \times \mathbb{S}^{d-1})} \\ &\leq \sup_{\mathbf{x} \in \mathbb{R}^d, \boldsymbol{\xi} \in \mathbb{S}^{d-1}} \left((1 + \|\mathbf{x}\|)^{-\alpha} S_d \|h_{\mathbf{x}}(\cdot, \boldsymbol{\xi})\|_{L_p(\mathbb{R})} \right) \|g\|_{L_q(\mathbb{R} \times \mathbb{S}^{d-1})} \\ &\leq S_d C_{p, \alpha} \|g\|_{L_q(\mathbb{R} \times \mathbb{S}^{d-1})}, \end{aligned}$$

where $S_d = \frac{2\pi^{d/2}}{\Gamma(d/2)}$ is the surface of the unit hypersphere \mathbb{S}^{d-1} . This proves the desired result with $\|T\| = \|T^*\| \leq S_d C_{p, \alpha}$. \square

5.2 Characterization of the Predual Space \mathcal{X}_{L_R}

The application of the general representer in [58] requires that the native space $\mathcal{X}'_{L_R} = (\mathcal{X}_{L_R})'$ be identifiable as the topological dual of some primary Banach space \mathcal{X}_{L_R} . The specification of the proper predual space is achieved constructively through a completion process that ensures that \mathcal{X}_{L_R} is a complete normed space (Banach property).

Theorem 9 (Construction of the predual Banach space \mathcal{X}_{L_R}) *Let L be an isotropic LSI operator with radial frequency profile \widehat{L}_{rad} that is spline-admissible with a polynomial null space \mathcal{P} of degree n_0 . Then, $L_R^{*\dagger} = RL^{-1*}(\text{Id} -$*

$\text{Proj}_{\mathcal{P}'}) : L_{1,n_0}(\mathbb{R}^d) \xrightarrow{c} \mathcal{X}_{\text{Rad}}$ is bounded, and admits the unique extension $L_{\text{R}}^{*\dagger} : \mathcal{X}_{L_{\text{R}}}(\mathbb{R}^d) \xrightarrow{c} \mathcal{X}_{\text{Rad}}$ with the properties

$$\forall v \in \mathcal{X}_{\text{Rad}} : L_{\text{R}}^{*\dagger} L_{\text{R}}^* \{v\} = v \quad (86)$$

$$\forall p_0^* \in \mathcal{P}' : L_{\text{R}}^{*\dagger} \{p_0^*\} = 0 \quad (87)$$

$$\forall f \in \mathcal{X}_{L_{\text{R}}}(\mathbb{R}^d) : L_{\text{R}}^* L_{\text{R}}^{*\dagger} \{f\} = (\text{Id} - \text{Proj}_{\mathcal{P}'}) \{f\} = \text{Proj}_{\mathcal{U}} \{f\}, \quad (88)$$

where $L_{\text{R}}^* = L^* R^* K_{\text{rad}}$ and $\mathcal{X}_{L_{\text{R}}}(\mathbb{R}^d) = L_{\text{R}}^*(\mathcal{X}_{\text{Rad}}) \oplus \mathcal{P}'$ is equipped with the norm $\|f\|_{\mathcal{X}_{L_{\text{R}}}} \triangleq \max(\|L_{\text{R}}^{*\dagger} \{f\}\|_{\mathcal{X}}, \|\text{Proj}_{\mathcal{P}'} \{f\}\|_{\mathcal{P}'})$. The space $\mathcal{X}_{L_{\text{R}}}$ is complete and isomorphic to $\mathcal{X}_{\text{Rad}} \times \mathcal{P}'$ with $f = L_{\text{R}}^* \{v\} + p_0^* \mapsto (v, p_0^*) = (L_{\text{R}}^{*\dagger} \{f\}, \text{Proj}_{\mathcal{P}'} \{f\})$. Moreover, $\mathcal{S}(\mathbb{R}^d) \xrightarrow{d} L_{1,n_0}(\mathbb{R}^d) \xrightarrow{d} \mathcal{X}_{L_{\text{R}}}(\mathbb{R}^d) \xrightarrow{d} \mathcal{S}'(\mathbb{R}^d)$ with the embedding being continuous and dense.

