

A Unified Definition and Computation of Laplacian Spectral Distances

Giuseppe Patané

*Consiglio Nazionale delle Ricerche
Istituto di Matematica Applicata e Tecnologie Informatiche
Genova, Italy*

Abstract

Laplacian spectral kernels and distances (e.g., biharmonic, heat diffusion, wave kernel distances) are easily defined through a filtering of the Laplacian eigenpairs. They play a central role in several applications, such as dimensionality reduction with spectral embeddings, diffusion geometry, image smoothing, geometric characterisations and embeddings of graphs. Extending the results recently derived in the discrete setting [38, 39] to the continuous case, we propose a novel definition of the Laplacian spectral kernels and distances, whose approximation requires the solution of a set of inhomogeneous Laplace equations. Their discrete counterparts are equivalent to a set of sparse, symmetric, and well-conditioned linear systems, which are efficiently solved with iterative methods. Finally, we discuss the optimality of the Laplacian spectrum for the approximation of the spectral kernels, the relation between the spectral and Green kernels, and the stability of the spectral distances with respect to the evaluation of the Laplacian spectrum and to multiple Laplacian eigenvalues.

Keywords: Laplacian spectrum, spectral distances, spectral kernels, heat kernel, diffusion distances and geometry, shape and graph analysis

1. Introduction

Spectral kernels and distances are easily defined through a filtering of the Laplacian eigenpairs and include random walks [41], biharmonic [28, 45], heat diffusion [10, 11, 17, 25], wave [12] kernels and distances. Laplacian spectral distances have been ap-

plied to shape segmentation [19] and comparison [11, 17, 35, 50] with multi-scale and isometry-invariant signatures [33, 32, 31]. Main applications of the heat kernel include dimensionality reduction [7, 53, 16, 4], diffusion geometry [6, 47], graphs' embeddings [30] and analysis [18], shape comparison [9, 13, 43], data representation [54] and classification [34, 46, 49].

Overview and contribution

Extending the results recently derived in the discrete setting [38, 39] to the continuous case, we propose a novel definition of the Laplacian spectral kernels and distances, whose approximation requires the solution of a set of inhomogeneous Laplace equations. Their discrete counterparts are equivalent to a set of sparse linear systems, which are efficiently solved with iterative methods. To this end, the Laplacian spectral distances $d^2(\mathbf{p}_i, \mathbf{p}_j) := \sum_{n=0}^{+\infty} |\phi_n(\mathbf{p}_i) - \phi_n(\mathbf{p}_j)|^2 / \rho^2(\lambda_n)$ are defined by filtering the Laplacian spectrum $(\lambda_n, \phi_n)_{n=0}^{+\infty}$, where $\rho : \mathbb{R}^+ \rightarrow \mathbb{R}$ is a proper filter function.

The proposed spectral kernels and distances are independent of the data dimensionality (e.g., surfaces, volumes, n -dimensional data), of the discretisation of the input domain, and of the Laplace-Beltrami operator. The Laplacian spectral distances generalise the bi-harmonic [28] ($\rho(s) := s^2$), wave [12] ($\rho_t(s) := \exp(ist)$), and diffusion [14, 21, 36] ($\rho_t(s) := \exp(st)$) distances, Mexican hat wavelets [23] ($\rho(s) := s^{1/2} \exp(s^2)$), and random walks [41]. The behaviour of the filter ρ , and in particular the convergence of $1/\rho$ to zero, determines the main properties of the resulting spectral distances, such as smoothness, encoding of local and global shape properties (e.g., Gaussian curvature, geodesic distance), localisation in time and frequency [21]. Through a proper filter design and a learning process [2, 8], we can obtain spectral kernels and distances that are also invariant with respect to isometric transformations or discriminative of a given shape class.

Recalling the definition of the spectral operator in [39, 37, 38], we review equivalent representations of the spectral kernels and distances in terms of the spectral norm of the δ -functions, the Laplacian spectrum, the spectral operator, and the spectral kernel (Sect. 2). Through these relations, the Laplacian spectral kernels and distances can

be computed by approximating the filter with a rational function and converting their evaluation to the solution of a set of differential equations that involve only the Laplace-Beltrami operator.

The proposed discrete approximation of the spectral kernels (Sect. 3) is equivalent to solving r sparse, symmetric, and well-conditioned linear systems, where r is the degree of the rational polynomial approximation of the filter. Applying an iterative linear solver, the computation of the spectral kernel takes $\mathcal{O}(r\tau(n))$ time, where $\tau(n)$ typically varies from $\tau(n) = n$ to $\tau(n) = n \log n$ and depends on the number n of shape samples and on the sparsity of the coefficient matrix. In a similar way, the computation of the spectral distance between two points has the same order of complexity of the evaluation of the corresponding spectral kernel at the same input points.

Through the spectrum-free approximation, the computed spectral kernel is not affected by the Gibbs phenomenon, as a consequence of the high accuracy of the rational approximation with respect to the low-pass filter associated with the truncated spectral approximation. Furthermore, the spectrum-free approximation is free of user-defined parameters (e.g., the number of Laplacian eigenpairs).

Finally (Sects. 4, 5), we experimentally verify the accuracy, efficiency, and numerical robustness of the definition and spectrum-free computation of the spectral kernels and distances.

Novelties with respect to previous work

With respect to previous work and our recent results on the definition and computation of discrete spectral distances [39, 38], the main novelties of this paper are

- the discussion of the *optimality of the Laplacian spectrum* for the approximation of the spectral kernel (Sect. 2.1);
- the study of the *relation between the spectral and Green kernels* (Sect. 2.4), associated with differential operators (Sect. 2.2);
- the *analysis of the stability of the spectral distances* with respect to the evaluation of the Laplacian spectrum and to multiple Laplacian eigenvalues (Sect. 2.3).

In fact, the evaluation of the Laplacian spectrum is generally affected by small perturbations of the Laplace-Beltrami operator;

- a *novel spectrum-free computation* of the spectral distances (Sect. 2.5), which is achieved by properly approximating the associated filter with a rational polynomial and is expressed in terms of the canonical basis;
- the *computation of the discrete spectral distances* (Sect. 3.2), which reduces to the solution of a set of sparse, symmetric, and well-conditioned linear systems. This analogy between the continuous and discrete cases shows the generality of the proposed approach and it has not been addressed by previous work. Furthermore, these results are complementary and more general than the spectrum-free approaches presented in [39, 38];
- a new *upper bound to the conditioning number* of the linear systems associated with the spectrum-free approximation, which confirms the numerical stability and well-conditioned computation of the class of the spectrum-free approaches (Sect. 3.1);
- new experiments (Sect. 4), which enrich the tests initially presented in [39, 38] and address the *robustness* of the computation of the Laplacian spectral distances and kernels with respect to the data resolution, partial sampling, geometric or topological noise, and deformations.

2. Laplacian spectral distances

Firstly, we present the Laplace-Beltrami operator, its spectrum, and the optimality of the Laplacian eigenbasis (Sect. 2.1). Recalling the definition of the spectral operator and kernel introduced in [39, 37, 38], we review equivalent representations of the spectral distances in terms of the spectral norm of the δ -functions, the Laplacian spectrum, the spectral operator, and the spectral kernel (Sect. 2.2). As novel contribution, we discuss the stability of the spectral distances with respect to the evaluation of the Laplacian spectrum and to multiple Laplacian eigenvalues, which are generally affected by small perturbations of the Laplace-Beltrami operator (Sect. 2.3). Finally, we introduce

a novel relation between the spectral and Green kernels (Sect. 2.4) and propose a novel spectrum-free computation of the spectral distances in the continuous case, which is presented in a way analogous to the discrete case [38, 39] and shows the generality of the proposed approach (Sect. 2.5).

