Volume xx (200y), Number z, pp. 1–13

Accurate and Efficient Computation of Laplacian Spectral Distances and Kernels

Giuseppe Patané, CNR-IMATI - Genova, Italy

Abstract

This paper introduces the Laplacian spectral distances, as a function that resembles the usual distance map, but exhibits properties (e.g., smoothness, locality, invariance to shape transformations) that make them useful to processing and analyzing geometric data. Spectral distances are easily defined through a filtering of the Laplacian eigenpairs and reduce to the heat diffusion, wave, bi-harmonic, and commute-time distances for specific filters. In particular, the smoothness of the spectral distances and the encoding of local and global shape properties depend on the convergence of the filtered eigenvalues to zero. Instead of applying a truncated spectral approximation or prolongation operators, we propose a computation of Laplacian distances and kernels through the solution of sparse linear systems. Our approach is free of user-defined parameters, overcomes the evaluation of the Laplacian spectrum, and guarantees a higher approximation accuracy than previous work.

Categories and Subject Descriptors (according to ACM CCS): Computer graphics [Computing methodologies]: Shape modeling—

1. Introduction

This paper introduces the Laplacian spectral distances, as a function that resembles the usual distance map, but exhibits properties (e.g., smoothness, locality, invariance to shape transformations) that make them useful to processing geometric data. Spectral distances are easily defined through a filtering of the Laplacian eigenpairs and include random walks [FPS05, RS13], heat diffusion [BBK*10, BBOG11, CL06, GBAL09, LKC06, LSW09], bi-harmonic [LRF10], and wave kernel [BB11, ASC11] distances. Laplacian spectral distances have been applied to shape segmentation [dGGV08] and comparison [BBOG11, GBAL09, Mem09, OMMG10, SOG09]. In fact, they are intrinsic to the input shape, invariant to isometries, multi-scale, and robust to noise and tessellation. Biharmonic [LRF10] distances provide a trade-off between a nearly geodesic behavior for small distances and the encoding of global surface properties for large distances, thus guaranteeing an intrinsic, multi-scale characterization of the input shape. The heat kernel [BBG94] is also central in diffusion geometry [BN03, CL06, GK06, Sin06], dimensionality reduction [BN03, XHW10], and data classification [SK03].

For a surface sampled with n points, the evaluation of the spectral distances and kernels generally requires a high computational cost, which varies from $\mathcal{O}(n)$ to $\mathcal{O}(n^3)$ time (e.g., QR factorization), according to the sparsity of the Laplacian matrix. Even though iterative solvers of sparse eigenproblems reduce the computational cost to super linear time [VL08], only a part of the Laplacian spectrum is computed. To speed-up the computation, spectral distances are approximated by prolongating their values evaluated on a sub-sampling of the input shape [VBCG10]; applying multi-resolution decompositions [CL06]; or considering only the eigenvectors related to smaller eigenvalues. In these cases, the number of selected eigenpairs and the targeted approximation accuracy are heuristically adapted to shape details and parameters (e.g., time for the heat kernel). This approximation is accurate only if the coefficients decay fast (e.g., large scales in the diffusion kernel and distances); otherwise (e.g., periodic or increasing decaying filters), a larger number of eigenpairs is needed, with a $\mathcal{O}(n^2)$ computational cost. Similarly to signal smoothing [CPS13, DMSB99, KR05, Tau95, TZG96, ZF03], geodesic [CWW13], bi-harmonic [LRF10], and diffusion [HVG11, Pat14, Pat13, PS13] distances have been rewritten in terms of the corresponding kernels, thus bypassing the evaluation of the Laplacian spectrum.

Overview and contribution We propose an accurate, computationally efficient approximation of spectral distances, through the solution of sparse linear systems, which is independent of the evaluation of the Laplacian spectrum. As *(filtered) Laplacian spectral distance*, we refer to *any* dis-

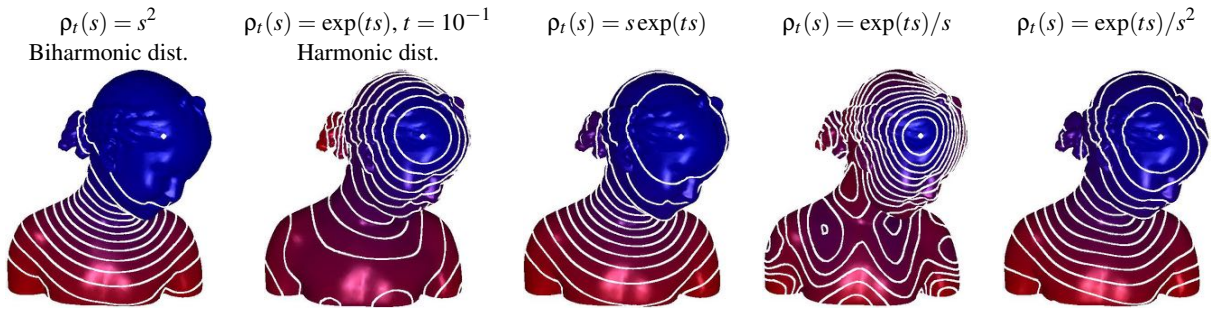

Figure 1: *Level-sets of the spectral distances from a source point (white dot) induced by the filter* ρ *and evaluated with the Padé-Chebyshev approximation* ($r = 5$)*. The level-sets are associated to iso-values uniformly sampled in the range of the distances, whose minimum and maximum are depicted in blue and red. The color coding represents the same scale for multiple images.*

tance $d^2(\mathbf{p}_i, \mathbf{p}_j) := \sum_{l=1}^{n} |x_l(\mathbf{p}_i) - x_l(\mathbf{p}_j)|^2/\rho^2(\lambda_l)$, which is a combination of the Laplacian eigenpairs $\{(\lambda_l, \mathbf{x}_l)\}_{l=1}^{n}$ of a surface or volume $\mathcal{N}$ sampled with n points $\mathcal{P} := \{\mathbf{p}_i\}_{i=1}^{n}$ and induced by a positive filter $\rho : \mathbb{R}^+ \to \mathbb{R}$. With respect to previous work on Laplacian distances and smoothing based on a filtering of the Laplacian spectrum, our approach takes this idea further by incorporating the Padé-Chebyshev approximation, which allows us to define more general distances. Furthermore, we provide an extensive analysis of error behavior in order to support a practical approach to the computation of the spectral distances. Finally, our work ties rational filtering to distance computation and demonstrates that it pays-off in terms of efficiency and stability.