Theorem 9 obviously also applies to scenarios where the null space is trivial by setting $\mathcal{P}' = \{0\}$ and only retaining (86), in which case $\mathcal{X}_{L_{\text{R}}}(\mathbb{R}^d) = L_{\text{R}}^*(\mathcal{X}_{\text{Rad}})$ and $\mathcal{S}(\mathbb{R}^d) \xrightarrow{d} L_1(\mathbb{R}^d) \xrightarrow{d} \mathcal{X}_{L_{\text{R}}}(\mathbb{R}^d) \xrightarrow{d} \mathcal{S}'(\mathbb{R}^d)$.

Proof Since $R^* K_{\text{rad}}(\mathcal{S}_{\text{Rad}}) = \mathcal{S}(\mathbb{R}^d)$, the image of \mathcal{S}_{Rad} under L_{R}^* is a vector space denoted by $\mathcal{S}_{L_{\text{R}}^*}(\mathbb{R}^d) = L_{\text{R}}^*(\mathcal{S}_{\text{Rad}}) = L^*(\mathcal{S}(\mathbb{R}^d))$. The spline-admissibility of L implies that L^* is injective on $\mathcal{S}(\mathbb{R}^d)$ which, in turn, translates into the injectivity of L_{R}^* on \mathcal{S}_{Rad} . The latter statement is equivalent to the existence of a linear operator $L_{\text{R}}^{*-1}|_{\mathcal{S}_{L_{\text{R}}^*}} = L_{\text{R}}^{*-1}$ (for short) such that, for any $u = L_{\text{R}}^* \{\phi\} \in \mathcal{S}_{L_{\text{R}}^*}(\mathbb{R}^d)$ with $\phi \in \mathcal{S}_{\text{Rad}}$, it holds that

$$L_{\text{R}}^{*-1} \{u\} = L_{\text{R}}^{*-1} L_{\text{R}}^* \{\phi\} = \phi.$$

Hence, if $\phi \mapsto \|\phi\|_{\mathcal{X}}$ is a norm on \mathcal{S}_{Rad} , then the same holds true for $u \mapsto \|u\|_{\mathcal{U}} \triangleq \|L_{\text{R}}^{*-1} \{u\}\|_{\mathcal{X}}$ on $\mathcal{S}_{L_{\text{R}}^*}(\mathbb{R}^d)$. This means that the normed spaces $(\mathcal{S}_{\text{Rad}}, \|\cdot\|_{\mathcal{X}})$ and $(\mathcal{S}_{L_{\text{R}}^*}(\mathbb{R}^d), \|\cdot\|_{\mathcal{U}})$ are (isometrically) isomorphic. Moreover, since $\mathcal{S}_{L_{\text{R}}^*}(\mathbb{R}^d) \subset L_{1,n_0}(\mathbb{R}^d)$ (Condition 4 in Definition 3) and $\mathcal{S}_{\text{Rad}} \xrightarrow{d} \mathcal{X}_{\text{Rad}}$, the inverse operator $L_{\text{R}}^{*-1}|_{\mathcal{S}_{L_{\text{R}}^*}}$ coincides on $\mathcal{S}_{L_{\text{R}}^*}$ with $L_{\text{R}}^{*\dagger} : L_{1,n_0}(\mathbb{R}^d) \xrightarrow{c} \mathcal{X}_{\text{Rad}}$ whose impulse response is characterized in Theorem 7. The well-posedness and boundedness of $L_{\text{R}}^{*\dagger}$ on $L_{1,n_0}(\mathbb{R}^d)$ for $n_0 = \lceil \gamma_0 - 1 \rceil$ follows from Theorem 8 and (71) in Theorem 7, which provides the required stability condition. The other fundamental ingredient is $\langle m_{\mathbf{k}}, u \rangle = \langle m_{\mathbf{k}}, L_{\text{R}}^* \{\phi\} \rangle = \langle L_{\text{R}} \{m_{\mathbf{k}}\}, \phi \rangle_{\text{Rad}} = 0$ for all $u \in \mathcal{S}_{L_{\text{R}}^*}(\mathbb{R}^d)$ and $|\mathbf{k}| \leq n_0$, which implies that $\text{Proj}_{\mathcal{P}'} \{u\} = 0$. Consequently, we have that

$$L_{\text{R}}^{*\dagger} \{u\} = L_{\text{R}}^{*-1} (\text{Id} - \text{Proj}_{\mathcal{P}'}) \{u\} = L_{\text{R}}^{*-1} \{u\} = L_{\text{R}}^{*-1} L_{\text{R}}^* \{\phi\} = \phi,$$

which confirms the equivalence of $L_{\text{R}}^{*\dagger}$ and $L_{\text{R}}^{*-1}|_{\mathcal{S}_{L_{\text{R}}^*}}$.

