

Geometric harmonics: A novel tool for multiscale out-of-sample extension of empirical functions

Ronald R. Coifman*, Stéphane Lafon¹

Applied Mathematics Department, Yale University, New Haven, CT 06510, USA

Received 29 October 2004; revised 19 July 2005; accepted 29 July 2005

Available online 21 June 2006

Communicated by the Editors

Abstract

We describe a simple scheme, based on the Nyström method, for extending empirical functions f defined on a set X to a larger set \bar{X} . The extension process that we describe involves the construction of a specific family of functions that we term *geometric harmonics*. These functions constitute a generalization of the prolate spheroidal wave functions of Slepian in the sense that they are optimally concentrated on X . We study the case when X is a submanifold of \mathbb{R}^n in greater detail. In this situation, any empirical function f on X can be characterized by its decomposition over the intrinsic Fourier modes, i.e., the eigenfunctions of the Laplace–Beltrami operator, and we show that this intrinsic frequency spectrum determines the largest domain of extension of f to the entire space \mathbb{R}^n . Our analysis relates the complexity of the function on the training set to the scale of extension off this set. This approach allows us to present a novel multiscale extension scheme for empirical functions.

© 2006 Published by Elsevier Inc.

Keywords: Nyström method; Intrinsic and extrinsic geometries; Subsampling; Prolate functions

1. Introduction

In applications where a large amount of data is involved, the only way to perform certain tasks like clustering, regression and classification is to subsample the data set \bar{X} in order to reduce the size of the problem, process the new set X , and then extend the result to the original data \bar{X} . If in addition, the data are embedded in a high-dimensional space, then kernel methods like those based on radial functions constitute an interesting alternative to grid-based methods (which are unusable in high dimension).

In this paper, we describe a scheme for extending empirical functions defined on X to \bar{X} . This technique is inspired from the Nyström method, widely used in partial differential solvers, and recently employed in machine learning and in spectral graph theory as a way to subsample large data sets [3,9,13]. For instance, in order to perform an efficient (fast) image segmentation, it was noticed that one could subsample the pixels of an image, do the segmentation in this subgrid, and then extend the results to the entire image. More recently, Belkin et al. [4] presented the idea of manifold

* Corresponding author.

E-mail addresses: coifman@math.yale.edu (R.R. Coifman), stephane.lafon@gmail.com (S. Lafon).

¹ Now with Google Inc.

regularization, introducing an out-of-sample extension scheme that is closely related to the Nyström extension. Our work is also highly related to the Kriging technique widely employed in geostatistics [5]. However, with all these techniques, the critical question of the choice of a scale of extension remains unanswered.

The contribution of this article is threefold. First, we show how the Nyström method naturally leads to the construction of a particular set of functions that we term *geometric harmonics*. We show that these functions are solution of an eigenproblem and generalize the prolate spheroidal wave functions introduced in [11,12]. In particular, they are maximally concentrated on X . Secondly, our analysis relates the complexity of the empirical function, measured in terms of frequency content over the training set, to the distance to which it can be extended off this set. Indeed, in the case where X is a submanifold of \mathbb{R}^n , any function f can be decomposed as a sum of intrinsic Fourier modes, namely, the eigenfunctions of the Laplace–Beltrami operator. This allows us to define a notion of intrinsic bandwidth, and we show how the geometric harmonics relate the intrinsic oscillations of functions on X to the extrinsic frequencies of their extensions in the ambient space \mathbb{R}^n . More precisely, we show that the restriction and extension operations defined via the geometric harmonics preserve the bandwidth of signals, modulo a multiplicative constant quantifying the complexity of the embedding of X into \mathbb{R}^n . Last, we prove that, following some version of the Heisenberg principle, a function f with intrinsic bandlimit ν on X can be extended as a function with the same extrinsic bandlimit on \mathbb{R}^n and is numerically supported in a tubular neighborhood of radius $\mathcal{O}(\frac{1}{\nu})$ around X . This allows us to define a multiscale extension scheme for empirical functions in which each function is decomposed as a superposition of atoms with specific time–frequency localization.

The paper is organized as follows. In Section 2, we introduce the notation and explain the construction of the geometric harmonics and their properties. We also describe the associated extension algorithm of empirical functions. In Section 3, we give examples of geometric harmonics, and in particular we show that they generalize the prolate spheroidal wave functions of Slepian and Pollak [11,12]. Last, in Section 4, we focus on the case when X is a submanifold of \mathbb{R}^n and we relate the intrinsic and extrinsic Fourier analyses of empirical functions. This allows us to present a novel multiscale extension algorithm.

2. Geometric harmonics

2.1. Construction

Let X and \bar{X} be two sets such that $X \subset \bar{X}$. Let μ be a finite measure on X , i.e., $\mu(X) < +\infty$. Our goal is to be able to extend a function f defined on X to the set \bar{X} and for that to be done, we construct a set of functions, the geometric harmonics, that allow us to perform this extension. Notation: functions defined on \bar{X} will be denoted using capital letters, whereas functions on X will be represented using lower case letters. In this section, arbitrary points in \bar{X} will be denoted with a bar, for instance \bar{x} .

Our main ingredient is a “kernel” $k: \bar{X} \times \bar{X} \rightarrow \mathbb{R}$ satisfying

- k is symmetric: $k(\bar{x}, \bar{y}) = k(\bar{y}, \bar{x})$ for all \bar{x} and \bar{y} in \bar{X} .
- k is positive semi-definite, i.e., for any $m \geq 1$ and any choice of real numbers $\alpha_1, \dots, \alpha_m$, and of points $\bar{x}_1, \dots, \bar{x}_m$ in \bar{X} , we have

$$\sum_{i=1}^m \sum_{j=1}^m \alpha_i \alpha_j k(\bar{x}_i, \bar{x}_j) \geq 0.$$

This property is not necessary to be able to define geometric harmonics and extensions. However, it allows us to interpret the geometric harmonics as maximizing some concentration measure over X .

- k is bounded on $\bar{X} \times \bar{X}$ by a number $M > 0$. Although this assumption can be weakened, it is very convenient for the simplicity of the exposition.

Since k is positive semi-definite, there exists a unique reproducing kernel Hilbert space \mathcal{H} of functions defined on \bar{X} for which k is the reproducing kernel (see Appendix A). Let $\langle \cdot, \cdot \rangle_{\bar{X}}$ and $\| \cdot \|_{\bar{X}}$ denote the inner product and norm of this space. In contrast, we use $\langle \cdot, \cdot \rangle_X$ and $\| \cdot \|_X$ to represent the inner product and norm in $L^2(X, d\mu)$.

The kernel k can be restricted to X and employed to define a linear operator $\mathbf{K} : L^2(X, d\mu) \rightarrow \mathcal{H}$:

$$\mathbf{K}f(\bar{x}) = \int_X k(\bar{x}, x)f(x) d\mu(x), \quad \text{where } \bar{x} \in \bar{X}.$$

We have the following lemma:

Lemma 1. *The adjoint $\mathbf{K}^* : \mathcal{H} \rightarrow L^2(X, d\mu)$ is the restriction operator on X : for all $F \in \mathcal{H}$ and $x \in X$,*

$$\mathbf{K}^*F(x) = F(x).$$

Moreover, $\mathbf{K}^\mathbf{K} : L^2(X, d\mu) \rightarrow L^2(X, d\mu)$ is compact.*

Proof. First, it can be checked that

$$\begin{aligned} \langle F, \mathbf{K}g \rangle_{\bar{X}} &= \left\langle F(\cdot), \int_X k(\cdot, x)g(x) d\mu(x) \right\rangle_{\bar{X}} \\ &= \int_X \langle F(\cdot), k(\cdot, x) \rangle_{\bar{X}} g(x) d\mu(x) \\ &= \int_X F(x)g(x) d\mu(x) \quad \text{by identity (A.1)} \\ &= \langle f, g \rangle_X. \end{aligned}$$

This proves that $\mathbf{K}^*F = f$. We now show that $\mathbf{K}^*\mathbf{K}$ is Hilbert–Schmidt, which implies compactness:

$$\iint_{X \times X} |k(x, y)|^2 d\mu(y) d\mu(x) < +\infty.$$

Indeed, we have

$$\iint_{X \times X} |k(x, y)|^2 d\mu(y) d\mu(x) \leq M^2 \mu(X)^2 < +\infty. \quad \square$$

Since this operator is self-adjoint and compact on $L^2(X, d\mu)$, it has a discrete sequence of eigenvalues $\{\lambda_j\}$ (in nonincreasing order) and eigenvectors $\{\psi_j\}$ defined on X : for $d\mu$ -almost all $x \in X$,

$$\lambda_j \psi_j(x) = \int_X k(x, y)\psi_j(y) d\mu(y).$$

Note that, because the operator is positive semi-definite, for all j , $\lambda_j \geq 0$. Also, by the result of this lemma, ψ_j is obtained by diagonalizing the kernel $k(x, y)$ with x and y restricted to X .

The eigenfunction ψ_j is defined on X , but provided $\lambda_j \neq 0$, it can be extended to the set \bar{X} via a technique known as the Nyström method (see [10]). This technique consists in observing that $k(x, y)$ is also defined for $x \in \bar{X}$, and therefore so is the right-hand side of the above equation, and that it can be used for the definition ψ_j outside X .

Definition 2. When $\lambda_j \neq 0$, the eigenfunction ψ_j can be extended to $\bar{x} \in \bar{X}$ by

$$\Psi_j(\bar{x}) = \frac{1}{\lambda_j} \int_X k(\bar{x}, y)\psi_j(y) d\mu(y).$$

This extension is called *geometric harmonic*.

The term “geometric harmonics” is inspired from the fact that ψ_j is extended as an average of its values on the set X , and thus can be thought of as verifying a certain form of mean value theorem.

Numerically, this extension procedure is extremely ill-conditioned as one divides by the eigenvalues of a compact operator ($\lambda_j \rightarrow 0$ as $j \rightarrow +\infty$). Consequently, we introduce the following notations.