The proposed definition of the spectral distance includes the bi-harmonic [LRF10] ($\rho(s) := s^2$), wave [BB11, ASC11] ($\rho_t(s) := \exp(ist)$), diffusion [CWW13, HVG11, Pat13] ($\rho_t(s) := \exp(st)$) distances and Mexican hat wavelets [HQ12] ($\rho(s) := s\exp(s^2)$). The filter map is defined according to different properties of the resulting distances, such as isometry invariance, localization in both space and frequency [HVG11], or through a supervised learning [ABBK11]. The selection of ρ also provides a simple way to design new distances, whose smoothness and balance between the measure of both local and global properties depend on the convergence of the filtered eigenvalues $(1/\rho(\lambda_i))_{i=1}^{n}$ to zero or on their periodic behavior. Increasing the growth of the filter ρ, the influence of the larger eigenvalues and of the corresponding eigenvectors on the filtered spectral distance is negligible with respect to the contribution of the lower eigenvalues. The resulting distance characterizes the global shape properties and local features are encoded by reducing the filter values.

Our representation (Sect. 2) of the spectral distances is equivalent to solving r sparse and symmetric linear systems, where r is the degree of the rational polynomial approximation of the filter. In this way, the proposed computation (Sect. 3) is independent of the evaluation of the Laplacian spectrum, is robust to shape discretization, and applies to both surfaces and volumes. Furthermore, it is free of user-defined parameters (e.g., the selection of the eigenpairs with respect to the approximation accuracy), which are used by the truncated spectral approximation and multi-resolution prolongation operators.

Neglecting round-off errors, the computation of spectral kernels and distances induced by polynomial filter maps is exact (e.g., for bi-harmonic kernel [Rus11] and distances [LRF10]). For arbitrary (e.g., heat diffusion, wave kernel) distances, the approximation accuracy is estimated a-priori in terms of the polynomial degree r and it can be further reduced by slightly increasing r. For instance, the accuracy of the diffusion distances computed through the Padé-Chebyshev approximation is lower than 10^{-r}, where $r = 5, 7$ is the degree of the rational polynomial. Applying an iterative linear solver (e.g., Gauss-Seidel method, conjugate gradient) [GV89] (Ch. 10) that takes $\mathcal{O}(\tau(n))$ time, the computational cost for the evaluation of any Laplacian spectral kernel and distance between two points is $\mathcal{O}(r\tau(n))$. Here, the function $\tau(n)$, which depends on the number n of shape samples and the sparsity of the coefficient matrix, typically varies from $\tau(n) = n$ to $\tau(n) = n \log n$. Since we solve r sparse, symmetric linear systems with n different right-hand vectors, our computation of the one-to-all distances takes $\mathcal{O}(rn\tau(n))$ time, which is comparable with the $\mathcal{O}(kn^2)$ time of the truncated spectral approximation with k Laplacian eigenpairs.

According to our tests (Sects. 4, 5), the proposed approach is competitive with respect to previous work for the evaluation of spectral distances, especially those ones induced by slowly-increasing (e.g., diffusion distances at small scales) or periodic filters and provides a higher approximation accuracy. Our computation is also useful to evaluate the spectral distances on densely sampled surfaces or among seed points sampled on the input surface, which is typical in the evaluation of shape descriptors [OFCD02] and bags-of-features [BB11, BBOG11]. In this case, the number of seed points is much lower than the number of samples and the

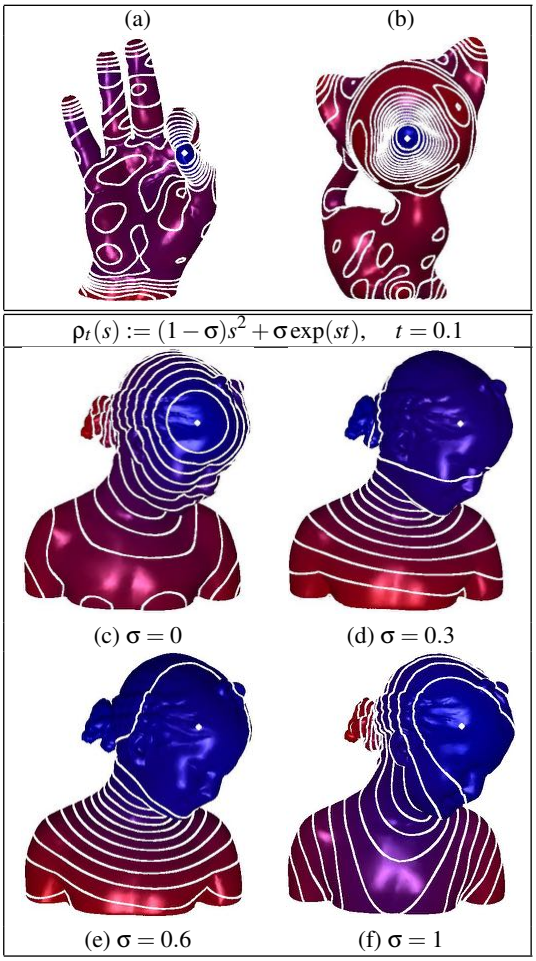

Figure 2: *Spectral distances induced by: (a,b) the filter in [HVG11] ($r = 5$); (c-f) a convex combination of bi-harmonic ($\sigma = 0$) and diffusion ($\sigma = 1$, $t = 0.1$) distances ($r = 7$).*

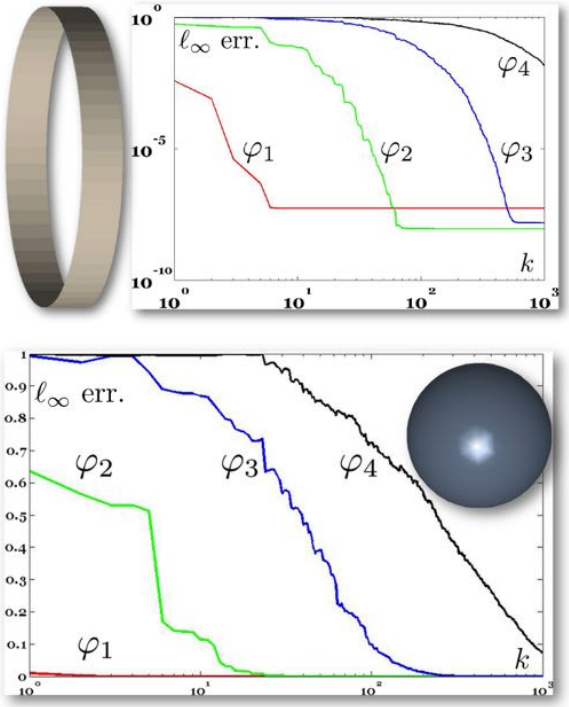

Figure 3: *ℓ_∞ error (y-axis) between the ground-truth distances induced by the filters $\{\rho_i\}_i$ and the truncated approximation with k (x-axis) eigenpairs. For the Padé-Chebyshev method ($r = 5$) and all the filters, the ℓ_∞ error with respect to the ground-truth is lower than 6.5×10^{-6}.*

higher accuracy improves the discrimination capabilities of shape descriptors based on spectral distances.