So far, we have shown that the operator $L_{\text{R}}^{*\dagger} : (\mathcal{S}_{L_{\text{R}}^*}, \|\cdot\|_{\mathcal{U}}) \rightarrow \mathcal{X}_{\text{Rad}}$ is an isometry. As next step, we invoke the bounded linear transformation (BLT)

theorem [44, Theorem I.7, p 9] to uniquely extend the operator to the completed space $\mathcal{U} = (\mathcal{S}_{L_R^*}, \|\cdot\|_{\mathcal{U}})$ which, by construction, is the Banach space equipped with the norm $\|\cdot\|_{\mathcal{U}}$. This extension argument also applies the other way around: Since $(\mathcal{S}_{L^*}(\mathbb{R}^d), \|\cdot\|_{\mathcal{U}}) \xrightarrow{\text{iso.}} \mathcal{U}$, the operator $L_R^* : (\mathcal{S}_{\text{Rad}}, \|\cdot\|_{\mathcal{X}}) \rightarrow \mathcal{U}$ has a unique (isometric) extension $L_R^* : \mathcal{X}_{\text{Rad}} \rightarrow \mathcal{U}$ with \mathcal{X}_{Rad} being the closure $\mathcal{X}_{\text{Rad}} = \overline{(\mathcal{S}_{\text{Rad}}, \|\cdot\|_{\mathcal{X}})}$. This proves that the spaces \mathcal{X}_{Rad} and \mathcal{U} are isometrically isomorphic with $\mathcal{U} = L_R^*(\mathcal{X}_{\text{Rad}})$ and $\mathcal{X}_{\text{Rad}} = L_R^{*\dagger}(\mathcal{U})$. In addition, we have that $\mathcal{U} \perp \mathcal{P}$, which means that $\text{Proj}_{\mathcal{P}'}\{u\} = 0$ for all $u \in \mathcal{U}$. Since \mathcal{U} and \mathcal{P}' are both Banach spaces, the direct-sum space $\mathcal{X}_{L_R} = \mathcal{U} \oplus \mathcal{P}'$, equipped with the composite norm $\|f\|_{\mathcal{X}_{L_R}} = \max(\|\text{Proj}_{\mathcal{U}}\{f\}\|_{\mathcal{U}}, \|\text{Proj}_{\mathcal{P}'}\{f\}\|_{\mathcal{P}'})$, is complete (Banach property) and isomorphic to $\mathcal{X}_{\text{Rad}} \times \mathcal{P}'$. The final element is to recognize that $\|\text{Proj}_{\mathcal{U}}\{f\}\|_{\mathcal{U}} = \|L_R^{*\dagger}\text{Proj}_{\mathcal{U}}\{f\}\|_{\mathcal{X}} = \|L_R^{*\dagger}\{f\}\|_{\mathcal{X}}$, where $\text{Proj}_{\mathcal{U}} = (\text{Id} - \text{Proj}_{\mathcal{P}'})$. This direct-sum decomposition has an equivalent description in terms of operators, which is the form given in the statement of Theorem 9. Specifically, the isomorphism between \mathcal{U} and \mathcal{X} is expressed by (86) and (88) for $f \in \mathcal{U}$. This is complemented by the null-space property (87), which ensures that the components of f that are in \mathcal{P}' are annihilated by L_R^* .

Embeddings: The denseness of \mathcal{S} in \mathcal{X}_{L_R} follows from the observation that $\mathcal{S}(\mathbb{R}^d) = (\text{Id} - \text{Proj}_{\mathcal{P}'}) (\mathcal{S}(\mathbb{R}^d)) \oplus \mathcal{P}'$. Since, by construction, one has that $(\text{Id} - \text{Proj}_{\mathcal{P}'}) (\mathcal{S}) \xrightarrow{\text{d.}} \mathcal{U}$ and $(\text{Id} - \text{Proj}_{\mathcal{P}'}) (L_{1,n_0}) \xrightarrow{\text{d.}} \mathcal{U}$, one also has that $\mathcal{S}(\mathbb{R}^d) \xrightarrow{\text{d.}} L_{1,n_0}(\mathbb{R}^d) \xrightarrow{\text{d.}} \mathcal{X}_{L_R}(\mathbb{R}^d)$.