Definition 3. For any $\delta > 0$, let

$$S_\delta = \{j \text{ such that } \lambda_j \geq \delta \lambda_0\}$$

and define the following finite-dimensional vector spaces:

$$L_\delta^2 = \text{Span}\{\psi_j, j \in S_\delta\}$$

and

$$\mathcal{H}_\delta = \text{Span}\{\Psi_j, j \in S_\delta\}.$$

The extension procedure from L_δ^2 to \mathcal{H}_δ has a condition number bounded from above by $\frac{1}{\delta}$. We summarize the algebraic relation between \mathbf{K} , \mathbf{K}^* , Ψ_j and ψ_j :

$$\mathbf{K}\psi_j = \lambda_j \Psi_j \quad (\text{extension}),$$

$$\mathbf{K}^*\Psi_j = \psi_j \quad (\text{restriction}).$$

We conclude this section with two remarks. First, an important class of positive kernels is generated by covariance kernels, i.e., by kernels of the form

$$k(x, y) = \int_{\xi \in I} e_\xi(x) e_\xi(y) p(\xi) d\xi,$$

where $\{e_\xi\}_{\xi \in I}$ is a family of functions defined on \bar{X} and $p(\xi) \geq 0$. Each function e_ξ , restricted to X , is interpreted as a vector whose coordinates are indexed by $x \in X$, and the kernel k represents the covariance of the cloud of points generated by the mass distribution $p(\xi) d\xi$. Finding the eigenfunctions and eigenvalues associated with this kernel is equivalent to computing the axes and moments of inertia of this cloud of points, which is also referred to principal component analysis.

The other remark concerns a variational interpretation of k . Let \mathbf{B} represent the orthogonal projector onto \mathcal{H} , defined by $\mathbf{B}F(x) = \langle F, k(x, \cdot) \rangle_{\mathcal{H}}$ and let \mathbf{D} be the restriction operator to X defined by $\mathbf{D}F(\bar{x}) = F(\bar{x})$ if $\bar{x} \in X$, and $\mathbf{D}F(\bar{x}) = 0$ otherwise. Then it can be checked that $\mathbf{K} = \mathbf{B}\mathbf{D}$ and if $x \in X$,

$$\mathbf{K}^*\mathbf{K}f(x) = \mathbf{D}\mathbf{B}\mathbf{D}f(x).$$

This decomposition of $\mathbf{K}^*\mathbf{K}$ as a product of orthogonal projection leads to a variational interpretation of the geometric harmonics that motivated Slepian and Pollak into introducing the prolate spheroidal wave functions [11,12].

2.2. Properties

The geometric harmonics feature two interesting properties: they are orthogonal on X and on \bar{X} , and among all functions of \mathcal{H} , they have maximum concentration on X .

Proposition 4. The system $\{\Psi_j\}_{j \in S_\delta}$ forms an orthogonal basis of \mathcal{H}_δ , and their restrictions $\{\psi_j\}_{j \in S_\delta}$ to X forms an orthogonal basis of L_δ^2 .

Proof. By definition, the ψ_j 's are obtained as the eigenfunctions of a self-adjoint operator, therefore they are orthogonal on X . In addition,

$$\langle \Psi_i, \Psi_j \rangle_{\bar{X}} = \frac{1}{\lambda_j} \langle \Psi_i, \mathbf{K}\mathbf{K}^*\Psi_j \rangle_{\bar{X}} = \frac{1}{\lambda_j} \langle \mathbf{K}^*\Psi_i, \mathbf{K}^*\Psi_j \rangle_{\bar{X}} = \frac{1}{\lambda_j} \langle \psi_i, \psi_j \rangle_X. \quad \square$$

Definition 5. For a function $F \in \mathcal{H}$, with restriction $f \in L^2(X, d\mu)$, we define the concentration of F over X to be the Rayleigh quotient

$$c_X(F) = \frac{\|f\|_X}{\|F\|_{\bar{X}}},$$

where $f = \mathbf{K}^*F$ is the restriction of F to X .

The geometric harmonics are also the function of \mathcal{H} that have maximum concentration on the set X :

Proposition 6. *The function Ψ_j is a solution to the problem*

$$\max_{F \in \mathcal{H}} c_X(F),$$

under the constraint that $F \perp \{\Psi_0, \Psi_1, \dots, \Psi_{j-1}\}$. In particular, Ψ_0 is the element of \mathcal{H} that is the most concentrated on X .

Proof. By homogeneity of the ratio, we can restrict our attention to all $F \in \mathcal{H}$ with norm 1. Thus we need to maximize

$$\langle f, f \rangle_X = \langle \mathbf{K}^*F, \mathbf{K}^*F \rangle_X = \langle F, \mathbf{K}\mathbf{K}^*F \rangle_{\bar{X}},$$

under the constraints $\langle F, F \rangle_{\bar{X}} = 0$, $\langle F, \Psi_0 \rangle_{\bar{X}} = 0$, \dots , $\langle F, \Psi_{j-1} \rangle_{\bar{X}} = 0$. Using the Lagrange multipliers technique, we conclude that there exist numbers λ and $\alpha_0, \dots, \alpha_{j-1}$ such that

$$\mathbf{K}\mathbf{K}^*F = \lambda F + \alpha_0\Psi_0 + \dots + \alpha_{j-1}\Psi_{j-1}.$$

Taking the inner product with Ψ_i for $i = 0, \dots, j - 1$, and invoking the constraints and the orthogonality of the geometric harmonics (see previous proposition), we obtain that

$$\alpha_i \|\Psi_i\|_{\bar{X}}^2 = \langle \mathbf{K}\mathbf{K}^*F, \Psi_i \rangle_{\bar{X}} = \langle F, \mathbf{K}\mathbf{K}^*\Psi_i \rangle_{\bar{X}} = \lambda_i \langle F, \Psi_i \rangle_{\bar{X}} = 0.$$

As a consequence, $\alpha_i = 0$, and F is a geometric harmonic associated with the eigenvalue λ and the functional to be minimized now takes the form

$$\langle f, f \rangle_X = \lambda.$$

The maximum is therefore achieved for $F = \Psi_j$. \square

2.3. Extension algorithm

We now describe the natural extension algorithm associated with the geometric harmonics obtained from a kernel k . This scheme only considers one scale of extension, namely that defined by the kernel k . Note that this scale is essentially arbitrary as it is a priori unrelated to the function f that one tries to extend.

We assume that ψ_j is normalized so that it has norm 1 on X .

Algorithm 1 (*Extension scheme*). Given a function $f \in L^2(X, d\mu)$,

- project f onto the space L^2_δ spanned by the orthonormal system $\{\psi_j\}_{j \in S_\delta}$:

$$f \mapsto \mathbf{P}_\delta f = \sum_{j \in S_\delta} \langle f, \psi_j \rangle_X \psi_j,$$

- use the extension Ψ_j of ψ_j to extend $\mathbf{P}_\delta f$ on \bar{X} as

$$\mathbf{E}f(\bar{x}) = \sum_{j \in S_\delta} \langle f, \psi_j \rangle_X \Psi_j(\bar{x}),$$

with $\bar{x} \in \bar{X}$.

This algorithm computes a truncated pseudo-inverse for K^* and is consistent in the sense that the restriction of $\mathbf{E}f$ to X is, again, extended as $\mathbf{E}f$ by the algorithm.

We now make three comments on this algorithm. First, this technique does not provide an extension for f but rather for a filtered version of it, namely, its orthogonal projection $\mathbf{P}_\delta f$ onto L^2_δ . This space is precisely that of functions that can be numerically extended to \bar{X} , since the condition number of the operator \mathbf{E} is bounded by $\frac{1}{\delta}$. Therefore, in order to be able to speak of an extension for f , the residual $\|f - \mathbf{P}_\delta f\|_X$ on X must be smaller than a prescribed error, which leads to the definition:

Definition 7. A function f defined on X is said to be (η, δ) -extendable if

$$\|\mathbf{P}_\delta f\|_X = \sum_{j \in S_\delta} |\langle f, \psi_j \rangle_X|^2 \geq (1 - \eta) \|f\|_{\bar{X}}^2.$$

The obvious consequence is that not all functions f can be extended at a given precision and this is where the scale of the kernel k comes into play. By scale, we mean the size of the numerical support of k in \mathbb{R}^n . A function that cannot be extended via a kernel k will eventually be extendable with a fine enough kernel. Indeed, as we reduce the scale of a kernel, the spectrum of the corresponding operator tend to be flatter (another form of the Heisenberg principle), making it possible to extend more geometric harmonics. Equivalently, if the kernel has a small scale, it is able to capture finer details of a function over the set X and to transcribe them in the extension process. *This observation is crucial for the multiscale analysis that we develop later in this article.* This fact can be used to relate the oscillations of f on X to the scale of extension. We develop this idea in Section 4, when we consider families of kernels parametrized by a scale parameter.

Another interesting aspect of the scheme is the relation to the manifold regularization framework introduced in [4]. Their work involves solving the following regularization problem:

$$\min_{F \in \mathcal{H}} \|f - \mathbf{K}^* F\|_{L^2(X)}^2 + \gamma \|F\|_{\bar{X}}^2,$$

where γ is a regularization parameter. This problem can be viewed as an instance of Ridge regression and is solved via

$$F_\gamma = (\mathbf{K}\mathbf{K}^* + \gamma \mathbf{Id})^{-1} \mathbf{K}f.$$

Therefore, as $\gamma \rightarrow 0$, and if $f = \psi_j$, one recovers the Nyström extension procedure.

It is also interesting to underline the similarity of our approach with the so-called Kriging technique [5]. The zeroth-order Kriging model is based upon extending f into $f_K = \mathbf{P}_0 f$, that is with $\delta = 0$. The Kriging model therefore computes the inverse of \mathbf{K} (provided it is invertible), whereas our approach is based on a regularized pseudo-inverse. For the Kriging model, $k(\cdot, \cdot)$ is usually a Gaussian kernel whose covariance is being optimized to maximize the fit on the training set. This optimization is global and so is the scale of extension. As we point out later in this paper, our technique also allows to work with multiple scales and to locally adapt the scale of extension.