2. Spectral distances and kernels

Let $\mathcal{L}_2(\mathcal{N})$ be the space of square integrable maps on the manifold $\mathcal{N}$ and $\{(\lambda_i, \phi_i)\}_{i=0}^{+\infty}$ the orthonormal Laplacian eigensystem $\Delta\phi_i = \lambda_i\phi_i$, $0 \le \lambda_i \le \lambda_{i+1}$. Given a strictly positive *filter map* $\rho : \mathbb{R}^+ \to \mathbb{R}$, the *spectral distances* on $\mathcal{N}$ are defined through a filtering of the Laplacian spectrum as

$$d^2(\mathbf{p},\mathbf{q}) = \sum_{n=0}^{+\infty} \frac{|\phi_n(\mathbf{p}) - \phi_n(\mathbf{q})|^2}{\rho^2(\lambda_n)}. \tag{1}$$

According to [Hoe68, Sog88], the Laplacian eigenfunction ϕ_n on a 2-manifold $\mathcal{N}$ and associated with the eigenvalue λ_n, $\lambda_n \neq 0$, satisfies the upper bound $\|\phi_n\|_\infty \le C\lambda_n^{1/4}\|\phi_n\|_2$, where C is a constant that depends only on the geometric properties (i.e., sectional curvature, injectivity radius) of $\mathcal{N}$. Selecting eigenfunctions with unitary norm, the spectral distance satisfies the upper bound

$$d^2(\mathbf{p},\mathbf{q}) \le 4\sum_{n=0}^{+\infty} \frac{\|\phi_n\|_\infty^2}{\rho^2(\lambda_n)} \le 4C^2 \sum_{n=0}^{+\infty} \frac{\lambda_n^{1/2}}{\rho^2(\lambda_n)}.$$

If the map $\tilde{\rho}(s) := s^{1/2}/\rho^2(s)$ is integrable on $\mathbb{R}^+$, then the series that defines the spectral distance is convergent. The nullity condition $d(\mathbf{p},\mathbf{q}) = 0$ holds if and only if $\phi_i(\mathbf{p}) = \phi_i(\mathbf{q})$, $\forall i \in \mathbb{N}$; i.e., $\langle 1_\mathbf{p} - 1_\mathbf{q}, \phi_i\rangle_2 = 0$, where $1_\mathbf{p}(\cdot)$ has value one at $\mathbf{p}$ and zero otherwise. Noting that $\{\phi_i\}_{i=0}^{+\infty}$ is a basis of $\mathcal{L}_2(\mathcal{N})$, the nullity relation is satisfied if and only if $(1_\mathbf{p} - 1_\mathbf{q})$ is the null function; i.e., $\mathbf{p} = \mathbf{q}$. The symmetry and triangular inequality follow from Eq. (1).

In a similar way, we define the *filtered kernel*

$$K(\mathbf{p},\mathbf{q}) := \sum_{n=0}^{+\infty} \frac{\phi_n(\mathbf{p})\phi_n(\mathbf{q})}{\rho(\lambda_n)}, \tag{2}$$

which is related to the spectral distances through the identity

Algorithm 1 Computation of the spectral distances.

Require: A surface or volume $\mathcal{N}$, a filter map $\rho : \mathbb{R} \to \mathbb{R}$.
Ensure: The spectral distance $d(\mathbf{p}_i, \mathbf{p}_j)$ in Eq. (1), $\mathbf{p}_i, \mathbf{p}_j \in \mathcal{N}$.
1: Compute $(\mathbf{L}, \mathbf{B})$, which define the Laplacian $\tilde{\mathbf{L}} := \mathbf{B}^{-1}\mathbf{L}$.
2: Define the vector $\mathbf{f} = \mathbf{e}_i - \mathbf{e}_j$.
3: *CASE I - Arbitrary filter: polynomial approximation*
4: Compute the polynomial approx. $p_r(s) = \sum_{i=0}^{r} \alpha_i s^i$ of $1/\rho$.
5: Compute $\mathbf{g}_1$: $\mathbf{B}\mathbf{g}_1 = \mathbf{L}\mathbf{f}$.
6: **for** $i = 1, \ldots, r-1$ **do**
7: Compute $\mathbf{g}_{i+1}$: $\mathbf{B}\mathbf{g}_{i+1} = \mathbf{L}\mathbf{g}_i$
8: **end for**
9: Compute $\mathbf{u} = \mathbf{K}_{1/\rho}\mathbf{f} \approx p_r(\tilde{\mathbf{L}}) = \alpha_0\mathbf{f} + \sum_{i=1}^{r} \alpha_i \mathbf{g}_i$ (c.f., Eq. (5)).
10: Compute the distance $d(\mathbf{p}_i, \mathbf{p}_j) = \|\mathbf{u}\|_{\mathbf{B}}$ (c.f., Eq. (3)).
11: *CASE II - Arbitrary filter: Padé-Chebyshev approximation*
12: Compute the P.C. approx. $p_r(s) = \sum_{i=1}^{r} \alpha_i (1 + \beta_i s)^{-1}$ of $1/\rho$.
13: **for** $i = 1, \ldots, r$ **do**
14: Compute $\mathbf{g}_i$: $(\mathbf{B} + \beta_i \mathbf{L})\,\mathbf{g}_i = \mathbf{B}\mathbf{f}$ (c.f., Eq. (6))
15: **end for**
16: Compute $\mathbf{u} = \mathbf{K}_{1/\rho}\mathbf{f} \approx p_r(\tilde{\mathbf{L}})\mathbf{f} = \sum_{i=1}^{r} \alpha_i \mathbf{g}_i$.
17: Compute the distance $d(\mathbf{p}_i, \mathbf{p}_j) = \|\mathbf{u}\|_{\mathbf{B}}$ (c.f., Eq. (3)).

$d(\mathbf{p},\mathbf{q}) = \|K(\mathbf{p},\cdot) - K(\mathbf{q},\cdot)\|_2$. From the upper bound

$$K(\mathbf{p},\mathbf{q}) \leq \sum_{n=0}^{+\infty} \frac{\|\phi_n\|_\infty^2}{\rho(\lambda_n)} \leq C^2 \sum_{n=0}^{+\infty} \frac{\lambda_n^{1/2}}{\rho(\lambda_n)},$$

we get that the integrability of $\tilde{\rho}(s) := s^{1/2}/\rho(s)$ on $\mathbb{R}^+$ guarantees the well-posedness of the spectral kernel. The symmetry and self-adjointness of the kernel follow from its definition.

In particular, if the Laplacian eugenfunctions are uniformly bounded (i.e., $\|\phi_n\|_\infty \leq C_\infty, \forall n \in \mathbb{N}$) then the integrability of $1/\rho$ guarantees the well-posedness of the spectral kernel and distance. This hypothesis applies only to the continuous case; in fact, their discrete counterparts (c.f., Eqs. (3), (4)) are defined through a finite sum.