As for the relation $\mathcal{X}_{L_R} = \mathcal{U} \oplus \mathcal{P}' \xrightarrow{\text{d.}} \mathcal{S}'(\mathbb{R}^d)$, we already have that $\mathcal{P}' \hookrightarrow \mathcal{S}(\mathbb{R}^d) \hookrightarrow \mathcal{S}'(\mathbb{R}^d)$, by construction. To show that $\mathcal{U} \hookrightarrow \mathcal{S}'(\mathbb{R}^d)$, we invoke the intertwining relation and express $\mathcal{U} = L^*R^*K_{\text{rad}}(\mathcal{X}_{\text{Rad}})$ as $\mathcal{U} = R^*Q_{\text{rad}}(\mathcal{X}_{\text{Rad}})$, where $Q_{\text{rad}} = L_{\text{rad}}K_{\text{rad}} : \mathcal{X}_{\text{Rad}} \rightarrow \mathcal{S}'(\mathbb{R} \times \mathbb{S}^{d-1})$. Our hypotheses on \tilde{L}_{rad} ensure that $Q_{\text{rad}}\{\phi\}$ is well-defined for every $\phi \in \mathcal{X}$, with the operator being continuous in the weak topology of $\mathcal{S}'(\mathbb{R} \times \mathbb{S}^{d-1})$. (As the latter is a complete nuclear space, the weak (sequential) convergence also ensures continuity in the strong topology [55].) Since $R^* : \mathcal{S}'(\mathbb{R} \times \mathbb{S}^{d-1}) \xrightarrow{\text{c.}} \mathcal{S}'(\mathbb{R}^d)$, as implied by (27), the composed operator $R^*Q_{\text{rad}} : \mathcal{X}_{\text{Rad}} \rightarrow \mathcal{S}'(\mathbb{R}^d)$ is continuous as well, which proves that \mathcal{U} and \mathcal{X}_{L_R} are both continuously embedded in $\mathcal{S}'(\mathbb{R}^d)$. The latter embedding is also dense by transitivity since $\mathcal{S}(\mathbb{R}^d) \subset \mathcal{X}_{L_R}$ and $\mathcal{S}(\mathbb{R}^d) \xrightarrow{\text{d.}} \mathcal{S}'(\mathbb{R}^d)$ [51, 19].

□

A Appendix: Duality Maps

The generalization of the Cauchy-Schwarz inequality for any dual pair $(\mathcal{X}, \mathcal{X}')$ of Banach spaces is the so-called *duality inequality*

$$\forall (f, x) \in \mathcal{X}' \times \mathcal{X} : \langle f, x \rangle_{\mathcal{X}' \times \mathcal{X}} \leq \|f\|_{\mathcal{X}'} \|x\|_{\mathcal{X}}, \quad (89)$$

which is tightly linked to the definition of the dual norm given by

$$\|f\|_{\mathcal{X}'} \triangleq \sup_{x \in \mathcal{X} : \|x\|_{\mathcal{X}} \leq 1} \langle f, x \rangle_{\mathcal{X}' \times \mathcal{X}}. \quad (90)$$

Definition 8 (Strict convexity) A Banach space \mathcal{X} is said to be strictly convex if, for all $x_1, x_2 \in \mathcal{X}$ such that $\|x_1\|_{\mathcal{X}} = \|x_2\|_{\mathcal{X}} = 1$ and $x_1 \neq x_2$, one has that $\|\theta x_1 + (1-\theta)x_2\|_{\mathcal{X}} < 1$ for all $\theta \in (0, 1)$.