The last point we wish to discuss concerns the interpretation of the extension. As we shall see in the next section, for a given f defined on X , there generally exist infinitely many possible extensions $F \in \mathcal{H}$ as functions in \mathcal{H} might not be determined by their values on X . However, according to Proposition 6, $\mathbf{E}f$ is the extension that has maximal concentration over X . From the point of view of Statistics, where extension means regression, this is related to the predictive power of Nyström extensions. For instance, it is known that if a random variable $(U, V) \in \mathbb{R}^{m+n}$ has the following covariance matrix:

$$C(U, V) = \begin{pmatrix} A & B \\ B^T & C \end{pmatrix},$$

then the covariance of V given U is

$$C(V|U) = \begin{pmatrix} A & B \\ B^T & C - B^T A^{-1} B \end{pmatrix},$$

while the Nyström extension method (taking X to be the set of first m indices, and \bar{X} being the set of $m + n$ indices) implicitly approximates $C(U, V)$ as

$$C(U, V) \simeq \begin{pmatrix} A & B \\ B^T & B^T A^{-1} B \end{pmatrix}.$$

Therefore, the norm of the Schur complement $\|C - B^T A^{-1} B\|$ is a measure of the average randomness of V that is left when one knows U . If the Nyström approximation is accurate, then the predictive power of U over V is large.

3. Examples of geometric harmonics

In this section we give some examples of geometric harmonics.

3.1. The prolate spheroidal wave functions—bandlimited extension

In [11], Slepian and Pollak introduce the prolate spheroidal wave functions as the solution of finding functions optimally concentrated in time and frequency. By construction, the prolates are bandlimited functions of unit energy that have maximum energy within an interval of the time domain. They also generalize their result to higher dimension [12] by defining \mathcal{H}^B to be the space of functions of $L^2(\mathbb{R}^n)$ whose Fourier is supported in the ball centered at the origin and of radius B . This space is a reproducing kernel Hilbert space with kernel (see Appendix B)

$$k_B^{(n)}(x, y) = \left(\frac{B}{2}\right)^{\frac{n}{2}} \frac{J_{\frac{n}{2}}(\pi B \|x - y\|)}{\|x - y\|^{\frac{n}{2}}},$$

where J_ν is the Bessel function of the first kind and of order ν . Such a kernel is referred to as a Bessel kernel. The prolate spheroidal wave functions are defined as the eigenfunctions of the integral operator of kernel k_B restricted to a certain space domain $X \subset \bar{X} = \mathbb{R}^n$. In the prolate setting, X is a set of positive measure (and the kernel is nonintegrable) and functions of \mathcal{H}^B are determined by their values on X . On the contrary, we are mainly interested in the case where X is singular in \mathbb{R}^n . In this situation, there are infinitely many bandlimited extensions for a function f defined on X , as two such extensions differ by a (nonzero) bandlimited function that vanishes on X . However, among all these extensions, the operator \mathbf{E} produces an extension with maximum concentration on X , or equivalently, that has minimal energy outside X .

From now on, we will refer to \mathbf{E}_B as the Bessel kernel with bandlimit $B > 0$ (we drop the dimension n in the notation); this operator computes the bandlimited extension of band B that is maximally concentrated on X . In Fig. 1, we show examples of bandlimited extensions (at a fixed band) of the functions $\cos(j\theta)$ for $j = 1, 2, 4, 8$, from the unit circle to the plane. The Bessel kernel tends to generate a large amount of oscillations outside the circle, whereas inside the circle the oscillation seem to cancel out.

3.2. Gaussian extension

Widely used in the Machine Learning community, the Gaussian kernel $\exp(-B^2 \|x - y\|^2)$ also arises as a limiting case of the Bessel kernels. Indeed, as proven in [6],

$$\lim_{n \rightarrow +\infty} \frac{k_B^{(n)}\left(\sqrt{\frac{n}{\pi}} \|x - y\|\right)}{V_{B,n}} = e^{-B^2 \|x - y\|^2},$$

where $V_{B,n}$ is the volume of the unit ball of radius B in \mathbb{R}^n . This seems to suggest that in high dimension, Bessel and Gaussian kernels are equivalent. In lower dimension, unlike Bessel kernels, Gaussian kernels do not exhibit oscillations.

3.3. Harmonic extension

Another example comes from potential theory. Consider the single-layer Newtonian potential in $\bar{X} = \mathbb{R}^n$,

$$k(x, y) = \begin{cases} -\log(\|x - y\|) & \text{if } n = 2, \\ \frac{1}{\|x - y\|^{n-2}} & \text{if } n \geq 3. \end{cases}$$

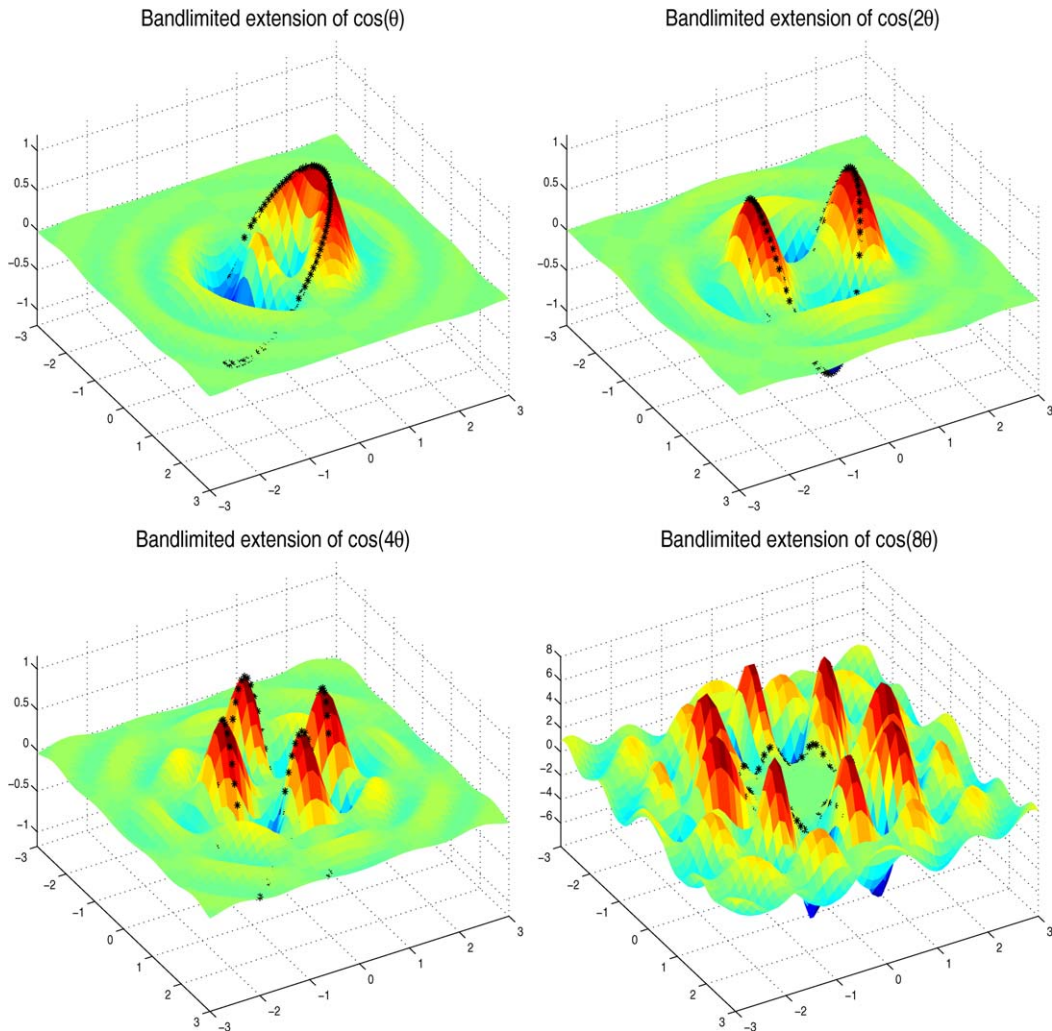


Fig. 1. Bandlimited extensions of the functions $\cos(j\theta)$ for $j = 1, 2, 4, 8$, from the unit circle to the plane.

By definition, this kernel is the Green's function for the Laplace operator on \mathbb{R}^n , and as a consequence, the potential is positive semi-definite. Assuming that $n = 3$ and that X is a Lipschitz surface in \mathbb{R}^3 , then \mathcal{H} is the space of potentials

$$F(x) = \int_X \frac{d\rho(y)}{\|x - y\|},$$

where ρ is a signed measure supported on X , representing a distribution of charges. In this space, the inner product is the electrostatic energy of interaction between two distributions of charges:

$$\langle F_1, F_2 \rangle_{\bar{X}} = \int_X \int_X \frac{d\rho_1(y) d\rho_2(x)}{\|x - y\|}.$$

Note that all elements of \mathcal{H} are harmonic and that the geometric harmonics are harmonic functions that minimize their electrostatic self-energy. The operator \mathbf{E} allows one to construct harmonic extensions for empirical functions defined on X .