Discrete spectral distances On the space of maps $f : \mathcal{P} \to \mathbb{R}$, $\mathbf{f} := (f(\mathbf{p}_i))_{i=1}^n$, on the point set $\mathcal{P} := \{\mathbf{p}_i\}_{i=1}^n$, we represent the Laplacian matrix as $\tilde{\mathbf{L}} := \mathbf{B}^{-1}\mathbf{L}$, where $\mathbf{B}$ is a positive definite matrix (e.g., mass matrix [RWP06, VL08], diagonal matrix of Voronoi areas [DMSB99]) and $\mathbf{L}$ is symmetric, positive semi-definite (e.g., cotangent [PP93, AW11] or exponential weights [BSW08, BN03, BN08, CL06, LPG12] on polygonal meshes or point sets) (Sect. 4). Let $\mathbf{L}\mathbf{X} = \mathbf{B}\mathbf{X}\Lambda$ be the generalized orthonormal eigensystem, $\mathbf{X}^\top\mathbf{B}\mathbf{X} = \mathbf{I}$, where $\mathbf{X} := [\mathbf{x}_1, \ldots, \mathbf{x}_n]$ and $\Lambda := \mathrm{diag}(\lambda_i)_{i=1}^n$ are the eigenvectors' and eigenvalues' matrices. Let $\mathbf{e}_i$ be the canonical vector whose i-th entry is 1 and 0 otherwise. Rewriting Eq. (2) as $\langle K(\mathbf{p},\cdot), \phi_l\rangle_2 = \langle \phi_l, 1_\mathbf{p}\rangle_2/\rho(\lambda_l)$ and sampling it on $\mathcal{P}$, the weak formulation $\langle K(\mathbf{p}_i,\cdot), \phi_l\rangle_2 = \langle \phi_l, \mathbf{e}_i\rangle_2/\rho(\lambda_l)$ of the kernel is equivalent to

$$\langle \mathbf{K}_{1/\rho}\mathbf{e}_i, \mathbf{x}_l\rangle_\mathbf{B} = \frac{\langle \mathbf{x}_l, \mathbf{e}_i\rangle_\mathbf{B}}{\rho(\lambda_l)} \iff \mathbf{K}_{1/\rho}\mathbf{e}_i = \sum_{l=1}^{n} \frac{\langle \mathbf{x}_l, \mathbf{e}_i\rangle_\mathbf{B}}{\rho(\lambda_l)}\mathbf{x}_l,$$

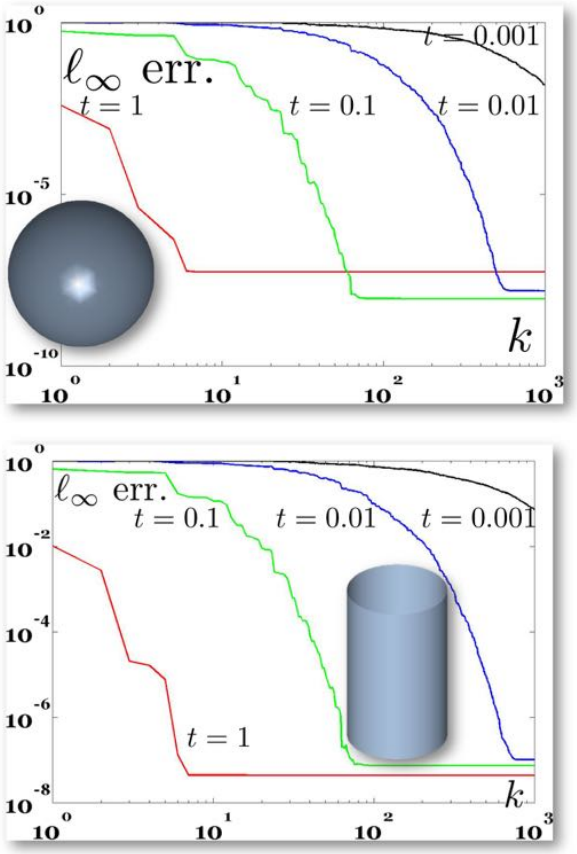

Figure 4: *ℓ_∞ error (y-axis) between the diffusion distance computed with the truncated spectral approximation with k (x-axis) Laplacian eigenpairs. For the Padé-Chebyshev method ($r = 5$) and all the scales, the ℓ_∞ error with respect to the ground-truth is lower than 8.9×10^{-6}.*

where $\langle \mathbf{f}, \mathbf{g}\rangle_\mathbf{B} := \mathbf{f}^\top\mathbf{B}\mathbf{g}$ is the inner product induced by $\mathbf{B}$. Indeed, the *filtered kernel matrix* is $\mathbf{K}_{1/\rho} = \mathbf{X}\rho^\dagger(\Lambda)\mathbf{X}^\top\mathbf{B}$, $\rho^\dagger(\Lambda) := \mathrm{diag}\,(1/\rho(\lambda_i))_{i=1}^n$, and the *spectral distances* are

$$d^2(\mathbf{p}_i, \mathbf{p}_j) = \|\mathbf{K}_{1/\rho}(\mathbf{e}_i - \mathbf{e}_j)\|_\mathbf{B}^2 = \sum_{l=1}^{n} \frac{|\langle \mathbf{x}_l, \mathbf{e}_i - \mathbf{e}_j\rangle_\mathbf{B}|^2}{\rho^2(\lambda_l)}. \quad (3)$$

If $\rho(\lambda_i) = 0$, then the corresponding entry in $\rho^\dagger(\Lambda)$ is chosen equal to zero. In the last equality of Eq. (3), we have applied the identity $\mathbf{K}_{1/\rho}^\top\mathbf{B}\mathbf{K}_{1/\rho} = \mathbf{B}\mathbf{X}[\rho^\dagger(\Lambda)]^2\mathbf{X}^\top\mathbf{B}$.

Filtered kernel and pseudo-inverse Writing the filter map as the power series $\rho(s) = \sum_{n=0}^{+\infty} \alpha_n s^n$ and noting that $\tilde{\mathbf{L}} = \mathbf{X}\Lambda\mathbf{X}^\top\mathbf{B}$, we get that

$$\rho(\tilde{\mathbf{L}}) = \mathbf{X}\sum_{n=0}^{+\infty} \alpha_n \Lambda^n \mathbf{X}^\top\mathbf{B} = \mathbf{X}\rho(\Lambda)\mathbf{X}^\top\mathbf{B} = \mathbf{K}_\rho; \quad (4)$$