It is obvious from (90) that (89) is sharp. Moreover, when \mathcal{X} is reflexive and strictly convex, there is a single element $x^* \in \mathcal{X}'$ (the Banach conjugate of x) such that $\|x^*\|_{\mathcal{X}'} = \|x\|_{\mathcal{X}}$ (isometry) and $\langle x^*, x \rangle_{\mathcal{X}' \times \mathcal{X}} = \|x^*\|_{\mathcal{X}'} \|x\|_{\mathcal{X}}$ (sharp duality bound) [15]. This leads to the definition of the corresponding duality map $J_{\mathcal{X}} : \mathcal{X} \rightarrow \mathcal{X}'$ as

$$J_{\mathcal{X}}\{x\} = x^*. \quad (91)$$

Since the dual of \mathcal{X}' is strictly convex as well, we have that $J_{\mathcal{X}'}^{-1} = J_{\mathcal{X}} : \mathcal{X}' \rightarrow \mathcal{X}$ with $J_{\mathcal{X}'}\{x^*\} = x$, where $(x^*)^* = x \in \mathcal{X}'' = \mathcal{X}$ is the unique Banach conjugate of $x^* \in \mathcal{X}'$.

A relevant example of reflexive and strictly convex Banach space is $\mathcal{X} = L_q(\mathbb{R} \times \mathbb{S}^{d-1})$ for $q \in (1, \infty)$. Its topological dual is $\mathcal{X}' = L_p(\mathbb{R} \times \mathbb{S}^{d-1})$ with $p = q/(q-1)$ being the conjugate exponent of q . For that particular pair, (89) reduces to the Hölder inequality for hyper-spherical functions. From [15, Chapter 4], the corresponding duality map $J_q : L_q(\mathbb{R} \times \mathbb{S}^{d-1}) \rightarrow L_p(\mathbb{R} \times \mathbb{S}^{d-1})$ is

$$J_q\{\nu\}(\mathbf{z}) = \nu^*(\mathbf{z}) = \frac{|\nu(\mathbf{z})|^{q-1}}{\|\nu\|_{L_q}^{q-2}} \text{sign}(\nu(\mathbf{z})), \quad (92)$$

which establishes a one-to-one isometric mapping between the spaces L_q and $L_p = (L_q)'$ with the property that $J_q^{-1} = J_p$.

Proposition 4 (Banach isometries) Let $(\mathcal{X}, \mathcal{X}')$ be a dual pair of reflexive and strictly convex Banach spaces with corresponding duality map $J_{\mathcal{X}} : \mathcal{X} \rightarrow \mathcal{X}'$. We consider two generic types of linear isometries.

1) *One-to-one map:* Let $T : \mathcal{X} \rightarrow \mathcal{Y} = T(\mathcal{X})$ be an injective operator whose inverse is denoted by T^{-1} with $T^{-1}T = \text{Id}$ on \mathcal{X} . Then, $\mathcal{Y} = T(\mathcal{X}) = \{y = T\{x\} : x \in \mathcal{X}\}$ equipped with the norm $\|y\|_{\mathcal{Y}} \triangleq \|T^{-1}\{y\}\|_{\mathcal{X}}$ is a Banach space. Its continuous dual is the Banach space $\mathcal{Y}' = T^{-1*}(\mathcal{X}')$ with $\|y'\|_{\mathcal{Y}'} = \|T^*\{y'\}\|_{\mathcal{X}'}$, while the corresponding duality map is $J_{\mathcal{Y}} = (T^*)^{-1} \circ J_{\mathcal{X}} \circ T^{-1} : \mathcal{Y} \rightarrow \mathcal{Y}'$.

2) *Projection:* Let $P : \mathcal{X} \rightarrow \mathcal{U} \xrightarrow{\text{iso.}} \mathcal{X}$ be a continuous projection operator on \mathcal{X} with $\|P\| = 1$. Then, $(P(\mathcal{X}), P^*(\mathcal{X}')) = (\mathcal{U}, \mathcal{U}')$ is a dual pair of Banach subspaces with corresponding duality map $J_{\mathcal{U}} = P^* \circ J_{\mathcal{X}} \circ P : \mathcal{U} \rightarrow \mathcal{U}'$.

Proof We first recall that the dual of a reflexive and strictly convex Banach space is reflexive (by definition) and strictly convex as well.