An interesting feature of the harmonic extension is that, unlike the bandlimited extension, it does not involve a scale (or bandwidth) parameter. In fact, the size of the numerical support of the extension of f depends on the oscillations of f on X . Indeed, suppose that X is a closed curve in \mathbb{R}^3 , and that f oscillates p times along X . Then

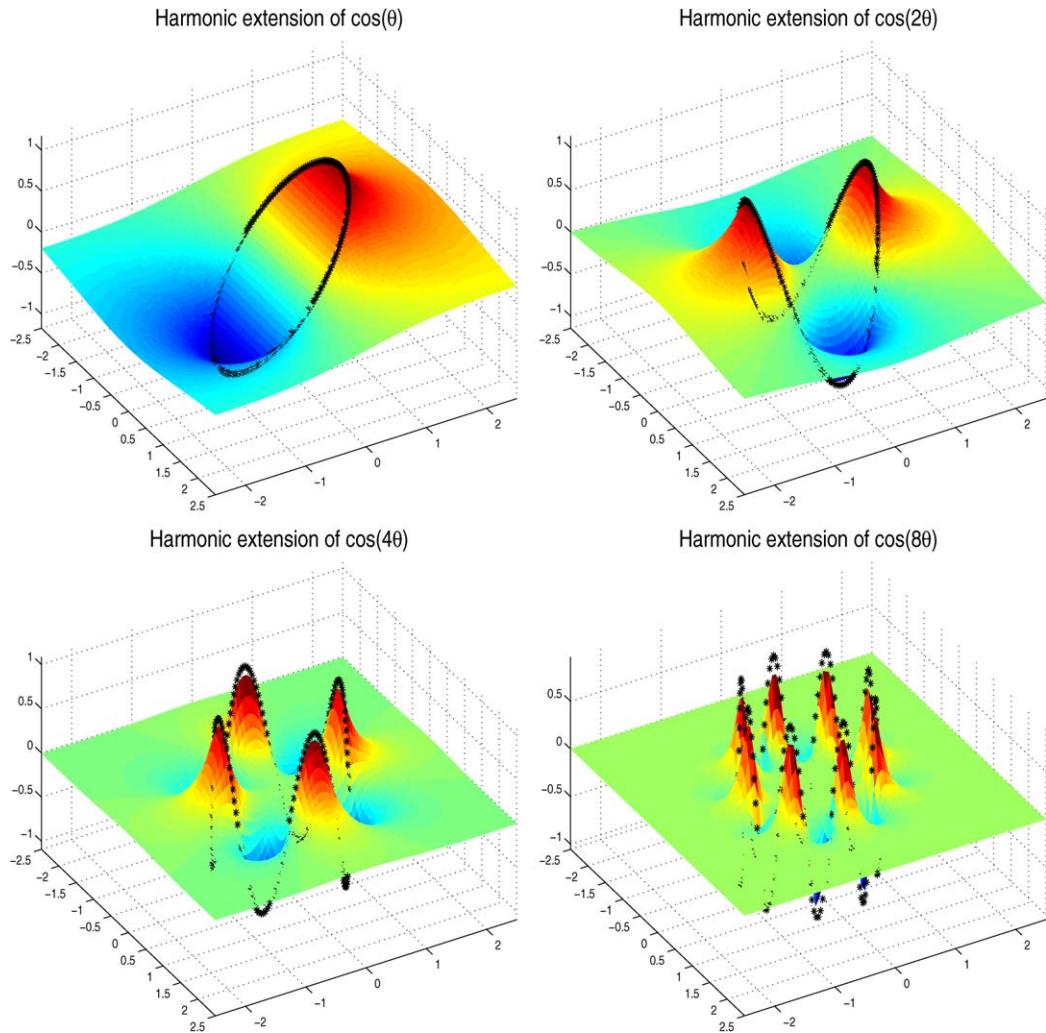


Fig. 2. Harmonic extensions of the functions $f(\theta) = \cos(j\theta)$ for $j = 1, 2, 4, 8$, from the unit circle to the plane. When the number of the oscillations increases, so does the order of the multipole extending f , and the extension decays faster.

since f is extended as a sum of elementary potentials generated by charges along X , it is roughly extended as the potential generated by a multipole of order p and therefore will decay as $\frac{1}{r^p}$, where r is the distance from the curve X . This relation between the intrinsic frequencies of f and its domain of extension is illustrate in Fig. 2 and is further investigated in Section 4.

3.4. Wavelet extension

Let $\{V_j\}_{j \in \mathbb{Z}}$ be a multiresolution analysis in \mathbb{R}^n and let \mathcal{H} be the space of scaling functions V_j at a fixed scale 2^{-j} . Let X be a subset of \mathbb{R}^n . One can define geometric harmonics corresponding to this space and use them to construct scaling functions adapted to the geometry of X and to compute extensions at a given scale.

We illustrate this idea in the simple case of the Haar multiresolution. Let Φ be the indicator function of the unit cube in \mathbb{R}^n and let X be a finite length curve. Therefore, for $\mathcal{H} = V_j$, the reproducing kernel is given by

$$\sum_{m \in \mathbb{Z}^n} 2^{-nj} \Phi(2^j x - m) \Phi(2^j y - m)$$




Fig. 3. Extension of the function $f(\theta) = \theta$ using the Haar scaling function at different scales.

and when one restricts it to the curve X ,

$$k(x, y) = \sum_{m \in Q_X} 2^{-nj} \Phi(2^j x - m) \Phi(2^j y - m),$$

where Q_X is the set of indices in \mathbb{Z}^n associated with unit cubes intersecting X . This formula allows to conclude that the geometric harmonics are the indicator of the dyadic cubes at scale 2^{-j} that intersect the curve X and it can be checked that the eigenvalues are given by the quantities $2^{-nj} |Q \cap X|$, where $|Q \cap X|$ is the length of the piece of X that intersect a given dyadic cube Q at scale 2^{-j} . As an illustration, in Fig. 3 we show the extension of the function $f(\theta) = \theta$ from the circle to the plane, at different scales.

The Haar case was simple as these scaling functions, when restricted to any curve, remain orthogonal. For more general scaling functions, the geometric harmonics provide scaling functions adapted to the geometry of X . Note that the same kind of construction can be done with wavelet spaces W_j .

4. Multiscale extension

We now restrict our attention to the case when $\bar{X} = \mathbb{R}^n$ and X is a C^∞ and compact submanifold of dimension d . To simplify the exposition, we assume that the measure $d\mu$ is now the Riemannian measure dx on X . In this case,

two different Fourier (or Paley–Littlewood) analyses can be performed on a function f defined on X . The first one is purely intrinsic and is obtained by using the eigenfunctions of the Laplace–Beltrami operator Δ on X . We recall that these eigenfunctions are the analogue of the Fourier basis to arbitrary submanifolds. The spectrum of Δ is discrete (as X is compact) and corresponds to pure frequency modes. The other frequency analysis is that of the classical Fourier transform in \mathbb{R}^n and that can be applied to the various extensions of f .

In this section, we investigate the relation between these two analyses by exploring the action of the restriction and extension operators. In particular, we ask the question of whether the intrinsic diffusion of heat is equivalent to the extrinsic diffusion. We show that this is true, provided the embedding of the manifold X into \mathbb{R}^n is not too “complicated.” We also develop a multiscale extension scheme for functions defined on X . The key point that we will introduce is to consider *several* kernels, i.e., several scales for the extension. Finding the optimal scale is obtained by adapting the analyzing the frequency content of the function over the training set.

4.1. Restriction operator

Since the set X is a smooth submanifold, the restriction of a function F to X is a well-behaved operation, in the sense that if F is smooth, then so is its restriction f .

We can give a precise meaning of this statement using a space characterization of smoothness: suppose that F is differentiable with a bounded derivative, then f is obviously differentiable as a map from X into \mathbb{R} and its (intrinsic) gradient at a point $x \in X$ is nothing else but the orthogonal projection of gradient of F onto the tangent plane at that same point x .

The same idea can be characterized by a frequency argument. Consider plane waves in \mathbb{R}^n :

$$F_\xi(x) = e^{2i\pi \langle \xi, x \rangle}.$$

Let Δ be the Laplace–Beltrami operator on X that we can assume to be a curve for the sake of simplicity. Let $\{v_j^2\}$ and $\{\phi_j\}$ be its eigenfunctions (in nondecreasing order) and eigenvalues:

$$\Delta \phi_j = v_j^2 \phi_j.$$

Definition 8 (Intrinsic bandwidth). A function $f \in L^2(X, dx)$ defined on X is said to be *intrinsically bandlimited* if it can be written as

$$f = \sum_{j=0}^m c_j \phi_j$$

for some $m \geq 0$, with $c_m \neq 0$. The largest pure frequency v_m in this decomposition is called *intrinsic bandwidth* of f .

To relate the extrinsic bandwidth ξ to the intrinsic one, we have the following result.

Proposition 9. *Suppose that the curvature of X is bounded by a number $M > 0$. Then if $\|\xi\| > \frac{M}{\pi}$, the intrinsic spectrum of f_ξ decays exponentially according to*

$$|\langle \phi_j, f_\xi \rangle| \leq \sqrt{\mu(X)} \left(\frac{4\pi^2 \|\xi\|^2}{v_j^2} \right)^m$$

for all $m \geq 0$. As a conclusion, a function with only low extrinsic frequencies is also a function with low intrinsic frequencies when restricted to X , with approximately the same band.

Proof. Locally, on the curve around $x \in X$, the function F_ξ has the form

$$f_\xi(y) = F_\xi(y) = \exp(2i\pi (\langle \xi, x \rangle + u\xi_T + a(x)u^2\xi_N)),$$

where u is the local coordinate in the tangent plane of the point y , $a(x)$ is the scalar curvature at x , and ξ_T and ξ_N are the tangent and normal projections of ξ in the osculatory plane at x (see configuration in Fig. 4). A Taylor expansion yields:

$$f_\xi(y) = e^{2i\pi \langle \xi, x \rangle} (1 + 2i\pi \xi_T u + 2i\pi (a(x)\xi_N + i\pi \xi_T^2) u^2 + \dots).$$




Fig. 4. Geometric configuration.