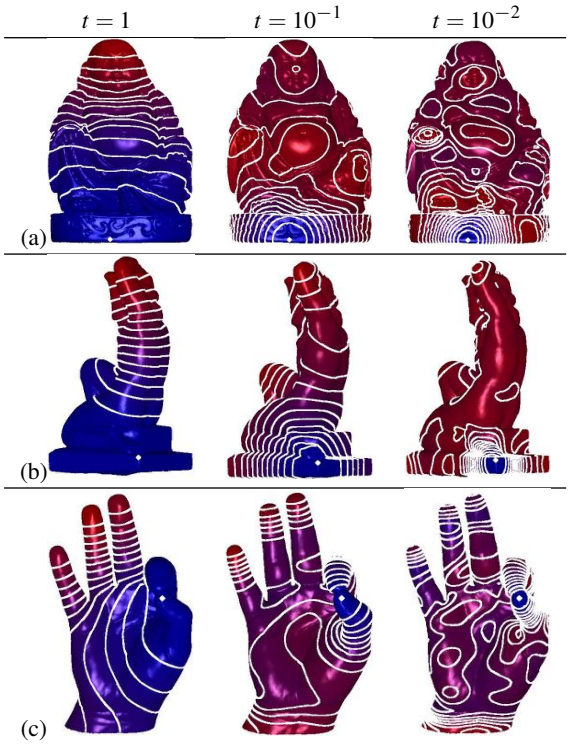

Figure 5: *Diffusion distances at different scales $t = 10^{-k}$, $k = 0,1,2$ (from left to right), and at the higher resolution ($n = 115K$); computational costs are reported in Table 1. The degree of the Padé-Chebyshev polynomial is $r = 5$.*

i.e., the spectral kernel $\mathbf{K}_\rho = \rho(\tilde{\mathbf{L}})$ is a filtered version of the Laplacian matrix. Furthermore, $\mathbf{K}_\rho$ is the pseudo-inverse of $\mathbf{K}_{1/\rho}$; i.e., $\mathbf{K}_\rho\mathbf{K}_{1/\rho}\mathbf{K}_\rho = \mathbf{K}_\rho$, $\mathbf{K}_{1/\rho}\mathbf{K}_\rho\mathbf{K}_{1/\rho} = \mathbf{K}_{1/\rho}$, and the matrices $\mathbf{K}_\rho\mathbf{K}_{1/\rho}$, $\mathbf{K}_{1/\rho}\mathbf{K}_\rho$ are $\mathbf{B}$-adjoint. In fact,

$$\mathbf{K}_\rho\mathbf{K}_{1/\rho}\mathbf{K}_\rho = \mathbf{X}\rho(\Lambda)\rho^\dagger(\Lambda)\rho(\Lambda)\mathbf{X}^\top\mathbf{B} = \mathbf{X}\rho(\Lambda)\mathbf{X}^\top\mathbf{B} = \mathbf{K}_\rho,$$
$$(\mathbf{K}_\rho\mathbf{K}_{1/\rho})^\top\mathbf{B} = \mathbf{B}\mathbf{X}\rho^\dagger(\Lambda)\rho(\Lambda)\mathbf{X}\mathbf{B} = \mathbf{B}(\mathbf{K}_\rho\mathbf{K}_{1/\rho}).$$

3. Distance computation

In the paper examples, the level-sets of the spectral distances are associated with iso-values uniformly sampled in the range of the spectral distances, whose minimum and maximum are depicted in blue and red, respectively. The color coding represents the same scale of values for multiple shapes. For the computation of the spectral distances, we distinguish two main cases, according to the polynomial or arbitrary representation of the filter.

If $\rho(s) := \sum_{i=0}^{r} \alpha_i s^i$ is a polynomial filter (e.g., commute-time, poly-harmonic [LRF10] distances), then $\mathbf{K}_\rho = \rho(\tilde{\mathbf{L}})$ is computed from the representation of ρ and $\tilde{\mathbf{L}}$. Applying Eq. (3) and the relation $\mathbf{K}_\rho^\dagger = \mathbf{K}_{1/\rho}$, the distance $d(\mathbf{p}_i,\mathbf{p}_j) = \|\mathbf{u}\|_\mathbf{B}$ is equal to the norm of the solution $\mathbf{u}$ to the linear system $\mathbf{K}_\rho\mathbf{u} = \mathbf{e}_i - \mathbf{e}_j$. If ρ is an arbitrary filter (e.g., diffusion, wave kernel distances), then it is approximated with a polynomial or rational polynomial of degree r and the evaluation of the corresponding distance reduces to the solution of r sparse, symmetric linear systems. Each linear system is solved with the *minres* procedure [GV89] (Ch. 10), which computes a minimum norm residual solution.

Table 1: *Accuracy and time (in seconds) of the diffusion distances with the (i) truncated spectral approximation (TSA) with $k = 50$ Laplacian eigenpairs, (ii) multi-resolution prolongation operators (MPO), (iii) Padé-Chebyshev (P-C) method ($r = 5$). Here, t is the scale, n is the number of vertices at a given resolution, RMSE is the root mean square deviation, and '$\times$' is the reduction rate of the computational cost of the Padé-Chebyshev method with respect to multi-resolution prolongation operators (i.e., MPO/P-C).*

t	n K	TSA	MPO	RMSE $\times 10^{-4}$	P-C	RMSE $\times 10^{-6}$	$\times$
Fig. 5(a)							
10^{-3}	115	6.85	26.17	0.2	7.04	1.1	**3.7**
10^{-2}	150	10.34	20.42	0.6	14.21	2.5	**1.4**
10^{-1}	7	1.52	19.51	19	1.23	7.4	**15.8**
1	6	1.27	15.50	5.7	0.98	5.5	**15.8**
Fig. 5(b)							
10^{-3}	115	5.58	24.17	0.3	8.42	4.1	**2.8**
10^{-2}	150	11.46	21.02	0.7	13.12	2.2	**1.6**
10^{-1}	7	1.68	18.12	2.7	1.20	7.3	**15.1**
1	6	1.16	16.19	7.3	0.89	1.5	**18.1**
Fig. 5(c)							
10^{-3}	115	6.38	20.13	0.2	7.57	5.2	**2.6**
10^{-2}	150	12.46	18.01	0.1	11.25	1.1	**1.6**
10^{-1}	7	0.98	17.11	0.6	1.61	5.2	**10.6**
1	6	1.42	12.18	1.4	1.01	0.1	**12.0**

Arbitrary filter: polynomial approximation Let $[0,\lambda]$ be an interval that contains the spectrum of $\tilde{\mathbf{L}}$, where λ is the maximum eigenvalue, which is computed by the Arnoldi method [GV89], or is set equal to the upper bound [LS96, Sor92]

$$\lambda_n \leq \min\{\max_i\{\sum_j \tilde{L}(i,j)\}, \max_j\{\sum_i \tilde{L}(i,j)\}\}.$$