2.1. Laplace-Beltrami operator

Let \mathcal{N} be a smooth surface, possibly with boundary, equipped with a Riemannian metrics and let us consider the *inner product* $\langle f, g \rangle_2 := \int_{\mathcal{N}} fg d\mu$ defined on the space $\mathcal{L}^2(\mathcal{N})$ of square integrable functions on \mathcal{N} and the corresponding norm $\|\cdot\|_2$. The *Laplace-Beltrami operator* Δ is *self-adjoint* $\langle \Delta f, g \rangle_2 = \langle f, \Delta g \rangle_2$, and *positive semi-definite* $\langle \Delta f, f \rangle_2 \geq 0, \forall f, g$, [44].

Laplacian eigenbasis and its optimality. Recalling that the Laplace-Beltrami operator is self-adjoint and positive semi-definite, it has an *orthonormal eigensystem* $(\lambda_n, \phi_n)_{n=0}^{+\infty}$, $\Delta \phi_n = \lambda_n \phi_n$, $\lambda_0 = 0$, $\lambda_n \leq \lambda_{n+1}$, in $\mathcal{L}^2(\mathcal{N})$. Expressing a function f in $\mathcal{L}^2(\mathcal{N})$ in terms of the Laplacian eigenfunctions, we have that

$$\begin{cases} f = \sum_{n=0}^{+\infty} \alpha_n \phi_n, & \alpha_n := \langle f, \phi_n \rangle_2; \\ \|f\|_2^2 = \sum_{n=0}^{+\infty} \alpha_n^2, & \|\nabla f\|_2^2 = \sum_{n=0}^{+\infty} \alpha_n^2 \lambda_n. \end{cases}$$

According to [1], the Laplacian eigenfunctions are an optimal basis for the representation of signals with bounded gradient magnitude. In fact, the spectral decomposition $f_n = \sum_{i=0}^n \langle f, \phi_i \rangle_2 \phi_i$ is optimal for the approximation of functions with \mathcal{L}^2 bounded gradient magnitude; i.e., the residual error $r_n := f - f_n$ is bounded as

$$\|r_n\|_2^2 \leq \frac{\|\nabla f\|_2^2}{\lambda_{n+1}}. \quad (1)$$

The spectral decomposition is also optimal for the approximation of functions with respect to the error estimation in (1). In fact, for any $0 \leq \alpha < 1$ there is no integer n and no sequence $(\psi_i)_{i=0}^n$ of linearly independent functions in \mathcal{L}^2 such that

$$\left\| f - \sum_{i=1}^n \langle f, \psi_i \rangle_2 \psi_i \right\|_2 \leq \alpha \frac{\|\nabla f\|_2}{\lambda_{n+1}}, \quad \forall f.$$

Indeed, the Laplacian eigenfunctions are an optimal basis for the definition of scalar functions on a given domain, such as the Laplacian spectral kernels and distances on 3D

shapes (Sect. 3). From the computational point of view, they have two main drawbacks: a high computational cost and storage overhead, which prevent the evaluation of a large number of Laplacian eigenpairs, and numerical instabilities with respect to the surface discretisation [38, 39].

2.2. Theoretical background: spectral kernels and distances

We briefly review the equivalent definitions of the spectral kernels (Sect. 2.2.1) and distances (Sect. 2.2.2), which will be useful to introduce their spectrum-free computation and main properties.

2.2.1. Spectral operator and kernel

According to [38, 39], let $\rho : \mathbb{R}^+ \rightarrow \mathbb{R}$ be a positive *filter map* that is square integrable and admits the power series $\rho(s) = \sum_{n=0}^{+\infty} \alpha_n s^n$. Considering the orthonormal Laplacian eigensystem $(\lambda_n, \phi_n)_{n=0}^{+\infty}$, the spectral representation of the functions $\Delta^i f$ and $(\Delta^\dagger)^i f$ is

$$\Delta^i f = \sum_{n=0}^{+\infty} \lambda_n^i \langle f, \phi_n \rangle_2 \phi_n, \quad (\Delta^\dagger)^i f = \sum_{n=1}^{+\infty} \frac{1}{\lambda_n^i} \langle f, \phi_n \rangle_2 \phi_n. \quad (2)$$

If $\lambda_k = 0$, then we neglect the corresponding entry in $(\Delta^\dagger)^i f$. In $\mathcal{L}^2(\mathcal{N})$, we define the *spectral operator*

$$\Phi_\rho f = \sum_{n=0}^{+\infty} \alpha_n \Delta^n f = \sum_{n=0}^{+\infty} \rho(\lambda_n) \langle f, \phi_n \rangle_2 \phi_n, \quad (3)$$

which is linear, continuous, self-adjoint, and $\Phi_\rho f = \langle K_\rho, f \rangle_2$, with

$$K_\rho(\mathbf{p}, \mathbf{q}) = \sum_{n=0}^{+\infty} \rho(\lambda_n) \phi_n(\mathbf{p}) \phi_n(\mathbf{q}) \quad (4)$$

spectral kernel. The spectral kernel operator and kernel are well-posed if the filter is bounded or square-integrable. The spectral kernel is also symmetric and self-adjoint, as a consequence of its definition. From these relations, it follows that

- the spectral kernel $K_\rho(\mathbf{p}, \cdot) = \Phi_\rho \delta_{\mathbf{p}}$ is achieved as the action of the spectral operator on the δ -function at \mathbf{p} ;
- the spectral kernel is the solution to the differential equation $\Phi_{1/\rho} K_\rho(\mathbf{p}, \cdot) = \delta_{\mathbf{p}}$. In Sect. 2.4, this property is further investigated in terms of the relation between the spectral and Green kernels.

To express the spectral distances in terms of the spectral operator (Sect. 2.5), we show that the pseudo-inverse of Φ_ρ is induced by the filter function $1/\rho$; i.e., $\Phi_\rho^\dagger = \Phi_{1/\rho}$. In fact, from the spectral representation (3) of $\Phi_{1/\rho}$ and Φ_ρ , we get that

$$\begin{aligned}\Phi_\rho \Phi_{1/\rho} \Phi_\rho &= \Phi_\rho, & \Phi_{1/\rho} \Phi_\rho \Phi_{1/\rho} &= \Phi_{1/\rho}, \\ \langle \Phi_{1/\rho} \Phi_\rho f, g \rangle_2 &= \langle f, \Phi_{1/\rho} \Phi_\rho g \rangle_2.\end{aligned}\tag{5}$$

In the following, we assume that both Φ_ρ and $\Phi_{1/\rho}$ are well-defined; for instance, this hypothesis is satisfied if ρ is not null only on a compact interval of \mathbb{R}^+ and it is valid in the discrete case (Sect. 3), where the Laplacian spectrum belongs to the interval $\mathcal{J} := [0, \lambda_{\max}(\tilde{\mathbf{L}})]$, with $\lambda_{\max}(\tilde{\mathbf{L}})$ maximum Laplacian eigenvalue.

2.2.2. Spectral distances

Through the spectral operator, we introduce the scalar product and the corresponding distance as [38, 39]

$$\begin{cases} \langle f, g \rangle := \langle \Phi_{1/\rho} f, \Phi_{1/\rho} g \rangle_2 = \sum_{n=0}^{+\infty} \frac{\langle f, \phi_n \rangle_2 \langle g, \phi_n \rangle_2}{\rho^2(\lambda_n)}, & (a) \\ d^2(f, g) := \|f - g\|^2 = \sum_{n=0}^{+\infty} \frac{|\langle f - g, \phi_n \rangle_2|^2}{\rho^2(\lambda_n)}. & (b) \end{cases}\tag{6}$$