Now, since for any function f defined on X , we have the Taylor expansion

$$f(y) = f(x) + u \frac{df}{ds}(x) + \frac{1}{2}u^2 \frac{d^2f}{ds^2}(x) + \dots$$

and since $\Delta = \frac{d^2}{ds^2}$, we identify

$$\Delta f_\xi(x) = 4i\pi (a(x)\xi_N + i\pi \xi_T^2) f_\xi(x).$$

Therefore, if the curvature is bounded by $M > 0$, we have the trivial estimate for $\|\xi\| > \frac{M}{\pi}$,

$$\|\Delta f_\xi\|_X \leq 4\pi^2 \|\xi\|^2 \|f_\xi\|_X.$$

In fact it is easily seen that for all $m \geq 0$,

$$\|\Delta^m f_\xi\|_X \leq (4\pi^2 \|\xi\|^2)^m \|f_\xi\|_X.$$

Since

$$\Delta^m f_\xi(x) = \sum_{j \geq 0} v_j^{2m} \langle \phi_j, f \rangle \phi_j(x)$$

and

$$\|f_\xi\|_X^2 \leq \mu(X),$$

we must have, by Parseval,

$$\sum_{j \geq 0} v_j^{4m} |\langle \phi_j, f_\xi \rangle|^2 \leq \mu(X) (4\pi^2 \|\xi\|^2)^{2m}.$$

In particular,

$$|\langle \phi_j, f_\xi \rangle| \leq \sqrt{\mu(X)} \left(\frac{4\pi^2 \|\xi\|^2}{v_j^2} \right)^m$$

for all $m \geq 0$. Therefore the coefficients of expansions in the eigenfunctions of the Laplace–Beltrami operator are negligible, except for finitely many, namely those for which the eigenvalue v_j^2 is less than $4\pi^2 \|\xi\|^2$. \square

4.2. Extension operator

The extension algorithm described in Section 2.3 is a two-step procedure

- the function f on the data is first prefiltered by a projection on the geometric harmonics that numerically admit an extension,
- then the extension is computed.

However, we know that not all functions f defined on X are (η, δ) -extendable. This means that some functions cannot be extended to some bandwidth, but for a given (η, δ) and any function f , if we choose the band B sufficiently large, f will be (η, δ) -extendable as a B -bandlimited function. This leads to the natural definition for the notion of extrinsic bandwidth as the infimum over all such $B > 0$. However, we prefer the following alternate definition which is easier to deal with:

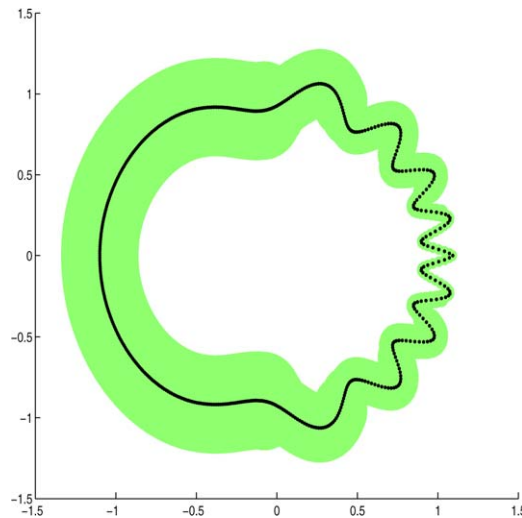


Fig. 5. Plot of r_x as a function of x , and a fixed value of $j = 10$ and $t = 2$. In regions where the embedding is “complicated,” the domain of extension is thinner.

Definition 10 (*Extrinsic bandwidth*). For a fixed value of $\varepsilon > 0$, the extrinsic bandwidth of a function f defined on X is the infimum of all $B > 0$ such that there exists a function F defined on \mathbb{R}^n satisfying:

- F is an extension of f , i.e.,

$$F(x) = f(x) \quad \text{for all } x \in X,$$

- F can be approximated to relative precision ε by a B -bandlimited function G on \mathbb{R}^n :

$$\frac{(\int_{\mathbb{R}^n} |F(x) - G(x)|^2 dx)^{\frac{1}{2}}}{(\int_{\mathbb{R}^n} |F(x)|^2 dx)^{\frac{1}{2}}} \leq \varepsilon.$$

For the study of the restriction operator, we looked at the restriction of the Fourier modes in \mathbb{R}^n , namely the plane waves. For the extension problem, it is natural to extend the Fourier modes on X , i.e., the set of eigenfunctions of the Laplace–Beltrami operator on X . In order to relate the intrinsic frequencies that these functions represent to the extrinsic spectrum (provided by the Fourier analysis in \mathbb{R}^n), it can be instructive to compute the extrinsic bandwidth of ϕ_j , or equivalently how far ϕ_j can be extended away from the set.

A simple example displayed in Fig. 5 shows that if the embedding of the set X in \mathbb{R}^n is complicated, then the extrinsic bandwidth can be very different from the intrinsic bandwidth. The curve X was chosen to have low frequencies on the left part and steep variations on the right part. We considered the following intrinsic wave packets:

$$w_{j,x}^t(y) = w_{x,t}(y)\phi_j(y),$$

where $w_{x,t}$ is a window function centered at x and of width \sqrt{t} . For instance,

$$w_{x,t}(y) = \exp(-d^2(x, y)/t)$$

is a Gaussian window along the curve, where $d(x, y)$ is the geodesic distance between x and y . In some sense, $w_{j,x}^t$ defines an intrinsic local cosine waveform: x represents the time location parameter, j is the intrinsic frequency location, and \sqrt{t} is the time–width. The function $w_{j,x}^t$ can be extended to a distance r_x that depends on the intrinsic frequency j as well as the local complexity of the embedding of X around x . Figure 5 shows r_x as a function of x , where x is located on the manifold (for fixed value of $j = 10$ and $t = 2$). This plot shows that in regions where the embedding is “complicated,” the domain of extension is thinner. In the following, we prove that $r_x = C_x j$, where C_x measures the complexity of the embedding of X around x .

Clearly, this elementary example shows that the extrinsic bandwidth can be much larger than the intrinsic one, especially at locations where the curve has wild oscillations. In the following proposition, we show that a function f

with intrinsic bandlimit B on X admits an approximate extension that is a bandlimited function with band CB , where C is some universal constant that depends on the geometry of X . In order to adopt a broader point of view, instead of considering the eigenfunctions of the Laplace–Beltrami operator Δ , we will study the extension of eigenfunctions of the following elliptic operator:

$$\mathbf{\Delta} = \Delta + Q,$$

where Q is a bounded potential function. In [7], the authors show that this type of differential operator arises naturally as the small-scale limit of several families of kernel operators. We keep the same notations for the eigenfunctions and eigenvalues, namely,

$$\mathbf{\Delta}\phi_j = v_j^2\phi_j.$$

We assume that the curvature of X is bounded by a number M .

Proposition 11. *Let $\varepsilon > 0$ be a preset accuracy. There exists a constant $C > 0$ such that for all $j \geq 0$, one can construct a function F_j defined on \mathbb{R}^n satisfying:*

- F_j is an extension of ϕ_j , i.e.,

$$F_j(x) = \phi_j(x) \quad \text{for all } x \in X,$$

- F_j can be approximated to relative precision ε by a bandlimited function G_j of band Cv_j :

$$\frac{(\int_{\mathbb{R}^n} |F_j(x) - G_j(x)|^2 dx)^{\frac{1}{2}}}{(\int_{\mathbb{R}^n} |F_j(x)|^2 dx)^{\frac{1}{2}}} \leq \varepsilon.$$

In particular, this proposition asserts that the extrinsic bandwidth is less than or equal to the intrinsic bandwidth times a constant C that depends on the precision ε and on the geometry of X . In some way, it is a measure of the complexity of the embedding of X into \mathbb{R}^n . For instance, in the example of Fig. 5, C can be quite large because of the oscillations of the curve.

Proof. If $v_j = 0$, then the result is trivial. We now assume that $v_j > 0$.

For any $x \in \mathbb{R}^n$, let $x' \in X$ be any point in \mathbb{R}^n such that

$$\|x - x'\| = \inf_{y \in X} \|x - y\|.$$

The point x' is not necessarily uniquely defined, which makes this definition ambiguous. However, since X is assumed to be C^∞ and with bounded curvature, in a tubular neighborhood \tilde{X} of X , $x \mapsto x'$ is an actual C^∞ function. For $u \in X$, let $r(u) > 0$ be the radius of the corresponding slice of \tilde{X} that is orthogonal to X . The maximum radius $R = \max_{u \in X} r(u)$ of \tilde{X} will be later adjusted (decreased) in order to account for the curvature bound M .

Let $W : \mathbb{R}^n \rightarrow \mathbb{R}$ be a C^∞ window function localized around X , that is, $W(x) = 1$ if $x \in X$, $W(x) = 0$ if $x \notin \tilde{X}$ and $0 \leq W(x) \leq 1$ for all $x \in \mathbb{R}^n$. Define F_j by

$$F_j(x) = W(x)e^{-v_j^2\|x-x'\|^2} \phi_j(x').$$

The function F_j is a C^∞ extension of ϕ_j to \mathbb{R}^n that numerically vanishes at a distance proportional to $\min(1/v_j, R)$. To estimate the decay of its spectrum, we need to bound its extrinsic derivatives $\nabla^m F_j(x)$ (tensor of order m). The Leibnitz formula implies that

$$\begin{aligned} \|\nabla^m F_j(x)\| &\leq \sum_{q=0}^m \binom{m}{q} \|\nabla^{m-q} W(x)\| \cdot \|\nabla^q (e^{-v_j^2\|x-x'\|^2} \phi_j(x'))\| \\ &\leq \sum_{q=0}^m K_{m,q} \|\nabla^q (e^{-v_j^2\|x-x'\|^2} \phi_j(x'))\| \cdot 1_{\{x \in \tilde{X}\}}, \end{aligned}$$

where the last inequality was obtained by noting that W and its derivatives are uniformly bounded and are supported on \tilde{X} .