Computing the best r-degree polynomial approximation $p_r(s) = \sum_{i=0}^{r} \alpha_i s^i$ of ρ with respect to the $\mathcal{L}_\infty$ norm (Fig. 1, $\rho(s) = s^k \exp(ts)$, $k \in \{-2,-1,1\}$, $r = 7$) and introducing the vector $\mathbf{f} = \mathbf{e}_i - \mathbf{e}_j$, Eq. (3) is rewritten as

$$\mathbf{u} = \mathbf{K}_{1/\rho}\mathbf{f} \approx p_r(\tilde{\mathbf{K}})\mathbf{f} = \alpha_0\mathbf{f} + \sum_{i=1}^{r} \alpha_i \mathbf{g}_i, \qquad (5)$$

where $\mathbf{g}_i := (\mathbf{B}^{-1}\mathbf{L})^i\mathbf{f}$ is calculated recursively. At the first iteration $\mathbf{B}\mathbf{g}_1 = \mathbf{L}\mathbf{f}$ and at the $(i+1)$-th step we apply the relation $\mathbf{g}_{i+1} := (\mathbf{B}^{-1}\mathbf{L})^{i+1}\mathbf{f} = \mathbf{B}^{-1}\mathbf{L}\mathbf{g}_i$; i.e., $\mathbf{g}_{i+1}$ solves the sparse, symmetric, positive-definite linear system $\mathbf{B}\mathbf{g}_{i+1} = \mathbf{L}\mathbf{g}_i$, $i = 1,\ldots,r-1$.

Arbitrary filter: Padé-Chebyshev approximation For an arbitrary filter (Fig. 1), we consider the rational Padé-

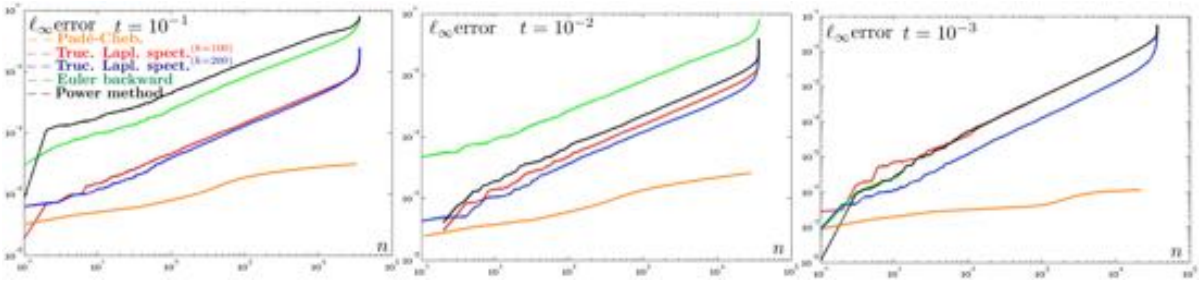

Figure 6: *ℓ_∞ error (y-axis) between the ground-truth diffusion distances on the cylinder, with a different sampling (x-axis). For different scales, the accuracy of the Padé-Chebyshev method ($r = 5$, orange) remains almost unchanged and higher than the truncated approximation with* 100 *and* 200 *eigenpairs (red, blue), the Euler backward (green) and power (black) methods.*

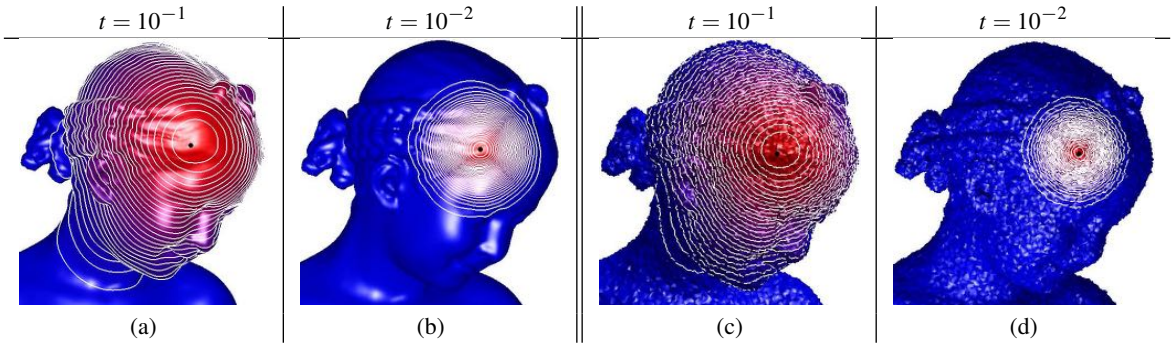

Figure 7: *Robustness of the Padé-Chebyshev approximation ($r = 7$) of the (a,c) diffusion kernel $\mathbf{K}_t\mathbf{e}_i$ and (b,d) distance at $\mathbf{p}_i$ (black dot) on a smooth and noisy triangulated surface.*

Chebyshev approximation $p_r(s) = \frac{a_r(s)}{b_r(s)}$ of $1/\rho$ [GV89] (Ch. 11) with respect to the $\mathcal{L}_\infty$ norm. Here, $a_r(\cdot)$ and $b_r(\cdot)$ are polynomials of degree equal to or lower than r. Let $p_r(s) = \sum_{i=1}^{r} \alpha_i (1+\beta_i s)^{-1}$ be the partial form of the Padé-Chebyshev approximation, where $(\alpha_i)_{i=1}^r$ are the weights and $(\beta_i)_{i=1}^r$ are the nodes of the r-point Gauss-Legendre quadrature rule [GV89] (Ch. 11). The weights and nodes are precomputed for any degree of the rational polynomial [CRV84]. Applying this approximation to the spectral kernel, we get that

$$\mathbf{u} = \mathbf{K}_{1/\rho}\mathbf{f} \approx p_r(\tilde{\mathbf{L}})\mathbf{f} = \sum_{i=1}^{r} \alpha_i \left(\mathbf{I} + \beta_i \tilde{\mathbf{L}}\right)^{-1} \mathbf{f} = \sum_{i=1}^{r} \alpha_i \mathbf{g}_i,$$

where $\mathbf{g}_i$ solves the symmetric and sparse linear system

$$(\mathbf{B} + \beta_i \mathbf{L})\,\mathbf{g}_i = \mathbf{B}\mathbf{f}, \quad i = 1, \ldots, r. \tag{6}$$

The Padé-Chebyshev approximation generally provides an accuracy higher than the polynomial approximation, as a matter of its uniform convergence to the filter.

Applying the Padé-Chebyshev rational polynomial $p_r(s) = \alpha_0 + \sum_{i=1}^{r} \alpha_i/(s-\theta_i)$ of the exponential map to the heat kernel [HVG11, Pat13] $\mathbf{K}_t = \exp(-t\tilde{\mathbf{L}})$ (Fig. 1), Eq. (6) reduces to $(t\mathbf{L} + \theta_i \mathbf{B})\mathbf{g}_i = \alpha_i \mathbf{B}\mathbf{f}$, $i = 1, \ldots, r$. This approximation also provides an efficient way to compute geodesics on surfaces, through the relation between geodesics and heat kernel recently presented in [CWW13]. In spectral graph theory [OSV12], the Padé-Chebyshev and the Lanczos methods have been applied to the approximation of $\exp(-\mathbf{A})\mathbf{f}$, where $\mathbf{A}$ is a symmetric and positive semi-definite matrix. Finally, these approximations are included in numerical libraries for signal processing.