1) *Injective operator:* For the first property, we refer to [58, Proposition 1]. The key observation is that the operators $T : \mathcal{X} \rightarrow \mathcal{Y}$ and $T^{-1} : \mathcal{Y} \rightarrow \mathcal{X}$, as well as their adjoints, are isometries with $(T^*)^{-1} = T^{-1*}$. The argument then primarily rests upon the duality inequality

$$\begin{aligned} \langle y', y \rangle_{\mathcal{Y}' \times \mathcal{Y}} &= \langle y', TT^{-1}\{y\} \rangle_{\mathcal{Y}' \times \mathcal{Y}} = \langle T^*\{y'\}, T^{-1}\{y\} \rangle_{\mathcal{X}' \times \mathcal{X}} \\ &\leq \|T^*\{y'\}\|_{\mathcal{X}'} \|T^{-1}\{y\}\|_{\mathcal{X}}, \end{aligned} \quad (93)$$

which is sharp if and only if $x = T^{-1}\{y\}$ and $x' = T^*\{y'\}$ (resp., y' and y) are Banach conjugates, so that $x' = J_{\mathcal{X}}\{x\}$.

2) *Projection operator:* The first part is obtained by using a standard argument with (projected) Cauchy sequences. For the second part, let $u \in \mathcal{U} \xrightarrow{\text{iso.}} \mathcal{X}$ with Banach conjugate $u^* \in \mathcal{X}'$. Then,

$$\begin{aligned} \|u^*\|_{\mathcal{X}'} \|u\|_{\mathcal{X}} &= \langle u^*, u \rangle_{\mathcal{X}' \times \mathcal{X}} = \langle u^*, P^2u \rangle_{\mathcal{X}' \times \mathcal{X}} = \langle P^*u^*, Pu \rangle_{\mathcal{U}' \times \mathcal{U}} \\ &\leq \|P^*\| \|u^*\|_{\mathcal{U}'} \|P\| \|u\|_{\mathcal{U}}, \end{aligned} \quad (94)$$

from which we deduce that $\langle P^*u^*, Pu \rangle_{\mathcal{U}' \times \mathcal{U}} = \|u^*\|_{\mathcal{U}'} \|u\|_{\mathcal{U}}$. We conclude that $u = Pu$ and $u^* = P^*u^* = P^* \circ J_{\mathcal{X}}\{Pu\}$ are (unique) Banach conjugates of each other. \square

Acknowledgments

The research was supported by the European Commission under Grant ERC-2020-AdG FunLearn-101020573. The author is thankful to Sebastian Neumayer for very helpful discussions.

References

1. Alvarez, M.A., Rosasco, L., Lawrence, N.D.: Kernels for vector-valued functions: A review. *Foundations and Trends in Machine Learning* **4**(3), 195–266 (2012)
2. Aronszajn, N.: Theory of reproducing kernels. *Transactions of the American Mathematical Society* **68**(3), 337–404 (1950)
3. Aziznejad, S., Unser, M.: Multikernel regression with sparsity constraint. *SIAM Journal on Mathematics of Data Science* **3**(1), 201–224 (2021). DOI 10.1137/20m1318882
4. Bach, F.: Breaking the curse of dimensionality with convex neural networks. *Journal of Machine Learning Research* **18**, 1–53 (2017)
5. Barron, A.R.: Universal approximation bounds for superpositions of a sigmoidal function. *IEEE Transactions on Information Theory* **39**(3), 930–945 (1993). DOI 10.1109/18.256500
6. Bartolucci, F., Vito, E.D., Rosasco, L., Vigogna, S.: Understanding neural networks with reproducing kernel Banach spaces. *Applied and Computational Harmonic Analysis* **62**, 194–236 (2023)
7. Berlinet, A., Thomas-Agnan, C.: *Reproducing Kernel Hilbert Spaces in Probability and Statistics*, vol. 3. Kluwer Academic Boston (2004)
8. Bishop, C.M.: *Pattern Recognition and Machine Learning*. Springer (2006)
9. de Boor, C.: On “best” interpolation. *Journal of Approximation Theory* **16**(1), 28–42 (1976)
10. de Boor, C., Lynch, R.E.: On splines and their minimum properties. *Journal of Mathematics and Mechanics* **15**(6), 953–969 (1966)
11. Boyer, C., Chambolle, A., De Castro, Y., Duval, V., De Gournay, F., Weiss, P.: On representer theorems and convex regularization. *SIAM Journal of Optimization* **29**(2), 1260–1281 (2019)
12. Bredies, K., Carioni, M.: Sparsity of solutions for variational inverse problems with finite-dimensional data. *Calculus of Variations and Partial Differential Equations* **59**(14), 26 (2020)
13. Buhmann, M.D.: *Radial Basis Functions*. Cambridge University Press (2003)
14. Chen, Z., Haykin, S.: On different facets of regularization theory. *Neural Computation* **14**(12), 2791–2846 (2002)
15. Cioranescu, I.: *Geometry of Banach Spaces, Duality Mappings and Nonlinear Problems*, vol. 62. Springer Science & Business Media (2012)
16. Cybenko, G.: Approximation by superpositions of a sigmoidal function. *Mathematics of Control, Signals and Systems* **2**(4), 303–314 (1989)
17. Duchon, J.: Splines minimizing rotation-invariant semi-norms in Sobolev spaces. In: W. Schempp, K. Zeller (eds.) *Constructive Theory of Functions of Several Variables*, pp. 85–100. Springer-Verlag, Berlin (1977)
18. Fisher, S.D., Jerome, J.W.: Spline solutions to L_1 extremal problems in one and several variables. *Journal of Approximation Theory* **13**(1), 73–83 (1975)
19. Gelfand, I.M., Shilov, G.: *Generalized Functions*. Vol. 1. Properties and Operations. Academic Press, New York (1964)
20. Gelfand, I.M., Shilov, G.: *Generalized Functions*. Vol. 5. Integral Geometry and Representation Theory. Academic Press, New York (1966)
21. Grafakos, L.: *Classical Fourier Analysis*. Springer (2008)
22. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, pp. 770–778 (2016)
23. Helgason, S.: *Integral Geometry and Radon Transforms*. Springer (2011)