The triangle inequality yields

$$\left(\int_{\mathbb{R}^n} \|\nabla^m F_j(x)\|^2 dx \right)^{\frac{1}{2}} \leq \sum_{q=0}^m K_{m,q} \left(\int_{\tilde{X}} \|\nabla^q (e^{-v_j^2 \|x-x'\|^2} \phi_j(x'))\|^2 dx \right)^{\frac{1}{2}}. \tag{1}$$

Another application of the Leibnitz formula gives, for $x \in \tilde{X}$,

$$\|\nabla^q (e^{-v_j^2 \|x-x'\|^2} \phi_j(x'))\| \leq \sum_{i=0}^q \binom{q}{i} \|\nabla^i (e^{-v_j^2 \|x-x'\|^2})\| \cdot \|\nabla^{q-i}(\phi_j(x'))\|.$$

Therefore, by the triangle inequality

$$\left(\int_{\tilde{X}} \|\nabla^q (e^{-v_j^2 \|x-x'\|^2} \phi_j(x'))\|^2 dx \right)^{\frac{1}{2}}$$

is bounded from above by

$$\sum_{i=0}^q \binom{q}{i} \left(\int_{\tilde{X}} \|\nabla^i (e^{-v_j^2 \|x-x'\|^2})\|^2 \cdot \|\nabla^{q-i}(\phi_j(x'))\|^2 dx \right)^{\frac{1}{2}}. \tag{2}$$

To evaluate each term of this bound, we use the following lemma:

Lemma 12. *Let f_v be a function on \tilde{X} of the form*

$$f_v(x) = g(v\|x - x'\|)h_v(x'),$$

where $g : [0, +\infty) \rightarrow \mathbb{R}$ has an exponential decay. Again, let $M > 0$ be a bound on the curvature of X . Then, as $v \rightarrow +\infty$,

$$\int_{\tilde{X}} |f(x)| dx \asymp v^{-(n-d)} \int_0^{+\infty} |g(r)| r^{n-d-1} dr \int_X |h_v(u)| du.$$

Proof. The proof of this lemma relies on a change of variable obtained by reparametrizing \tilde{X} as follows: the slice of \tilde{X} that is orthogonal to X at $u \in X$ is a ball of radius $r(u)$ in the subspace orthogonal to the tangent space of X at u . Therefore any point in this slice can be parameterized by the ball of radius $r(u)$ in \mathbb{R}^{n-d} . In addition, the collection of slices is parametrized by $u \in X$. Therefore we apply the change of variable $x \mapsto (u, t)$ from \tilde{X} to $X \times \mathbb{R}^{n-d}$.

Let $J(u, t)$ denote the Jacobian of the change of variable $(u, t) \mapsto x$. The lemma follows from the fact that J is bounded from below and above for all $x \in \tilde{X}$, provided that $R = \max_{u \in X} r(u)$ be chosen small enough. Indeed, first, a variation dt of t entails the same variation of x . Second, a variation du in the tangent plane at u entails a variation of x of order $(1 \pm 2\alpha(u)\|t\|) du$ in the tangent direction (where $\alpha(u)$ is the curvature of X at u), and of order $\|t\| du^2$ in the normal direction. To conclude, since $\|t\| < r(u) \leq R$ and $|\alpha(u)| < M$, it suffices to choose R so that $2RM < \frac{1}{2}$, and we obtain that $1 - 2\alpha(u)\|t\| > 1 - \frac{1}{2}$ and $1 + 2\alpha(u)\|t\| < 1 + \frac{1}{2}$. Finally,

$$\frac{1}{2} < J(u, t) < \frac{3}{2}.$$

Therefore, with the same constants,

$$\int_{\tilde{X}} |f_v(x)| dx \asymp \int_X |h_v(u)| \int_{t \in \mathbb{R}^{n-d}, \|t\| \leq r(u)} g(v\|t\|) dt du.$$

By going to polar coordinates $r = v\|t\|$, we obtain

$$\int_{\tilde{X}} |f_v(x)| dx \asymp v^{-(n-d)} \int_X |h_v(u)| \int_0^{vr(u)} g(r)r^{n-d-1} dr du.$$

Last, due to the decay of g and the fact that $v \rightarrow +\infty$, up to exponentially small terms, the inner integral can be computed on the half-line $[0, +\infty)$, which ends the proof of the lemma. \square

Going back to the proof of the proposition, the lemma implies that

$$\int_{\tilde{X}} \|\nabla^i (e^{-v_j^2\|x-x'\|^2})\|^2 \cdot \|\nabla^{q-i}(\phi_j(x'))\|^2 dx \leq K'_{q,i} v_j^{2i-(n-d)} \int_X \|\nabla^{q-i}\phi_j(u)\|^2 du. \tag{3}$$

What remains to be done is to bound the $L^2(X)$ -norm of the ambient derivatives of ϕ_j . To do so, we start by bounding the norm of the intrinsic derivatives:

Lemma 13. *For all $s \geq 0$, there exists C'_s such that*

$$\|\phi_j\|_s = \left(\sum_{|\alpha| \leq s} \int_X |\partial^\alpha \phi_j(u)|^2 du \right)^{\frac{1}{2}} \leq C'_s v_j^i.$$

This lemma follows from the classical theory of elliptic operators, which says that since we can bound the norm of $\Delta^k \phi_j$, we have a bound on all derivatives of order less than or equal to $2k$.

Let

$$C_Q = \sup_{u \in X} |Q(u)|.$$

For $s = 0$, the lemma is trivial, and for $s = 1$, it results from an integration by parts as

$$\begin{aligned} \|\nabla \phi_j\|_X^2 &= \langle \Delta \phi_j, \phi_j \rangle_X \quad \text{by the Stokes formula} \\ &= \langle \mathbf{\Delta} \phi_j, \phi_j \rangle_X - \langle Q \phi_j, \phi_j \rangle_X \\ &\leq v_j^2 + C_Q \end{aligned}$$

and therefore

$$\|\phi_j\|_1^2 = \|\phi_j\|_X^2 + \|\nabla \phi_j\|_X^2 = 1 + v_j^2 + C_Q \leq C_1'^2 v_j^2.$$

In [8, p. 262], it is shown that if L is an elliptic operator of order k , then for all $i \geq k$ and all f defined on X , we have

$$\|f\|_i \leq C'(\|Lf\|_{i-k} + \|f\|_{i-1}). \tag{4}$$

We can now proceed by induction:

- for $i = k = 2s$ and $L = \mathbf{\Delta}^s$, identity (4) yields

$$\|\phi_j\|_{2s} \leq C'(\|\mathbf{\Delta}^s \phi_j\|_X + \|\phi_j\|_{2s-1}) \leq C'(v_j^{2s} + C'_{2s-1} v_j^{2s-1}) \leq C'_{2s} v_j^{2s},$$

- for $i = 2s + 1, k = 2s$ and $L = \mathbf{\Delta}^s$, identity (4) becomes

$$\|\phi_j\|_{2s+1} \leq C'(\|\mathbf{\Delta}^s \phi_j\|_1 + \|\phi_j\|_{2s}).$$

We have

$$\begin{aligned} \|\Delta^s \phi_j\|_1^2 &= \|\Delta^s \phi_j\|_X^2 + \|\nabla \Delta^s \phi_j\|_X^2 \quad \text{by definition} \\ &= v_j^{4s} + \langle \Delta \Delta^s \phi_j, \Delta^s \phi_j \rangle_X \quad \text{by the Stokes formula} \\ &= v_j^{4s} + \langle \Delta^{s+1} \phi_j, \Delta^s \phi_j \rangle_X - \langle Q \Delta^s \phi_j, \Delta^s \phi_j \rangle_X \\ &\leq v_j^{4s} + v_j^{2(2s+1)} + C_Q v_j^{4s} \end{aligned}$$

and finally,

$$\|\phi_j\|_{2s+1} \leq C' \left(\sqrt{(1 + C_Q) v_j^{4s} + v_j^{2(2s+1)}} + C'_{2s} v_j^{2s} \right) \leq C'_{2s+1} v_j^{2s+1}.$$

The lemma is now proven, and to finish the proof of the proposition, we need to bound the extrinsic (ambient) gradient of $x \mapsto \phi_j(x')$ on X using the intrinsic (covariant) derivatives of ϕ_j . Note that this function is constant in directions orthogonal to X . Applying the chain rule, we see that the extrinsic derivatives of this function can be expressed as a sum of products of derivatives of $x \mapsto x'$ and intrinsic derivatives of ϕ_j . The first type of derivative is uniformly bounded since the curvature of X is itself bounded, and as a consequence,

$$\left(\int_{\mathbb{R}^n} \|\nabla^{q-i} \phi_j(u)\|^2 du \right)^{\frac{1}{2}} \leq C_{q,i} \|\phi_j\|_{q-i} \leq C'_{q,i} v_j^{q-i}.$$

Combining this equation with Eqs. (1), (2) and (3), we obtain

$$\left(\int_{\mathbb{R}^n} \|\nabla^m F_j(x)\|^2 dx \right)^{\frac{1}{2}} \leq C_m v_j^{m - \frac{n-d}{2}}.$$

Now for a fixed value of m , define

$$\hat{G}_j(\xi) = \begin{cases} \hat{F}_j(\xi) & \text{if } \|\xi\| < C v_j, \\ 0 & \text{otherwise,} \end{cases}$$

then, by the Parseval identity, we have

$$\begin{aligned} \int_{\mathbb{R}^n} |F_j(x) - G_j(x)|^2 dx &= \int_{\|\xi\| > C v_j} |\hat{F}_j(\xi)|^2 d\xi \leq \int_{\|\xi\| > C v_j} |\hat{F}_j(\xi)|^2 \frac{\|\xi\|^{2m}}{(C v_j)^{2m}} d\xi \leq \frac{1}{(C v_j)^{2m}} \int_{\mathbb{R}^n} \|\nabla^m F_j(x)\|^2 dx \\ &\leq \frac{C_m^2}{C^{2m}} v_j^{-(n-d)}. \end{aligned}$$