Indicating with $\delta_{\mathbf{p}}$ the map that takes value 1 at \mathbf{p} and 0 otherwise, and selecting $f := \delta_{\mathbf{p}}$, $g := \delta_{\mathbf{q}}$ in Eq. (6b), the *spectral distance* (Fig. 1) on \mathcal{N} is defined as

$$\begin{aligned}d^2(\mathbf{p}, \mathbf{q}) &:= \|\delta_{\mathbf{p}} - \delta_{\mathbf{q}}\|^2 \\ &\stackrel{\text{Eq. (6b)}}{=} \sum_{n=0}^{+\infty} \frac{|\phi_n(\mathbf{p}) - \phi_n(\mathbf{q})|^2}{\rho^2(\lambda_n)} \\ &\stackrel{\text{Eq. (6a)}}{=} \|\Phi_{1/\rho}(\delta_{\mathbf{p}}) - \Phi_{1/\rho}(\delta_{\mathbf{q}})\|_2^2 \\ &= \|K_{1/\rho}(\mathbf{p}, \cdot) - K_{1/\rho}(\mathbf{q}, \cdot)\|_2^2 \\ &= K_{1/\rho}(\mathbf{p}, \mathbf{p}) - 2K_{1/\rho}(\mathbf{p}, \mathbf{q}) + K_{1/\rho}(\mathbf{q}, \mathbf{q});\end{aligned}$$

i.e., these equivalent formulations involve the Laplacian spectrum, the spectral operator and kernel. The third equality follows from the identity $\Phi_\rho(\delta_{\mathbf{p}}) = K_\rho(\mathbf{p}, \cdot)$ and will be applied to the computation of the spectral distances (Sect. 3.2); in fact, it is independent of the evaluation of the Laplacian spectrum. The last identity is achieved by applying

the relation

$$\begin{aligned}
d^2(\mathbf{p}, \mathbf{q}) &= \|K_{1/\rho}(\mathbf{p}, \cdot) - K_{1/\rho}(\mathbf{q}, \cdot)\|_2^2 \\
&= \sum_{n=0}^{+\infty} \rho^{-2}(\lambda_n) \phi_n(\mathbf{p}) \phi_n(\mathbf{p}) - 2 \sum_{n=0}^{+\infty} \rho^{-2}(\lambda_n) \phi_n(\mathbf{p}) \phi_n(\mathbf{q}) \\
&\quad + \sum_{n=0}^{+\infty} \rho^{-2}(\lambda_n) \phi_n(\mathbf{q}) \phi_n(\mathbf{q}).
\end{aligned}$$

2.3. Stability of the eigenpairs of the spectral operator

Generalising the results in [38, 39], we show that the computation of a single eigenvalue of the spectral operator is numerically stable and instabilities are generally due to repeated or close eigenvalues. Firstly, we notice that if λ is a Laplacian eigenvalue of multiplicity m then $\mu := \rho(\lambda)$ is an eigenvalue of Φ_ρ and its multiplicity is equal to or greater than m . The corresponding eigenfunction is the Laplacian eigenfunction associated with the eigenvalue λ ; i.e., $\Delta\phi = \lambda\phi$ and $\Phi_\rho\phi = \rho(\lambda)\phi$.

We perturb the spectral operator by $\delta\mathcal{E}$, $\delta \rightarrow 0$, and compute the eigenpair $(\mu(\delta), \phi(\delta))$ of the corresponding operator $\Phi_\rho + \delta\mathcal{E}$; i.e.,

$$(\Phi_\rho + \delta\mathcal{E})\phi(\delta) = \mu(\delta)\phi(\delta), \quad \phi(0) = \phi, \quad \mu(0) = \mu. \quad (7)$$

Recalling that the derivative of $\mu(\delta)$ measures the variation of the eigenvalue, deriving (7) with respect to δ , and evaluating the resulting relation at 0, we have that

$$\mathcal{E}\phi + \Phi_\rho\phi'(0) = \mu'(0)\phi + \mu\phi'(0). \quad (8)$$

From (8) and the self-adjointness of Φ_ρ , it follows that $\langle\phi, \Phi_\rho\phi'(0)\rangle_2 = \mu\langle\phi, \phi'(0)\rangle_2$.

Multiplying both sides of (8) by ϕ and applying the previous identity we get that

$$|\mu'(0)| = |\langle\phi, \mathcal{E}\phi\rangle_2| \leq \|\mathcal{E}\|_2 \|\phi\|_2^2 = \|\mathcal{E}\|_2;$$

i.e., the computation of the eigenvalue of Φ_ρ with multiplicity one is stable.

Assuming that $\mu := \rho(\lambda)$ is an eigenvalue of Φ_ρ with multiplicity m and rewriting the characteristic polynomial as $p(s) = (s - \mu)^m q(s)$, where $q(\cdot)$ is a polynomial of degree $n - m$ and $q(\mu) \neq 0$, we get that

$$(s - \mu)^m = \frac{p(s)}{q(s)} \approx \frac{O(\delta)}{q(s)}, \quad \delta \rightarrow 0, \quad \text{as } p(s) \rightarrow 0, \text{ when } s \rightarrow \mu;$$




Figure 1: Different filters ρ induce spectral distances from a seed point with a different behaviour, in terms of the locality, shape, and distribution of the level-sets and according to the decay of $1/\rho$ to zero. These distances have been computed with the Padé-Chebyshev approximation ($r = 5$).

i.e., $s \approx \mu + O(\delta^{\frac{1}{m}})$. Indeed, modifying the Laplacian matrix in such a way that the filtered eigenvalues are perturbed by $\delta := 10^{-m}$ corresponds to a change of order 0.1 in μ (i.e., $s \approx \mu + 0.1$) and this amplification becomes larger as the multiplicity of the eigenvalue increases.

While multiple eigenvalues are typically associated with symmetric shapes, numerically close or switched eigenvalues are present regardless of the surface regularity. This unstable computation of multiple eigenpairs of the spectral operator generally affects the accuracy of the *truncated spectral approximation*

$$d_k(\mathbf{p}, \mathbf{q}) := \sum_{n=0}^k \frac{|\phi_n(\mathbf{p}) - \phi_n(\mathbf{q})|^2}{\rho^2(\lambda_n)} \quad (9)$$

of the corresponding distances. In fact, each filtered eigenvalue $\rho(\lambda_n)$, $n = 1, \dots, k$, which appears at the denominator of d_k , can further accentuate the numerical error of λ_n in the distance computation. Furthermore, the computation of the Laplacian spectrum is time-consuming and it is difficult to properly select the number of eigenpairs that is necessary to accurately approximate the spectral distance. Indeed, we propose an evaluation of the spectral distance that is independent of the computation of the Laplacian spectrum and is equivalent to a set of differential equations involving only the Laplace-Beltrami operator and its pseudo-inverse. This novel spectrum-free computation is a generalisation of the discrete approach recently presented in [38, 39].

2.4. Relation between spectral and Green kernels

To study the relation between the Green and the spectral kernels, let us consider a linear differential operator \mathcal{L} and the corresponding *Green kernel* $K : \mathcal{N} \times \mathcal{N} \rightarrow \mathbb{R}$ such that $\mathcal{L}K(\mathbf{p}, \mathbf{q}) = \delta(\mathbf{p} - \mathbf{q})$. Then, the solution to the differential equation $\mathcal{L}u = f$ is expressed in terms of the Green kernel as $u(\mathbf{p}) = \langle K(\mathbf{p}, \cdot), f \rangle_2$.

Noting that the eigensystem $(\mu_n, \psi_n)_{n=0}^{+\infty}$ of the integral operator $\mathcal{A}_K f := \langle K(\cdot, \cdot), f \rangle_2$ induced by the Green kernel satisfies the relation $\mathcal{L}\psi_n = \mu_n^{-1}\psi_n$, $\mu_n \neq 0$, we have that the eigensystem of \mathcal{L} is $(\mu_n^{-1}, \psi_n)_{n=1}^{+\infty}$. Indeed, \mathcal{L} and \mathcal{A}_K have the same eigenfunctions and reciprocal eigenvalues. In particular, the spectral representation of the Green kernel $K(\mathbf{p}, \mathbf{q}) = \sum_{n=0}^{+\infty} \mu_n \psi_n(\mathbf{p}) \psi_n(\mathbf{q})$ is uniquely defined by the spectrum of \mathcal{L} . Combining the previous results with the properties of the spectral operator and kernel (Sect. 2.2.2), we get that

- for the *harmonic operator* $\mathcal{L} = \Delta$, the Green kernel is the *commute-time kernel*

$$K_\Delta(\mathbf{p}, \mathbf{q}) = \sum_{n=1}^{+\infty} \frac{1}{\lambda_n} \phi_n(\mathbf{p}) \phi_n(\mathbf{q});$$

- for the *bi-Laplacian operator* $\mathcal{L} := \Delta^2$, the Green kernel is the *bi-harmonic kernel*

$$K_{\Delta^2}(\mathbf{p}, \mathbf{q}) := \sum_{n=1}^{+\infty} \frac{1}{\lambda_n^2} \phi_n(\mathbf{p}) \phi_n(\mathbf{q});$$

- for the *diffusion operator* $\mathcal{L} = \exp(t\Delta)$, the Green kernel is the *diffusion kernel* (Fig. 2)

$$K_t(\mathbf{p}, \mathbf{q}) = \sum_{n=0}^{+\infty} \exp(-t\lambda_n) \phi_n(\mathbf{p}) \phi_n(\mathbf{q});$$

- for the *spectral operator* $\mathcal{L} = \Phi_{1/\rho}$, the Green kernel is the *spectral kernel* K_ρ defined in Eq. (4).