4. Results and discussion

We now discuss the filter selection (Sect. 4.1), the approximation accuracy (Sect. 4.2), stability (Sect. 4.3), and computational cost (Sect. 4.4).

4.1. Selection of the filter map

Increasing the growth of the filter ρ (i.e., the decay of the filtered eigenvalues $(1/\rho(\lambda_i))_{i=1}^n$ to zero), the effects of the larger eigenvalues and of the corresponding eigenvectors on the spectral distance are negligible with respect to the contribution of the lower eigenvalues. The resulting distance encodes the global shape properties; e.g., increasing the power k of $\rho(s) = s^k \log(1+s)$ in Fig. 1. Reducing the filter growth, local shape features are better characterized.

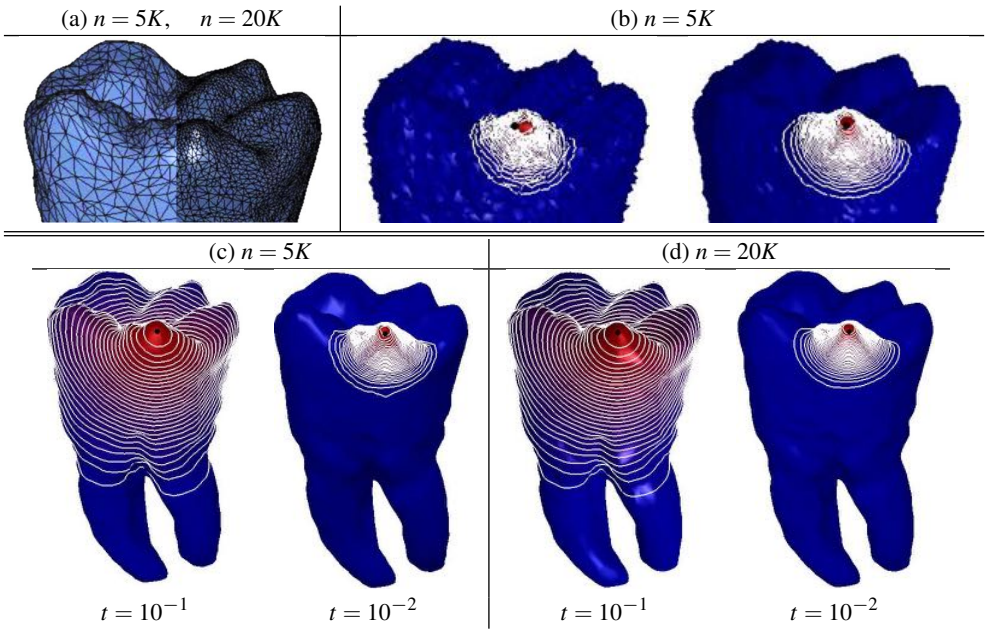

Figure 8: *(a) Input shapes at different resolution. Robustness of the Padé-Chebyshev approximation ($r = 5$) of the heat kernel at different scales ($t = 10^{-1}, 10^{-2}$) with respect to (b) noise and (c,d) surface sampling ($n = 5K, 20K$).*

Similarly to random walks [FPS05, RS13], we consider the filter map $\rho_t(s) = t^{-k} s^{-\alpha} \exp(ts^{\alpha})$, where k scales the diffusion rate and α controls the distance smoothness (Fig. 1). The filters $\rho_t(s) := [\cos^{-1}(\sqrt{s}t), s^{-1/2}\sin(\sqrt{s}t)]$ and $\rho(s,t) = \exp(s^r t)$ are associated with the diffusion equations $\left(\partial_t^2 + \Delta\right) F(\cdot,t) = 0$ and $(\partial_t + \Delta^r) F(\cdot,t) = 0$, respectively. To achieve localization of the signal content in both space and frequency [HVG11], we select a filter map ρ that has a power growth $\rho(s) = s_1^{-a} s^a$ on the interval $[0, s_1)$ and a power decay $\rho(s) = s_2^b s^{-b}$ in $(s_2, +\infty)$, $s_1 < s_2$, for large s. In $[s_1, s_2]$, ρ is a cubic polynomial such that the filter map and its first derivative are continuous on $\mathbb{R}^+$. A stronger or weaker encoding of surface details is achieved by enlarging or reducing the interval size. A common choice is $a = b = 1$, $s_1 = 1$, $s_2 = 2$, $\rho(s) = -5 + 11s - 6s^2 + s^3$ (Fig. 2(a,b)). In [ABBK11], a spline filter is defined through a supervised learning process that discriminates among shapes of a certain class and is insensitive to a selected family of transformations. Finally, for the definition of new distances we can consider a convex combination of the filters (Fig. 2(c-f)).

4.2. Approximation accuracy

Let λ be an upper bound to the Laplacian eigenvalues, p_r the polynomial or Padé-Chebyshev rational polynomial of degree r that approximates the filter map, and C_∞ the upper bound to the approximation between ρ and p_r. Since

$$\|\mathbf{K}_{1/\rho}\mathbf{f} - p_r(\tilde{\mathbf{L}})\mathbf{f}\|_{\mathbf{B}}^2 \leq \sum_{i=1}^{n} |\rho^{-1}(\lambda_i) - p_r(\lambda_i)|^2 |\langle \mathbf{f}, \mathbf{x}_i \rangle_{\mathbf{B}}|^2 \leq C_\infty^2 \|\mathbf{f}\|_{\mathbf{B}}^2,$$

the approximation error is controlled by the approximation accuracy of the filter map. For the polynomial approximation $C_\infty = \frac{n}{(r+1)!}\left[\frac{\lambda_{\max}(\mathbf{L})}{\lambda_{\min}(\mathbf{B})}\right]^{r+1} \|\rho^{(r+1)}(\tilde{\mathbf{L}})\|_2$ and for the Padé-Chebyshev approximation $C_\infty \leq 10^{-r}$. For the exponential filter [Var90], the ℓ_2 error between $\exp(-t\tilde{\mathbf{L}})$ and its rational approximation $p_r(t\tilde{\mathbf{L}})$ is lower than the uniform Chebyshev constant σ_{rr}. Since this constant is known, independent of t, and related to the degree of the rational Padé-Chebyshev approximation by the relation $\sigma_{rr} \approx 10^{-r}$, the degrees $r = 5, 7$ provide an error that is satisfactory for the evaluation of diffusion distances. If necessary, the degree is easily tuned to improve the approximation accuracy.