Last, one can check that there exists a constant K' such that

$$\int_{\mathbb{R}^n} |F_j(x)|^2 dx \geq K' \int_{\tilde{X}} |e^{-v_j^2 \|x-x'\|^2} \phi_j(x')|^2 dx$$

and thus, from Lemma 12, we have

$$\int_{\mathbb{R}^n} |F_j(x)|^2 dx \geq K^2 v_j^{-(n-d)}$$

for some $K > 0$, and we merely have to pick C so that

$$\frac{C_m}{K C^m} < \varepsilon. \quad \square$$

Recall that \mathbf{E}_B is the bandlimited extension operator corresponding to band B . From now on, we set $B_j = C v_j$. The consequence of this proposition is that, because of its optimal property, the extension provided by \mathbf{E}_{B_j} must have an energy on \mathbb{R}^n that is less than or equal to that of the extension that we constructed in the proof above. This means that the numerical support of $\mathbf{E}_{B_j} \phi_j$ will be included in a tube of radius proportional to $\frac{1}{v_j}$ around X . Theoretically, it could be much thinner, but because of the Heisenberg principle, then the support cannot really be smaller:

Lemma 14. *The standard deviation of the extension $\mathbf{E}_{B_j}\phi_j$ along any normal direction to X is at least equal to $\frac{C'}{v_j}$ for some $C' > 0$ independent of v_j .*

Proof. Let f be the restriction of $\mathbf{E}_{B_j}\phi_j$ on a line that is normal to X . Then f is a univariate bandlimited function of band $C\sqrt{d}v_j$. Let

$$\text{Var}(f) = \frac{\int_{\mathbb{R}}(x - \bar{x})^2|f(x)|^2 dx}{\int_{\mathbb{R}}|f(x)|^2 dx}$$

and

$$\text{Var}(\hat{f}) = \frac{\int_{\mathbb{R}}(\xi - \bar{\xi})^2|\hat{f}(\xi)|^2 d\xi}{\int_{\mathbb{R}}|\hat{f}(\xi)|^2 d\xi}$$

be the variances of f in the space and frequency domains, \bar{x} and $\bar{\xi}$ being the corresponding means. Then since f is bandlimited,

$$\int_{\mathbb{R}}\xi^2|\hat{f}(\xi)|^2 d\xi \leq (C\sqrt{d}v_j)^2 \int_{\mathbb{R}}|\hat{f}(\xi)|^2 d\xi,$$

and consequently, $\text{Var}(\hat{f}) \leq (C\sqrt{d}v_j)^2$ (the variance is always smaller than the second moment). The Heisenberg uncertainty principle implies that $\text{Var}(f) \geq C'^2v_j^{-2}$. \square

As a conclusion, the extension operation satisfies a certain version of the Heisenberg principle relating the spectrum of the operator Δ to the space and frequency localizations of the extensions of its eigenfunctions ϕ_j . This principle says that if $\Delta\phi_j = v_j^2\phi_j$, then the operator \mathbf{E}_{B_j} extends ϕ_j to a bandlimited function of band $\mathcal{O}(v_j)$ and localized in a tube of radius $\mathcal{O}(\frac{1}{v_j})$ around X . It is worthy to mention that similar results can be obtained for, say, Gaussian kernels.

4.3. Multiscale extension scheme

As a consequence of the previous section, where we have related the intrinsic and extrinsic Fourier analysis of a function on X , the size of the domain of extension and the bandwidth to which a given function f can be extended depends on int intrinsic spectrum. In particular, we know that each eigenfunction ϕ_j of Δ can be extended as a bandlimited function with bandwidth Cv_j^2 , at distance proportional to $\frac{1}{v_j}$ from X . This observation gives rise to a natural multiscale extension technique that we now describe.

Algorithm 2 (*Multiscale extension scheme*). Fix a precision η and a condition number $\frac{1}{\delta}$.

- Precomputation phase: for each eigenfunction ϕ_j of Δ , compute the minimal frequency band B_j to which it can be extended using \mathbf{E}_{B_j} . This step can also be done by grouping the eigenfunctions in dyadic packets and by computing a band for each packet.
- Extension phase: for any function f defined on X , compute its decomposition over $\{\phi_j\}$ and retain enough coefficients so that the relative error is of order η :

$$f = \sum_{j \in \mathcal{S}} \langle f, \phi_j \rangle_X \phi_j + \mathcal{O}(\eta \|f\|_X)$$

and use the precomputed extensions of ϕ_j to extend f as

$$F = \sum_{j \in \mathcal{S}} \langle f, \phi_j \rangle_X \mathbf{E}_{B_j}\phi_j.$$

This algorithm extends f as a linear combination of “atoms” with different localizations in time and frequency. More precisely, F is a sum of functions that oscillate at intrinsic frequency v_j^2 on X and that vanish at distance $\frac{1}{v_j}$

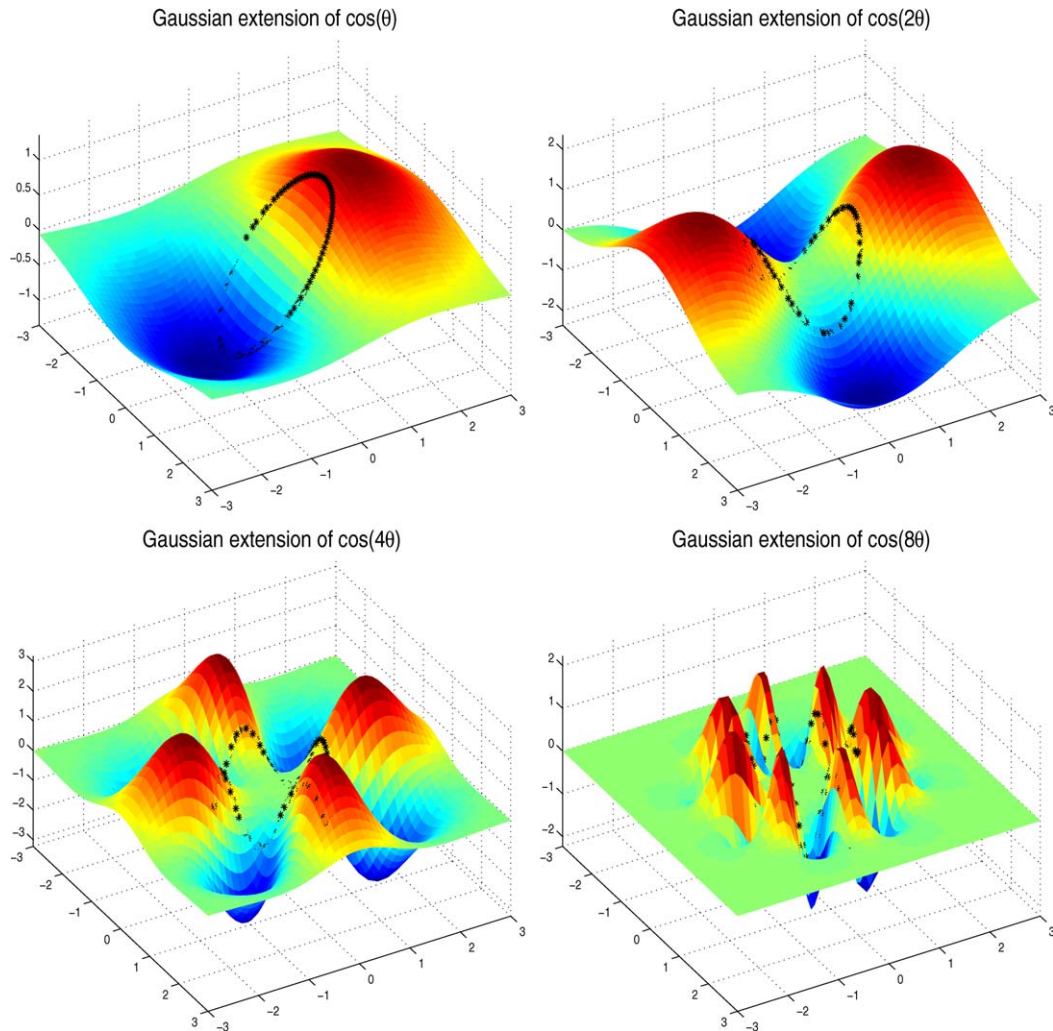


Fig. 6. Gaussian extensions of the functions $\cos(j\theta)$ for $j = 1, 2, 4, 8$, from the unit circle to the plane. Unlike the bandlimited extensions, these ones are much more localized in the plane and they do not oscillate off the circle. These plots also illustrate the Heisenberg principle as functions with high frequencies are extended as functions with a small support.

from this set. From the implementation point of view, it is to be noted that in [7], the authors provide an efficient way to compute the eigenfunctions of Δ on a submanifold X .

We illustrate the space–frequency localization principle in Fig. 6 where different Fourier modes of the unit circle are extended to the plane using Gaussian kernels with a scale adapted to intrinsic frequencies.

5. Conclusion

We have described the construction and properties of the geometric harmonics and we have shown how they can be used to perform out-of-sample extension of empirical function. The study of the case of submanifolds shows that these functions can be viewed as a simple yet powerful tool for relating the intrinsic and extrinsic Fourier analyses. The key point is to consider a multiscale approach and to adapt the scale of extension to the frequency content of the empirical function. This provides an answer to the problem of determining the scale of extension of empirical functions. In addition, there is no reason to consider a global scale of extension as the local complexity of the embedding of X may vary a lot along the manifold. A local version of our approach can be easily developed by using local windows.

Acknowledgments

We thank Ann Lee, Mauro Maggioni, Boaz Nadler and Naoki Saito for interesting discussions and comments regarding this work. We also express our thanks to the reviewers for their useful suggestions.

Appendix A. Reproducing kernel Hilbert spaces

In this section, we provide the basic background concerning reproducing kernel Hilbert spaces, and their connection to positive semi-definite kernels. A classical reference concerning this topic is [2].