2.5. Spectrum-free approximation of kernels and distances

Recalling the relation $\Phi_\rho^\dagger = \Phi_{1/\rho}$ in Eq. (5) and noting that

$$d(f, g) = \|u\|_2, \quad u = \Phi_{1/\rho}(f - g) \iff \Phi_\rho u = f - g,$$




Figure 2: (a,b) Level-sets of the diffusion kernel at two different scales and at different seed points, which confirm the locality, smoothness, and shape-awareness of the heat kernel.

the spectral distance is equivalent to the norm of (i) the function $u = \Phi_{1/\rho}(f - g)$ and (ii) the solution of the differential equation $\Phi_\rho u = f - g$. Through these relations, the spectral kernels will be computed by approximating the filter with a rational function (Sect. 2.5.1) and by converting the evaluation of the corresponding distances (Sect. 2.5.2) to the solution of a set of differential equations that involve the Laplace-Beltrami operator. The choice of one of these two equivalent representations depends on the selected filter and the simplicity of evaluating either Φ_ρ or $\Phi_{1/\rho}$.

2.5.1. Kernel approximation, convergence, and accuracy

Let us introduce the space of the Laplacian spectral kernels as

$$\mathcal{K}(\mathcal{N}) := \{K_\rho : \mathcal{N} \times \mathcal{N} \rightarrow \mathbb{R}, K_\rho \text{ spectral kernel in Eq. (4)}\},$$

and equipped with the $\mathcal{L}^2(\mathcal{N} \times \mathcal{N})$ scalar product. Since

$$\|K_\rho - K_\varphi\|_2^2 = \sum_{n=0}^{+\infty} |\rho(\lambda_n) - \varphi(\lambda_n)|^2 \leq \|\rho - \varphi\|_2^2,$$

the approximation of a given spectral kernel K_ρ with a new kernel K_φ in $\mathcal{K}(\mathcal{N})$ is reduced to the approximation of ρ by φ on a proper subspace of functions (e.g., the




Figure 3: Level-sets of the basis functions at \mathbf{p}_i associated with the rational polynomial approximation of the spectral distances.

space of polynomials or rational polynomials). The class of functions used for the approximation of the input filter ρ is selected in such a way that K_ϕ provides a good approximation of the input filter K_ρ and is easily computable.

2.5.2. Rational approximation based on the canonical basis

We apply the *Padé-Chebyshev rational approximation* to the filter map $1/\rho$. Let \mathcal{R}_r^l be the space of all rational functions

$$c_{rl}(s) := \frac{p_l(s)}{q_r(s)} = \frac{\beta_0 + \beta_1 s + \dots + \beta_l s^l}{\alpha_0 + \alpha_1 s + \dots + \alpha_r s^r}, \quad s \in [a, b].$$

We briefly recall that the representation of the rational approximation is not unique, unless we impose that it is irreducible; for example, by choosing $q_r(a) = 1$. Given a filter $\rho : [a, b] \rightarrow \mathbb{R}$, there exists a unique best approximation of $1/\rho$ in \mathcal{R}_r^l with respect to the ℓ_∞ norm [20] (Ch. 9), which is represented in terms of the *canonic basis*




Figure 4: Smoothness and locality of the diffusion kernel at a seed point (red dot) and at different scales, which have been computed with (a-c) the truncated spectral approximation (i.e., k Laplacian eigenpairs) and (d) the Padè-Chebyshev approximation. Since the input shape has $2K$ vertices, the spectral approximation (c) provides the ground-truth.

$$\mathcal{B} := \{s^i\}_{i=0}^l \cup \{s^{-i}\}_{i=1}^r.$$

Recalling that the dimension of the space \mathcal{R}_r^l of rational polynomial of degree (r, l) is $l + r + 1$, let us express the best rational approximation c_{rl} of ρ in terms of these basis functions. To this end, we can either apply algebraic rules or impose interpolating constraints at $(l + r + 1)$ points in order to compute the new coefficients through the identity

$$\sum_{i=0}^l a_i s^i + \sum_{i=1}^r b_i s^{-i} = c_{rl}(s).$$

The canonical basis used to represent the rational approximation of the input filter simplifies the computation of the spectral kernel; in fact, we need to evaluate only the functions $\Delta^i f$ and $(\Delta^\dagger)^i f$. Instead of applying the spectral representation (2), we define

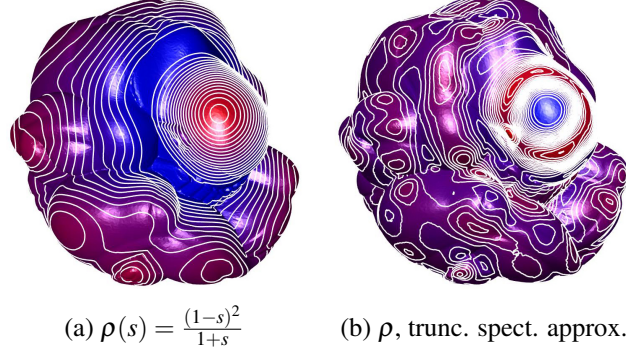


Figure 5: Comparison of the spectral distance induced by the same filter and approximated with the (a) spectrum-free and (b) truncated spectral approximation ($k = 500$). This last method produces some artefacts that become evident as we move far from the seed point.

a recursive procedure as follows. To compute $g_1 = \Delta^\dagger f$, let us multiply both sides by Δ and notice that

$$\Delta g_1 = \Delta \Delta^\dagger f = f - \langle f, \phi_0 \rangle_2 \phi_0, \quad (10)$$

where $\phi_0 = \mathbb{1}$ is the constant eigenfunction equal to 1. By definition of pseudo-inverse, g_1 is the least-squares solution to the equation $\Delta g_1 = \tilde{f}$, where \tilde{f} is equal to f minus its mean $\langle f, \phi_0 \rangle_2 \phi_0$. For the general case, we apply the recursive relation

$$g_i := (\Delta^\dagger)^i f = \Delta^\dagger (\Delta^\dagger)^{i-1} f = \Delta^\dagger g_{i-1}, \quad i \geq 2,$$

which reduces to the previous case. In a similar way, $g_i = \Delta^i f$ is calculated as $h_i = \Delta h_{i-1}$, $i = 1, \dots, r$, with $h_0 := f$. Indeed, the proposed approach requires the solution of $(l + r)$ Laplace equation with a different right-hand side.

Convergence and accuracy. To verify that the sequence

$$(\Phi_{1/\rho}^{(r)} f)_{r=0}^{+\infty}, \quad \Phi_{1/\rho}^{(r)} f := \sum_{n=0}^{+\infty} c_{rl}(\lambda_n) \langle f, \phi_n \rangle_2 \phi_n,$$

induced by the rational polynomial approximation c_{rl} of $1/\rho$, converges to $\Phi_{1/\rho} f$, we apply the upper bound

$$\begin{aligned} \left\| \Phi_{1/\rho}^{(r)} f - \Phi_{1/\rho} f \right\|_2^2 &\leq \|c_{rl} - 1/\rho\|_\infty^2 \sum_{n=0}^{+\infty} |\langle f, \phi_n \rangle_2|^2 \\ &= \sigma_{rl}^2 \|f\|_2^2, \quad \sigma_{rl} \approx \mathcal{O}(s^{r+l+1}), \quad s \rightarrow 0; \end{aligned}$$




Figure 6: Reducing the scale t , the corresponding diffusion distances become locally unstable, as we move far from the seed point (on the feet). See also Fig. 7.

where σ_{rl} is the approximation error between $1/\rho$ and c_{rl} . For the diffusion operator [52], $\sigma_{rr} = 10^{-5}$.