24. Hofmann, T., Schölkopf, B., Smola, A.J.: Kernel methods in machine learning. *Annals of Statistics* **36**(3), 1171–1220 (2008)
25. Hornik, K., Stinchcombe, M., White, H.: Multilayer feedforward networks are universal approximators. *Neural Networks* **2**(5), 359–366 (1989)
26. Kostadinova, S., Pilipović, S., Saneva, K., Vindas, J.: The ridgelet transform of distributions. *Integral Transforms and Special Functions* **25**(5), 344–358 (2014)
27. Ludwig, D.: The Radon transform on Euclidean space. *Communications on Pure and Applied Mathematics* **19**(1), 49–81 (1966). DOI 10.1002/cpa.3160190105
28. Mammen, E., van de Geer, S.: Locally adaptive regression splines. *Annals of Statistics* **25**(1), 387–413 (1997)
29. Meinguet, J.: Multivariate interpolation at arbitrary points made simple. *Zeitschrift für Angewandte Mathematik und Physik* **30**, 292–304 (1979)
30. Mhaskar, H., Micchelli, C.A.: Approximation by superposition of sigmoidal and radial basis functions. *Advances in Applied Mathematics* **13**(3), 350–373 (1992). DOI 10.1016/0196-8858(92)90016-p
31. Micchelli, C.A.: Interpolation of scattered data: Distance matrices and conditionally positive definite functions. *Constructive Approximation* **2**(1), 11–22 (1986)
32. Micchelli, C.A., Xu, Y., Zhang, H.: Universal kernels. *Journal of Machine Learning Research* **7**, 2651–2667 (2006)
33. Natterer, F.: *The Mathematics of Computed Tomography*. John Wiley & Sons Ltd. (1984)
34. Neumayer, S., Unser, M.: Explicit representations for Banach subspaces of Lizorkin distributions. *Analysis and Applications* (in press). DOI 10.1142/S0219530523500148
35. Ongie, G., Willett, R., Soudry, D., Srebro, N.: A function space view of bounded norm infinite width ReLU nets: The multivariate case. *International Conference on Representation Learning (ICLR)* (2020)
36. Parhi, R., Nowak, R.D.: The role of neural network activation functions. *IEEE Signal Processing Letters* **27**, 1779–1783 (2020). DOI 10.1109/LSP.2020.3027517
37. Parhi, R., Nowak, R.D.: Banach space representer theorems for neural networks and ridge splines. *Journal of Machine Learning Research* **22**(43), 1–40 (2021)
38. Parhi, R., Nowak, R.D.: What kinds of functions do deep neural networks learn? Insights from variational spline theory. *SIAM Journal on Mathematics of Data Science* **4**(2), 464–489 (2022). DOI 10.1137/21m1418642
39. Pinkus, A.: Approximation theory of the MLP model in neural networks. *Acta Numerica* **8**, 143–195 (1999). DOI 10.1017/s0962492900002919
40. Poggio, T., Girosi, F.: Networks for approximation and learning. *Proceedings of the IEEE* **78**(9), 1481–1497 (1990)
41. Poggio, T., Girosi, F.: Regularization algorithms for learning that are equivalent to multilayer networks. *Science* **247**(4945), 978–982 (1990)
42. Poggio, T., Smale, S.: The mathematics of learning: Dealing with data. *Notices of the AMS* **50**(5), 537–544 (2003)
43. Ramm, A.G., Katsevich, A.I.: *The Radon transform and local tomography*. CRC Press (2020)
44. Reed, M., Simon, B.: *Methods of Modern Mathematical Physics. Vol. 1: Functional Analysis*. Academic Press, San Diego (1980)
45. Rudin, W.: *Functional Analysis*, 2nd edn. McGraw-Hill, New York (1991). McGraw-Hill Series in Higher Mathematics
46. Samko, S.G.: Denseness of the Lizorkin-type spaces Φ_V in $L_p(\mathbb{R}^n)$. *Mathematical Notes of the Academy of Sciences of the USSR* **31**(6), 432–437 (1982)
47. Samko, S.G., Kilbas, A.A., Marichev, O.I.: *Fractional Integrals and Derivatives: Theory and Applications*. Gordon and Breach Science Publishers (1993)
48. Savarese, P., Evron, I., Soudry, D., Srebro, N.: How do infinite width bounded norm networks look in function space? In: *Conference on Learning Theory*. PMLR, pp. 2667–2690 (2019)
49. Schölkopf, B., Herbrich, R., Smola, A.J.: A generalized representer theorem. In: D. Helmbold, B. Williamson (eds.) *Computational Learning Theory*, pp. 416–426. Springer Berlin Heidelberg (2001)