Definition 15 (*Reproducing kernel Hilbert space*). A space Hilbert \mathcal{H} of functions defined on a set \bar{X} is said to be a reproducing kernel Hilbert space if there exists a function (“kernel”) $k: \bar{X} \times \bar{X} \rightarrow \mathbb{R}^n$ such that

- for almost every $x \in \bar{X}$, $k(x, \cdot) \in \mathcal{H}$,
- if $\langle \cdot, \cdot \rangle_{\bar{X}}$ is the inner product in \mathcal{H} , then for all function $f \in \mathcal{H}$ and almost all $x \in \bar{X}$,

$$\langle f, k(x, \cdot) \rangle_{\bar{X}} = f(x). \quad (\text{A.1})$$

In [2], it is shown that the concepts of reproducing kernels and positive kernels are identical in the following sense: any reproducing kernel is positive semi-definite and to any positive kernel k there corresponds a reproducing kernel Hilbert space \mathcal{H} for which k is the reproducing kernel. The construction is given in [2, pp. 150–151].

In the particular case when $\bar{X} = \mathbb{R}^n$ and $k(x, y) = h(x - y)$, then more can be said. By Bochner’s theorem, we know that the Fourier transform \hat{h} of h is a finite positive measure. Assuming that this measure has the form $\hat{h}(\xi) d\xi$, then it can be checked that \mathcal{H} is the space of functions f defined on \mathbb{R}^n such that

$$\int_{\hat{h}(\xi) > 0} |\hat{f}(\xi)|^2 \frac{d\xi}{\hat{h}(\xi)} < +\infty$$

and such that $\hat{f}(\xi) = 0$ if $\hat{h}(\xi) = 0$. Consequently, this space is the Fourier transform of a weighted L^2 space, where the weight penalizes high frequencies since \hat{h} is integrable.

In the case of the Bessel kernel, the function \hat{h} is the indicator of a ball, and therefore \mathcal{H} is imply a space bandlimited of bandlimited functions endowed with the classical inner product of $L^2(\mathbb{R}^n, dx)$.

Appendix B. Expression of the Bessel kernel

In what follows, we derive the form of the kernel corresponding to functions whose Fourier transform is the indicator function of the ball of radius $\frac{B}{2}$ centered at the origin, namely:

$$k_B^{(n)}(x, y) = \int_{\|\xi\| < \frac{B}{2}} e^{2i\pi \langle \xi, x-y \rangle} d\xi = \left(\frac{B}{2}\right)^{\frac{n}{2}} \frac{J_n(\pi B \|x - y\|)}{\|x - y\|^{\frac{n}{2}}},$$

where J_ν is the Bessel function of the first kind of order ν . This kernel will be termed “Bessel kernel.”

Since the kernel is really a function of $\|x - y\|$, we are looking for the form of the Fourier transform of the indicator of the unit ball in dimension n . To do so, we make use of a result known under the name of the Bochner–Coifman–Howe periodicity relations:

Lemma 16. *Let f be a radial function, and let $\mathcal{F}_n f(\xi) = h_n(\|\xi\|^2)$ be its Fourier transform in dimension n . Then the Fourier transforms of f in dimension n and $n + 2$ are related in the following manner:*

$$h_{n+2}(u) = -\frac{1}{\pi} h_n'(u).$$

In other words, to compute the Fourier transform $h_{n+2}(\|\xi\|^2)$ of f in \mathbb{R}^{n+2} , one can start from the Fourier transform $h_n(\|\xi\|^2)$ in dimension n , view this function as a function of $\|\xi\|^2$ and compute its derivative in this variable.

Proof. Since any radial function or tempered distribution can be approximated as a sum of Gaussians, one merely needs to verify the relation for $f(x) = e^{-\alpha r^2}$. In this case,

$$\mathcal{F}_n f(\xi) = \left(\frac{\pi}{\alpha}\right)^{\frac{n}{2}} e^{-\frac{\pi^2 \xi^2}{\alpha}}.$$

Thus

$$h_n(u) = \left(\frac{\pi}{\alpha}\right)^{\frac{n}{2}} e^{-\frac{\pi^2 u}{\alpha}}$$

and the identity is satisfied. \square

Using this lemma we can now conclude:

Proposition 17. *In dimension n , the Bessel kernel has the following form:*

$$k_B^{(n)}(x, y) = \left(\frac{B}{2}\right)^{\frac{n}{2}} \frac{J_{\frac{n}{2}}(\pi B \|x - y\|)}{\|x - y\|^{\frac{n}{2}}}.$$

Moreover, if n is odd, then the simpler formula can be used:

$$k_B^{(n)}(x, y) = M_{B,n} \left(\frac{1}{r} \frac{d}{dr}\right)^{\frac{n-1}{2}} \text{sinc}(Br),$$

where $r = \|x - y\|$ and

$$M_{B,n} = \left(\frac{B}{2}\right)^{\frac{n}{2}} \sqrt{2B} (-1)^{\frac{n-1}{2}}.$$

Proof. By a trivial scaling argument, we may assume that $B = 2$.

Then if $n = 1$, then

$$k_2^{(1)}(x, y) = \int_{-1}^1 e^{2i\xi\pi(x-y)} d\xi = 2\text{sinc}(2\|x - y\|) = \frac{J_{\frac{1}{2}}(2\pi\|x - y\|)}{\|x - y\|^{\frac{1}{2}}},$$

where the third equality is obtained using 10.1.1 and 10.1.11 in [1].

If $n = 2$, then in polar coordinates (ρ, θ) :

$$\begin{aligned} \int_{\|\xi\| < 1} e^{2i\pi(\xi, x-y)} d\xi &= \int_0^1 \int_0^{2\pi} e^{2i\pi r \rho \cos \theta} d\theta d\rho \\ &= 2\pi \int_0^1 J_0(2\pi r \rho) \rho d\rho \quad \text{by 9.1.21 in [1]} \\ &= \frac{1}{2\pi r^2} \int_0^{2\pi r} u J_0(u) du \\ &= -\frac{1}{r} J_0'(2\pi r) \quad \text{since by 9.1 in [1] } (u J_0(u))' = -u J_0(u) \\ &= \frac{J_1(2\pi r)}{r} \quad \text{by 9.1.28 in [1].} \end{aligned}$$

For higher orders we proceed by induction on n , noting that if

$$h(u) = \frac{J_{\frac{n-2}{2}}(2\pi\sqrt{u})}{u^{\frac{n-2}{4}}},$$

then

$$\begin{aligned} h'(u) &= \frac{J'_{\frac{n-2}{2}}(2\pi\sqrt{u}) \frac{\pi}{\sqrt{u}} u^{\frac{n-2}{4}} - J_{\frac{n-2}{2}}(2\pi\sqrt{u}) \frac{n-2}{4} u^{\frac{n-2}{4}-1}}{u^{\frac{n-2}{2}}} \\ &= \frac{2\pi\sqrt{u} J'_{\frac{n-2}{2}}(2\pi\sqrt{u}) - \frac{n-2}{2} J_{\frac{n-2}{2}}(2\pi\sqrt{u})}{2u^{\frac{n-2}{4}+1}} \\ &= -\frac{\pi\sqrt{u} J_{\frac{n}{2}}(2\pi\sqrt{u})}{u^{\frac{n-2}{4}+1}} \quad \text{according 9.1.27 in [1]} \\ &= -\frac{\pi J_{\frac{n}{2}}(2\pi\sqrt{u})}{u^{\frac{n}{2}}}. \end{aligned}$$

Now invoking Lemma 16 yields the result. Finally, to obtain a formula in terms of the variable r instead of r^2 , notice that $d(r^2) = r dr$, and this implies that for odd values of n

$$k_2^{(n)}(x, y) = 2(-1)^{\frac{n-1}{2}} \left(\frac{1}{r} \frac{d}{dr} \right)^{\frac{n-1}{2}} \text{sinc}(2r). \quad \square$$

References

- [1] M. Abramowitz, I. Stegun, *Handbook of Mathematical Functions*, Dover Publications, New York, 1965.
- [2] N. Aronszajn, Theory of reproducing kernels, *Trans. Amer. Math. Soc.* 68 (3) (1950).
- [3] Y. Bengio, J.-F. Paiement, P. Vincent, Out-of-sample extensions for LLE, isomap, MDS, eigenmaps, and spectral clustering, Technical report 1238, Université de Montréal, 2003.
- [4] M. Belkin, P. Nyoigi, V. Sindhwani, On manifold regularization, in: *Proc. 10th Int. Workshop on Artificial Intelligence and Statistics*, Society for Artificial Intelligence and Statistics, pp. 17–24.
- [5] H. Wackernagel, *Multivariate Geostatistics*, Springer-Verlag, 2003.
- [6] I.J. Schoenberg, Metric spaces and completely monotone functions, *Ann. of Math.* (2) 39 (4) (1938) 811–841.
- [7] R.R. Coifman, S. Lafon, Diffusion maps, in preparation.
- [8] G.B. Folland, *Introduction to Partial Differential Equations*, Princeton Univ. Press, Princeton, NJ, 1976.
- [9] C. Fowlkes, S. Belongie, J. Malik, Efficient spatiotemporal grouping using the Nyström method, *IEEE Comput. Vision Pattern Recogn.*, December 2001.
- [10] W.H. Press, S.A. Teulkosky, W.T. Vetterling, B.P. Flannery, *Numerical Recipes in C*, Cambridge Univ. Press, New York, chap. 18.1.
- [11] D. Slepian, H.O. Pollak, Prolate spheroidal wave functions, Fourier analysis and uncertainty I, *Bell System Tech. J.* 40 (1961) 43–64.
- [12] D. Slepian, Prolate spheroidal wave functions, Fourier analysis and uncertainty IV: Extensions to many dimensions; generalized prolate spheroidal wave functions, *Bell System Tech. J.* 43 (1964) 3009–3058.
- [13] C. Williams, M. Seeger, Using the Nyström method to speed up kernel machines, *Neural Inf. Process. Systems* 13 (2001) 682–688.