3. Spectrum-free approximation of kernels and distances

We briefly summarise the discretisation of the spectral kernel and distances [38, 39], which is then used to establish a simple relation between the spectral kernel and its pseudo-inverse (Sect. 3.1). Then, we introduce the spectrum-free approximation of the spectral kernels and distances (Sect. 3.2) and discuss the computational cost of the main steps of the proposed approach (Sect. 3.3).

3.1. Discretisation of spectral kernels and distances

Let us consider a (triangular, polygonal, volumetric) mesh $\mathcal{M} := (\mathcal{P}, \mathcal{T})$, which discretises a domain \mathcal{N} , where $\mathcal{P} := \{\mathbf{p}_i\}_{i=1}^n$ is the set of n vertices and \mathcal{T} is the connectivity graph. On \mathcal{M} , a scalar function $f: \mathcal{M} \rightarrow \mathbb{R}$ is identified with the vector $\mathbf{f} := (f(\mathbf{p}_i))_{i=1}^n$ of f -values at \mathcal{P} .

Let us introduce the *Laplacian matrix* $\tilde{\mathbf{L}} := \mathbf{B}^{-1}\mathbf{L}$, where \mathbf{L} and \mathbf{B} are the stiffness and mass matrix respectively (e.g., cotangent [40], Voronoi-cotg [15], linear FEM [42] weights), and the corresponding spectral decomposition is $\mathbf{L}\mathbf{X} = \mathbf{B}\mathbf{X}\mathbf{\Lambda}$, $\mathbf{X}^\top \mathbf{B}\mathbf{X} = \mathbf{I}$, where $\mathbf{X} := [\mathbf{x}_1, \dots, \mathbf{x}_n]$ is the eigenvectors' matrix and $\mathbf{\Lambda}$ is the diagonal matrix of the

eigenvalues $(\lambda_i)_{i=1}^n$. Analogous discretisations apply to polygonal [3, 22] and tetrahedral [27, 51] meshes, or point sets [29].

Under these assumptions, we introduce the spectral kernels and distances (Sect. 3.1.1), together with their spectrum-free approximation (Sect. 3.1.2), which is based on a rational approximation of the input filter.

3.1.1. Discrete spectral kernels and distances

The spectral operator $\Phi_{1/\rho}$ is discretised by the *spectral kernel matrix* $\mathbf{K}_{1/\rho}$ such that

$$\rho^{-1}(\lambda_i) \stackrel{\text{Eq. (3)}}{=} \langle \Phi_{1/\rho} \phi_i, \phi_j \rangle_2 = \mathbf{x}_i^\top \mathbf{K}_{1/\rho}^\top \mathbf{B} \mathbf{x}_j, \quad \forall i = 1, \dots, n.$$

Indeed, $\mathbf{K}_{1/\rho} = \mathbf{X} \rho^\dagger(\Lambda) \mathbf{X}^\top \mathbf{B}$ is the pseudo-inverse of the *spectral kernel* $\mathbf{K}_\rho = \rho(\tilde{\mathbf{L}})$, which is a filtered version of the Laplacian matrix. Here, $\rho^\dagger(\Lambda) = \text{diag}(1/\rho(\lambda_i))_{i=1}^n$ and its entry is null if $\rho(\lambda_i) = 0$. Then, the *discrete spectral distances* [38, 39] are

$$d^2(\mathbf{p}_i, \mathbf{p}_j) = \|\mathbf{K}_{1/\rho}(\mathbf{e}_i - \mathbf{e}_j)\|_{\mathbf{B}}^2 = \sum_{l=1}^n \frac{|\langle \mathbf{x}_l, \mathbf{e}_i - \mathbf{e}_j \rangle_{\mathbf{B}}|^2}{\rho^2(\lambda_l)},$$

where \mathbf{e}_i is the vector of the canonical basis of \mathbb{R}^n , $\langle \mathbf{f}, \mathbf{g} \rangle_{\mathbf{B}} := \mathbf{f}^\top \mathbf{B} \mathbf{g}$ and $\|\mathbf{f}\|_{\mathbf{B}}^2 := \mathbf{f}^\top \mathbf{B} \mathbf{f}$ are the scalar product and the norm induced by the mass matrix, respectively. In this case, we have derived the spectral distances by applying the continuous expression of the spectral operator instead of the spectral kernel, as done in [38, 39].

Analogously to the definitions in Sect. 2.2, the spectral distance is equal to the norm $d(\mathbf{f}, \mathbf{g}) := \|\mathbf{u}\|_{\mathbf{B}}$ of the solution to the linear system $\mathbf{K}_\rho \mathbf{u} = \mathbf{f} - \mathbf{g}$ or to the norm of the vector $\mathbf{u} = \mathbf{K}_{1/\rho}(\mathbf{f} - \mathbf{g})$. In a similar way, the spectral distance between two points reduces to

$$d(\mathbf{p}_i, \mathbf{p}_j) = \|\mathbf{K}_{1/\rho}(\mathbf{e}_i - \mathbf{e}_j)\|_{\mathbf{B}}.$$

Indeed, this approximation of the spectral distances involves the spectral kernel only and allows us to bypass numerical inaccuracies due to repeated or close Laplacian eigenvalues [20] (§ 7).

3.1.2. Spectral kernel approximation and conditioning

Since the eigenvalues of the spectral kernel matrix \mathbf{K}_ρ are $(\rho(\lambda_i))_{i=1}^n$, the approximation of \mathbf{K}_ρ with a new kernel \mathbf{K}_φ reduces to the approximation of the corresponding




Figure 7: Analogous behaviour of the heat kernel with respect to a different resolution (n vertices) of the input shape, in terms of shape and distribution of the level-sets.

filters with respect to the ℓ_∞ norm; in fact,

$$\|\mathbf{K}_\rho - \mathbf{K}_\varphi\|_2 = \|\mathbf{K}_{\rho-\varphi}\|_2 = \max_{i=1,\dots,n} \{|\rho(\lambda_i) - \varphi(\lambda_i)|\} \leq \|\rho - \varphi\|_\infty.$$

The approximation φ of ρ is computed on the interval $[0, \lambda_{\max}(\tilde{\mathbf{L}})]$, where the maximum Laplacian eigenvalue is evaluated by the Arnoldi method [20], or is set equal to the upper bound $\lambda_{\max}(\tilde{\mathbf{L}}) \leq \min\{\max_i\{\sum_j \tilde{L}(i, j)\}, \max_j\{\sum_i \tilde{L}(i, j)\}\}$ [26, 48].

Conditioning of the spectral kernel. We now analyse the conditioning of the spectral kernel. Assuming that ρ is an increasing function (i.e., $1/\rho$ is a low pass filter), the conditioning number of the filtered Laplacian matrix is bounded as

$$\kappa_2(\mathbf{K}_\rho) = \kappa_2(\rho(\tilde{\mathbf{L}})) = \frac{\max_{i=1,\dots,n} \{\rho(\lambda_i)\}}{\min_{i=1,\dots,n} \{\rho(\lambda_i)\}} = \frac{\|\rho\|_\infty}{\rho(0)},$$

and it is ill-conditioned when $\rho(0)$ is close to zero or ρ is unbounded. If ρ is bounded and $\rho(0)$ is not too close to 0, then the filtered Laplacian matrix is well-conditioned. If $\rho(0)$ is null, then we consider the smallest and not null filtered Laplacian eigenvalue at the denominator of the previous relation.