50. Schölkopf, B., Sung, K.K., Burges, C.J.C., Girosi, F., Niyogi, P., Poggio, T., Vapnik, V.: Comparing support vector machines with Gaussian kernels to radial basis function classifiers. *IEEE Transactions on Signal Processing* **45**(11), 2758–2765 (1997)
51. Schwartz, L.: *Théorie des Distributions*. Hermann, Paris (1966)
52. Shawe-Taylor, J., Cristianini, N., et al.: *Kernel methods for pattern analysis*. Cambridge university press (2004)
53. Sonoda, S., Ishikawa, I., Ikeda, M.: Ridge regression with over-parametrized two-layer networks converge to ridgelet spectrum. pp. 2674–2682. PMLR (2021)
54. Sonoda, S., Murata, N.: Neural network with unbounded activation functions is universal approximator. *Applied and Computational Harmonic Analysis* **43**(2), 233–268 (2017)
55. Trèves, F.: *Topological Vector Spaces, Distributions and Kernels*. Dover Publications, New York (2006)
56. Unser, M.: A unifying representer theorem for inverse problems and machine learning. *Foundations of Computational Mathematics* **21**(4), 941–960 (2021). DOI 10.1007/s10208-020-09472-x
57. Unser, M.: Ridges, neural networks, and the Radon transform. *Journal of Machine Learning Research* **24**, 1–33 (2023)
58. Unser, M., Aziznejad, S.: Convex optimization in sums of Banach spaces. *Applied and Computational Harmonic Analysis* **56**, 1–25 (2022). DOI 10.1016/j.acha.2021.07.002
59. Unser, M., Blu, T.: Fractional splines and wavelets. *SIAM Review* **42**(1), 43–67 (2000)
60. Unser, M., Fageot, J., Ward, J.P.: Splines are universal solutions of linear inverse problems with generalized-TV regularization. *SIAM Review* **59**(4), 769–793 (2017)
61. Unser, M., Tafti, P.D.: *An Introduction to Sparse Stochastic Processes*. Cambridge University Press (2014)
62. Wahba, G.: *Spline Models for Observational Data*. Society for Industrial and Applied Mathematics, Philadelphia (1990)
63. Wendland, H.: *Scattered Data Approximations*. Cambridge University Press (2005)