3.2. Spectrum-free approximation of distances: canonical basis

Analogously to the discussion in Sect. 2.5.2, the basis \mathcal{B} used to represent the rational approximation of the input filter simplifies the computation of the spectral kernel; in fact, we need to compute only the following vectors $\tilde{\mathbf{L}}^i \mathbf{f}$ and $(\tilde{\mathbf{L}}^i)^\dagger \mathbf{f}$. Applying the spectral representation of the Laplacian matrix

$$\tilde{\mathbf{L}} = \mathbf{X} \mathbf{\Lambda} \mathbf{X}^\top \mathbf{B}, \quad \mathbf{\Lambda} = \text{diag}(\lambda_i)_{i=1}^n, \quad \lambda_1 = 0, \quad \lambda_i \leq \lambda_{i+1}, \quad i = 2, \dots, n-1,$$

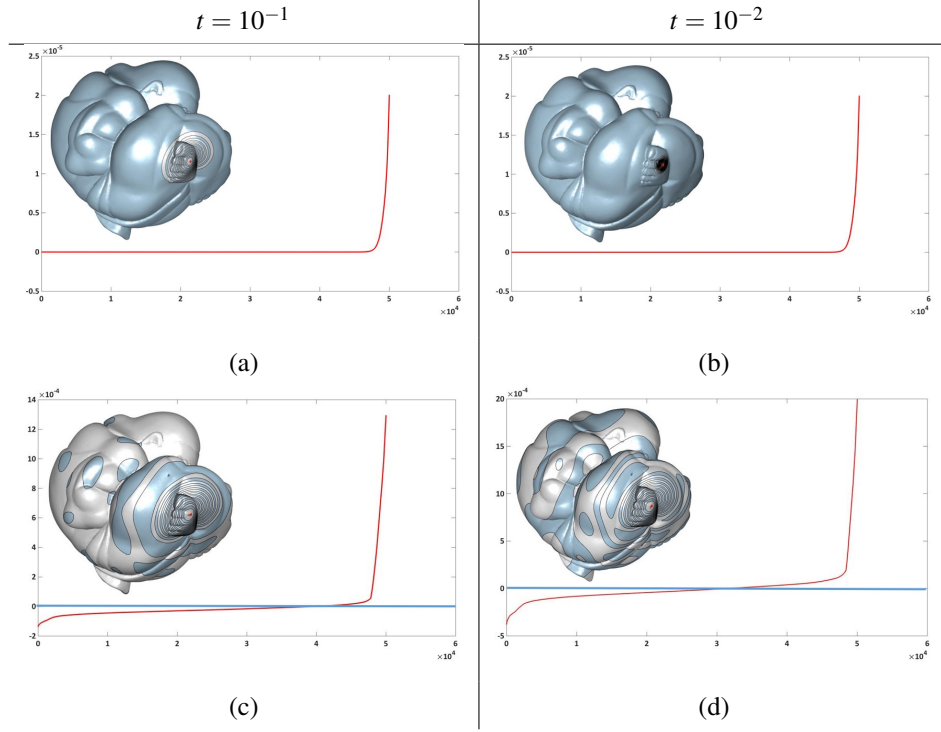


Figure 8: (c,d) Small undulations and negative values of the truncated spectral approximation at small scales. These local undulations are not present in the Padé-Chebyshev approximation (a,b), whose values are always positive.

we can represent its powers and the corresponding pseudo-inverse matrices as

$$\tilde{\mathbf{L}}^i = \mathbf{X} \Lambda^i \mathbf{X}^\top \mathbf{B}, \quad (\tilde{\mathbf{L}}^i)^\dagger = \mathbf{X} (\Lambda^i)^\dagger \mathbf{X}^\top \mathbf{B}, \quad \Lambda^\dagger = \text{diag}(0, \lambda_i^{-1})_{i=2}^n.$$

To compute $\mathbf{g}_1 = \tilde{\mathbf{L}}^\dagger \mathbf{f}$, let us multiply both sides by $\tilde{\mathbf{L}}$ and notice that

$$\begin{aligned} \tilde{\mathbf{L}} \mathbf{g}_1 &= \tilde{\mathbf{L}} \tilde{\mathbf{L}}^\dagger \mathbf{f} \\ &= \mathbf{X} \Lambda \Lambda^\dagger \mathbf{X}^\top \mathbf{B} \mathbf{f}, \quad \mathbf{X}^\top \mathbf{B} \mathbf{X} = \mathbf{I}, \\ &= \mathbf{X} (\mathbf{I} - \mathbf{e}_1 \mathbf{e}_1^\top) \mathbf{X}^\top \mathbf{B} \mathbf{f}, \quad \mathbf{e}_1 := [1, 0, \dots, 0]^\top, \\ &= \mathbf{f} - (\mathbf{1}^\top \mathbf{B} \mathbf{f}) \mathbf{1}, \quad \mathbf{1} := [1, 1, \dots, 1]^\top; \end{aligned}$$

i.e., \mathbf{g}_1 is the least-squares solution to the sparse and symmetric linear system

$$\mathbf{L} \mathbf{g}_1 = \mathbf{B} \mathbf{g}_0, \quad \mathbf{g}_0 := \mathbf{f} - (\mathbf{1}^\top \mathbf{B} \mathbf{f}) \mathbf{1},$$




Figure 9: *Stability of the spectrum-free computation of the diffusion kernel with respect to partially-sampled surfaces.*

where \mathbf{g}_0 is achieved by subtracting to \mathbf{f} its mean value $\langle \mathbf{f}, \mathbf{1} \rangle_{\mathbf{B}}$ (c.f., Eq. (10)). For the general case, we apply the recursive relation

$$\mathbf{g}_i := (\tilde{\mathbf{L}}^\dagger)^i \mathbf{f} = \tilde{\mathbf{L}}^\dagger (\tilde{\mathbf{L}}^\dagger)^{i-1} \mathbf{f} = \tilde{\mathbf{L}}^\dagger \mathbf{g}_{i-1}, \quad i \geq 2,$$

which reduces to the previous case (Fig. 3). In a similar way, $\mathbf{g}_i = \tilde{\mathbf{L}}^i \mathbf{f}$ is calculated by recursively solving the sparse, symmetric, and positive-definite linear systems $\mathbf{B}\mathbf{g}_{i+1} = \mathbf{L}\mathbf{g}_i$, $i = 1, \dots, r-1$, with $\mathbf{g}_0 := \mathbf{f}$.

3.3. Computational cost

The truncated spectral approximation, whose computational cost depends on the sparsity degree of the Laplacian matrix, takes from $\mathcal{O}(kn \log n)$ to $\mathcal{O}(kn^2)$ time, where k is the number of selected eigenpairs. Selecting a rational approximation of the input filter of degree (r, l) , the evaluation of the corresponding spectral kernel is reduced to solve l linear systems whose coefficient matrix is \mathbf{B} and r linear systems whose coefficient matrix is \mathbf{L} . Through iterative solvers, the computational cost is $\mathcal{O}((r+l)\tau(n))$, where $\tau(n)$ is the cost for the solution of a sparse linear system, which varies from $\mathcal{O}(n)$ to $\mathcal{O}(n^2)$, according to the sparsity of the coefficient matrix, and it is $\mathcal{O}(n \log n)$ in the average case.

In a similar way, the computation of the spectral distance between two points has the same order of complexity of the evaluation of the corresponding spectral kernel at the same input points. For the computation of the one-to-all spectral distance, we lump or pre-factorise \mathbf{B} (if not already diagonal) and pre-factorise \mathbf{L} . Then, the overall

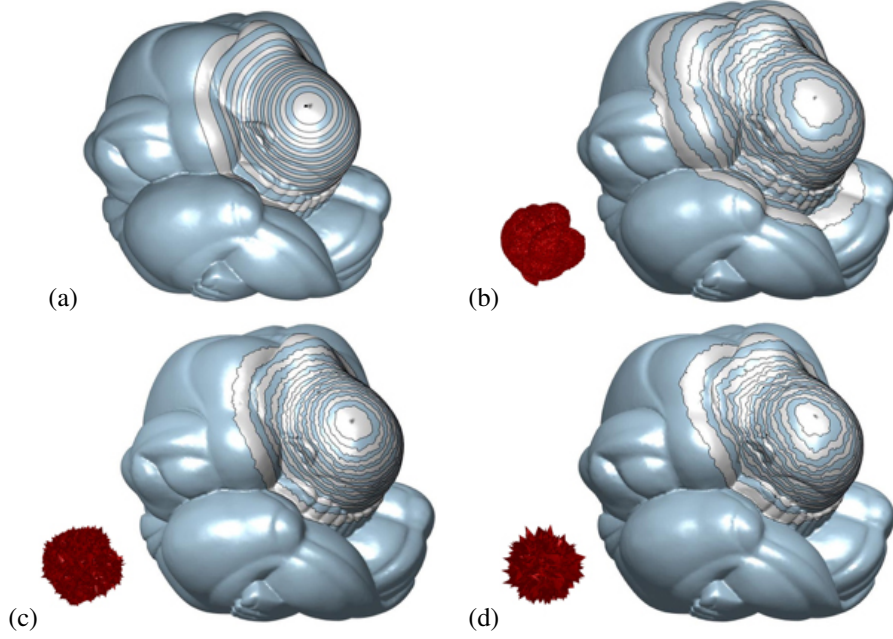


Figure 10: Behaviour and robustness of the spectrum-free approximation of the diffusion kernel on (b-d) the noisy shapes (bottom part, red surfaces), which have been plotted on (a) the initial shape.

computational cost varies from $\mathcal{O}(rn)$ (\mathbf{B} diagonal) to $\mathcal{O}(n \log n + rn)$ (\mathbf{B} not diagonal), where $\mathcal{O}(n \log n)$ is the time for the factorisation of \mathbf{L} .

In spite of this different computational cost, the truncated spectral approximation (Fig. 4) is affected by small geometric undulations (especially at small scales), the use of heuristics for the selection of the number of Laplacian eigenpairs with respect to the target approximation accuracy, and the scale of features of the input shape. The spectrum-free computation (Fig. 5) generally provides better results in terms of smoothness, regularity, and accuracy of the computed spectral basis. For further comparison examples, we refer the reader to Sect. 4.

4. Discussion

We now discuss the selection of the filter map and main examples on the computation of the Laplacian spectral distances.




Figure 11: Robustness of the diffusion kernel at different scales and on self-intersecting surfaces.

The spectral kernels and distances depend only on the behaviour of the filter in the spectral domain and on the Laplacian spectrum: increasing or decreasing the decay of the filter to zero encodes global or local shape details, respectively. Recalling that an arbitrary filter $1/\rho$ can be approximated in \mathcal{R}_r^l with an accuracy of order $\mathcal{O}(s^{l+r+1})$ with respect to the ℓ_∞ -norm (Sect. 3.2), we can use the space \mathcal{R}_r^l to define any filter and the corresponding spectral distances. In this way, we reduce the degree of freedom in the definition of the filter to the selection of $(l+r+1)$ coefficients and without losing the richness of the resulting spectral distances. For the selection of the coefficients of the rational filter and the filter frequencies, we can apply the rules proposed in [24] for




Figure 12: *Behaviour of the diffusion kernel on (almost) isometric 3D shapes.*

Laplacian spectral smoothing. In particular, the filter frequencies are derived from the dimension of a bounding box placed around the chosen feature \mathcal{F} and whose axis are aligned with the eigenvectors of the covariance matrix of \mathcal{F} .

The truncated spectral approximation (c.f., Eq. (9)) of the diffusion, and more generally of the spectral distances, is affected by (i) small undulations (especially at small scales, Fig. 6), (ii) heuristics for the selection of the number of Laplacian eigenpairs with respect to the target approximation accuracy and the scale of the shape features, and (iii) the overall computational cost and storage overhead for the computation of a subpart of the Laplacian spectrum. The comparison of these results with the ones induced by the spectrum-free computation (Fig. 7) shows the improvement of the proposed approach in terms of smoothness, regularity, and accuracy of the computed distances.

At small scales (Fig. 8(e,f)), the truncated spectral approximation of the diffusion kernel generally have small negative values as a matter of the slow decay of the exponential filter to zero for small eigenvalues. For the Padé-Chebyshev approximation (Fig. 8(a-d)), the kernel values are positive at all the scales, as we approximate the filter with a higher accuracy with respect to the truncated spectral approximation, which applies a low-pass filter.

Finally, the analogous distribution and shape of the level-sets of the heat kernel and diffusion distance confirm the robustness of the Padé-Chebyshev of the approximation with respect to surface resolution (Fig. 7), partial sampling (Fig. 9), geometric (Fig. 10) and topological (Fig. 11) noise, almost isometric deformations (Fig. 12).

5. Conclusions and future work

We have presented a unified approach to the definition, discretisation, and computation of the spectral kernels and distances, which are defined by filtering the Laplacian spectrum and generalise the commute-time, bi-harmonic, diffusion, and wave kernel and distances. Even though analytic filters are easily defined and encode global/local shape properties, the corresponding spectral kernels and distances are not capable of characterising a specific shape class (e.g., humans with respect to four-legs animals, chairs with respect to tables). Indeed, the main topic for future research is a deeper analysis of the constraints on the filter in order to define “optimal” spectral kernels and distances for shape comparison. To achieve this goal, we plan to apply and specialise learning methods, as initially investigated in [2, 8]. Another interesting aspect is the study of spectral kernels for the adaptive hash retrieval [5] and the identification of those PDEs, whose Green kernels are useful for shape analysis, as already demonstrated by the bi-harmonic and diffusion kernels.

Acknowledgements. We thank the Reviewers for their thorough review and constructive comments, which helped us to improve the technical part and presentation of the revised paper. This work is supported by the H2020 ERC Advanced Grant CHANGE, grant agreement No. 694515. Shapes are courtesy of the AIM@SHAPE Repository, the SHREC 2010 and 2016 data sets; tet-meshes have been generated by the TETGEN software <http://wias-berlin.de/software/tetgen/>.

References

- [1] Aflalo, Y., Brezis, H., R.Kimmel, 2015. On the optimality of shape and data representation in the spectral domain. *SIAM Journal Imaging Sciences* 8, 1141–1160.

- [2] Aflalo, Y., Bronstein, A.M., Bronstein, M.M., Kimmel, R., 2011. Deformable shape retrieval by learning diffusion kernels, in: *Scale space and variational methods in Computer Vision*, pp. 689–700.
- [3] Alexa, M., Wardetzky, M., 2011. Discrete Laplacians on general polygonal meshes. *ACM Trans. on Graphics* 30.
- [4] Bahonar, H., Mirzaei, A., Wilson, R.C., 2018. Diffusion wavelet embedding: a multi-resolution approach for graph embedding in vector space. *Pattern Recognition* 74, 518 – 530.
- [5] Bai, X., Yan, C., Yang, H., Bai, L., Zhou, J., Hancock, E.R., 2018. Adaptive hash retrieval with kernel-based similarity. *Pattern Recognition* 75, 136 – 148. *Distance Metric Learning for Pattern Recognition*.
- [6] Belkin, M., Niyogi, P., 2003. Laplacian eigenmaps for dimensionality reduction and data representation. *Neural Computations* 15, 1373–1396.
- [7] Berard, P., Besson, G., Gallot, S., 1994. Embedding Riemannian manifolds by their heat kernel. *Geometric and Functional Analysis* 4, 373–398.
- [8] Boscaini, D., Masci, J., Melzi, S., Bronstein, M.M., Castellani, U., Vandergheynst, P., 2015. Learning class-specific descriptors for deformable shapes using localized spectral convolutional networks. *Computer Graphics Forum* 34, 13–23.
- [9] Boscaini, D., Masci, J., Rodolà, E., Bronstein, M.M., Cremers, D., 2016. Anisotropic diffusion descriptors. *Computer Graphics Forum* 35, 431–441.
- [10] Bronstein, A., Bronstein, M., Kimmel, R., Mahmoudi, M., Sapiro, G., 2010. A Gromov-Hausdorff framework with diffusion geometry for topologically-robust non-rigid shape matching. *Intern. Journal of Computer Vision* 2-3, 266–286.
- [11] Bronstein, A.M., Bronstein, M.M., Ovsjanikov, M., Guibas, L.J., 2011. Shape Google: geometric words and expressions for invariant shape retrieval. *ACM Trans. on Graphics* 30.

- [12] Bronstein, M., Bronstein, A., 2011. Shape recognition with spectral distances. *IEEE Trans. on Pattern Analysis and Machine Intelligence* 33, 1065 –1071.
- [13] Cosmo, L., Rodolá, E., Albarelli, A., Mémoli, F., Cremers, D., 2017. Consistent partial matching of shape collections via sparse modeling. *Computer Graphics Forum* 36, 209–221.
- [14] Crane, K., Weischedel, C., Wardetzky, M., 2013. Geodesics in heat: A new approach to computing distance based on heat flow. *ACM Trans. on Graphics* 32, 152:1–152:11.
- [15] Desbrun, M., Meyer, M., Schröder, P., Barr, A.H., 1999. Implicit fairing of irregular meshes using diffusion and curvature flow, in: *ACM Siggraph*, pp. 317–324.
- [16] ElGhawalby, H., Hancock, E.R., 2015. Heat kernel embeddings, differential geometry and graph structure. *Axioms* 4, 275–293.
- [17] Gebal, K., Bærentzen, J.A., Aanaes, H., Larsen, R., 2009. Shape analysis using the auto diffusion function. *Computer Graphics Forum* 28, 1405–1413.
- [18] Ghosh, S., Das, N., Gonzalves, T., Quaresma, P., Kundu, M., 2018. The journey of graph kernels through two decades. *Computer Science Review* 27, 88 – 111.
- [19] de Goes, F., Goldenstein, S., Velho, L., 2008. A hierarchical segmentation of articulated bodies. *Computer Graphics Forum* 27, 1349–1356.
- [20] Golub, G., VanLoan, G., 1989. *Matrix Computations*. John Hopkins University Press, 2nd Edition.
- [21] Hammond, D.K., Vandergheynst, P., Gribonval, R., 2011. Wavelets on graphs via spectral graph theory. *Applied and Computational Harmonic Analysis* 30, 129 – 150.
- [22] Herholz, P., Kyprianidis, J.E., Alexa, M., 2015. Perfect Laplacians for polygon meshes. *Computer Graphics Forum* 34, 211–218.

- [23] Hou, T., Qin, H., 2012. Continuous and discrete Mexican hat wavelet transforms on manifolds. *Graphical Models* 74, 221–232.
- [24] Kim, B., Rossignac, J., 2005. Geofilter: Geometric selection of mesh filter parameters. *Computer Graphics Forum* 24, 295–302.
- [25] Lafon, S., Keller, Y., Coifman, R.R., 2006. Data fusion and multicue data matching by diffusion maps. *IEEE Trans. on Pattern Analysis Machine Intelligence* 28, 1784–1797.
- [26] Lehoucq, R., Sorensen, D.C., 1996. Deflation techniques for an implicitly restarted Arnoldi iteration. *SIAM Journal of Matrix Analysis and Applications* 17, 789–821.
- [27] Liao, S., Tong, R., Dong, J., Zhu, F., 2009. Gradient field based inhomogeneous volumetric mesh deformation for maxillofacial surgery simulation. *Computers & Graphics* 33, 424 – 432.
- [28] Lipman, Y., Rustamov, R.M., Funkhouser, T.A., 2010. Biharmonic distance. *ACM Trans. on Graphics* 29, 27:1–27:11.
- [29] Liu, Y., Prabhakaran, B., Guo, X., 2012. Point-based manifold harmonics. *IEEE Trans. on Visualization and Computer Graphics* 18, 1693 –1703.
- [30] Luo, B., Wilson, R.C., Hancock, E.R., 2003. Spectral embedding of graphs. *Pattern Recognition* 36, 2213–2230.
- [31] Mahmoudi, M., Sapiro, G., 2009. Three-dimensional point cloud recognition via distributions of geometric distances. *Graphical Models* 71, 22–31.
- [32] Memoli, F., 2011. A spectral notion of Gromov-Wasserstein distance and related methods. *Applied and Computational Harmonic Analysis* 30, 363 – 401.
- [33] Mèmolì, F., Sapiro, G., 2005. A theoretical and computational framework for isometry invariant recognition of point cloud data. *Foundations of Computational Mathematics* 5, 313–347.

- [34] Ng, A.Y., Jordan, M.I., Weiss, Y., 2001. On spectral clustering: analysis and an algorithm, in: *Advances in Neural Information Processing Systems 14*, MIT Press. pp. 849–856.
- [35] Ovsjanikov, M., Mèrigot, Q., Mèmolli, F., Guibas, L., 2010. One point isometric matching with the heat kernel. *Computer Graphics Forum* 29, 1555–1564.
- [36] Patanè, G., 2013. wFEM heat kernel: discretization and applications to shape analysis and retrieval. *Computer-Aided Geometric Design* 30, 276–295.
- [37] Patanè, G., 2014. Laplacian spectral distances and kernels on 3D shapes. *Pattern Recognition Letters* 47, 102–110.
- [38] Patanè, G., 2016. STAR - Laplacian spectral kernels and distances for geometry processing and shape analysis. *Computer Graphics Forum* 35, 599–624.
- [39] Patanè, G., 2017. Accurate and efficient computation of Laplacian spectral distances and kernels. *Computer Graphics Forum* 36, 184–196.
- [40] Pinkall, U., Polthier, K., 1993. Computing discrete minimal surfaces and their conjugates. *Experimental Mathematics* 2, 15–36.
- [41] Ramani, K., Sinha, A., 2013. Multiscale kernels using random walks. *Computer Graphics Forum* 33, 164–177.
- [42] Reuter, M., Wolter, F.E., Peinecke, N., 2006. Laplace-Beltrami spectra as Shape-DNA of surfaces and solids. *Computer-Aided Design* 38, 342–366.
- [43] Rodolà, E., Cosmo, L., Bronstein, M.M., Torsello, A., Cremers, D., 2017. Partial functional correspondence. *Computer Graphics Forum* 36, 222–236.
- [44] Rosenberg, S., 1997. *The Laplacian on a Riemannian Manifold*. Cambridge University Press.
- [45] Rustamov, R.M., 2011. Multiscale biharmonic kernels. *Computer Graphics Forum* 30, 1521–1531.

- [46] Shi, J., Malik, J., 2000. Normalized cuts and image segmentation. *IEEE Trans. on Pattern Analysis and Machine Intelligence* 22, 888–905.
- [47] Singer, A., 2006. From graph to manifold Laplacian: the convergence rate. *Applied and Computational Harmonic Analysis* 21, 128 – 134.
- [48] Sorensen, D.C., 1992. Implicit application of polynomial filters in a k-step arnoldi method. *SIAM Journal of Matrix Analysis and Applications* 13, 357–385.
- [49] Spielman, D.A., Teng, S.H., 2007. Spectral partitioning works: planar graphs and finite element meshes. *Linear Algebra and its Applications* 421, 284–305.
- [50] Sun, J., Ovsjanikov, M., Guibas, L.J., 2009. A concise and provably informative multi-scale signature based on heat diffusion. *Computer Graphics Forum* 28, 1383–1392.
- [51] Tong, Y., Lombeyda, S., Hirani, A.N., Desbrun, M., 2003. Discrete multiscale vector field decomposition. *ACM Trans. on Graphics* 22, 445–452.
- [52] Varga, R., 1990. Scientific computation on mathematical problems and conjectures. *SIAM, CBMS-NSF regional conference series in applied mathematics*.
- [53] Xiao, B., Hancock, E.R., Wilsonb, R., 2010. Geometric characterization and clustering of graphs using heat kernel embeddings. *Image and Vision Computing* 28, 1003 – 1021.
- [54] Zhu, X., Ghahramani, Z., Lafferty, J., 2003. Semi-supervised learning using gaussian fields and harmonic functions, in: *Intern. Conf. on Machine Learning*, pp. 912–919.