For the evaluation of the approximation accuracy, we compare the ground-truth distances on cylinders and spheres with the proposed approach and the truncated spectral approximation, where the Laplacian eigenpairs have been computed through a variant [VL08] of the Arnoldi iteration method [LS96, Sor92]. According to [RBG*09], we consider a rectangular domain with edge length $a = 1$, $b = 2$ and the corresponding isometric cylinder $(\cos x, y, \sin x)$. Introducing Neumann boundary conditions, their Laplacian eigenpairs are $\phi_{m,n}(x,y) := (\cos(\frac{m\phi}{a}x), \cos(\frac{n\phi}{b}y))$, $m, n \in \mathbb{N}$, and

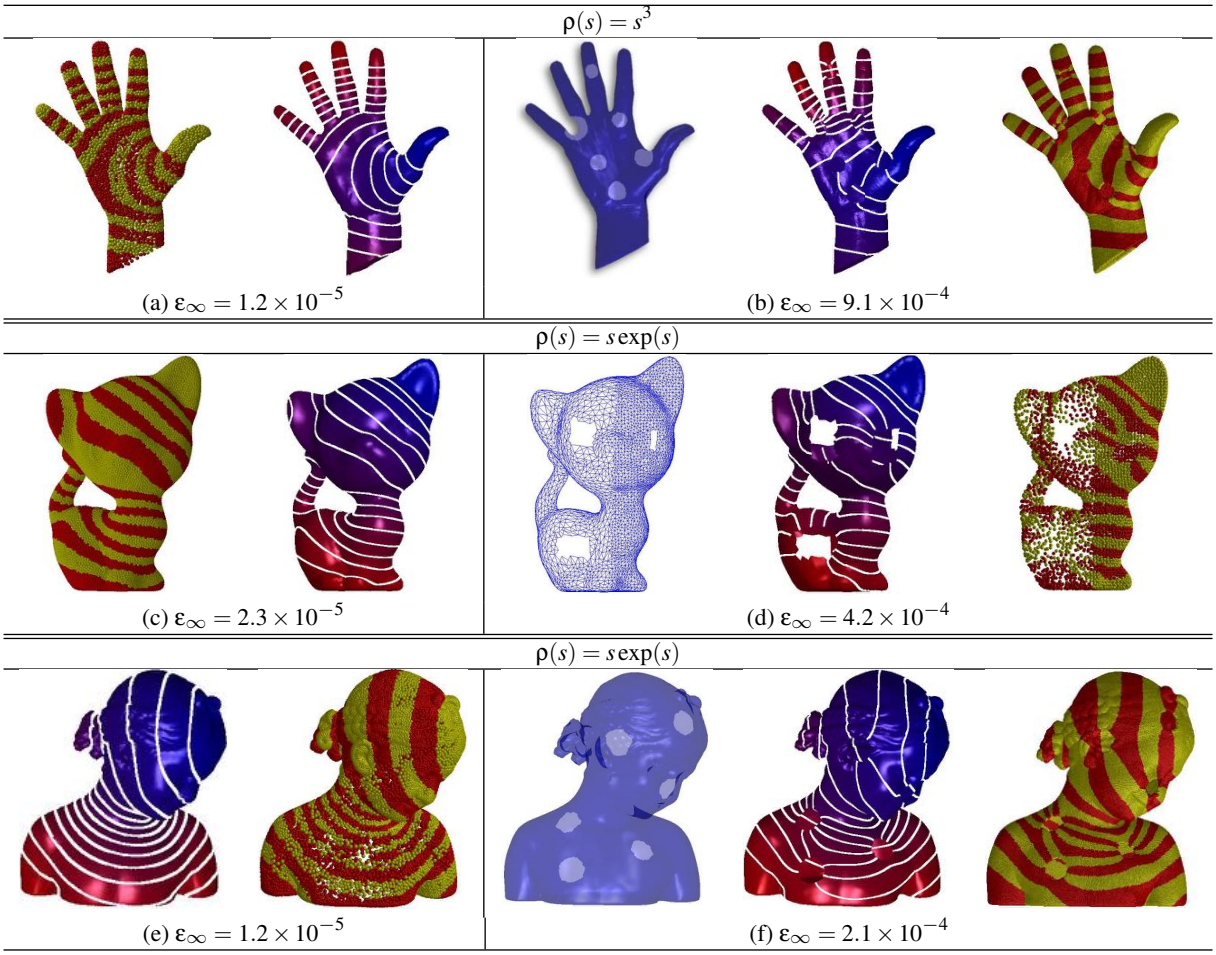

Figure 9: *Distances computed with the Padé-Chebyshev method ($r = 5$) on (a,c,e) regularly-sampled and (b,d,f) irregularly-sampled (left) meshes and (right) point sets with holes. To improve the visualization, points are represented as spheres.*

$\lambda_{m,n} := \pi^2(m^2/a^2 + n^2/b^2)$. For the sphere, the spherical harmonics are $\phi_{m,n}(\theta,\rho) = N\exp(im\rho)p_l^m(\cos\theta)$, where N is a normalization constant and $p_l^m(\cdot)$ is an associated Legendre function. Since we have an infinite number of eigenpairs, we select k such that the spectral distance $d_k(\mathbf{p},\mathbf{q}) := \sum_{n=0}^{k} \frac{|\phi_n(\mathbf{p})-\phi_n(\mathbf{q})|^2}{\rho^2(\lambda_n)}$ becomes stationary; i.e., $|d_{k+1}(\mathbf{p},\mathbf{q}) - d_k(\mathbf{p},\mathbf{q})| < \varepsilon$, where ε is equal to the 1%.

Fig. 3 reports the ℓ_∞ error (y-axis) between the ground-truth distances induced by four filters and their approximation with the truncated spectral method with k Laplacian eigenpairs (x-axis) and our approach. For filters with a fast growth (e.g., $\rho_1 = s^2\exp(st)$, $\rho_2 = s\exp(st)$), the truncated spectral approximation provides a good accuracy (i.e., lower than 10^{-5}, with $k \geq 85$ for the cylinder, and $k \geq 137$ for the sphere). Slowly increasing filters generally require a large number of eigenpairs (i.e., $k \geq 300$ for the cylinder, $k \geq 1K$ for the sphere) to achieve an accuracy lower than 10^{-1}.

Fig. 4 reports the ℓ_∞ discrepancy (y-axis) between the diffusion distance on the sphere/cylinder and its approximation computed with the Padé-Chebyshev method and the truncated spectral approximation. In this case, the analytical expression of the Laplacian eigenfunctions on the sphere and cylinder has been used to compute the ground-truth distances. For small scales (e.g., $t = 10^{-2}, 10^{-3}$), the approximation error remains higher than 10^{-2}, with $k \leq 280$ eigenpairs; in fact, local shape features encoded by the heat kernel are recovered for a small t using the eigenvectors associated with high frequencies, thus requiring the computation of a large part of the Laplacian spectrum. For large scales (e.g., $t = 1, 10^{-1}$), increasing k strongly reduces the approximation error until it becomes almost constant and close to zero. In this case, the behavior of the heat kernel is mainly influenced by the Laplacian eigenvectors related to the smaller eigenvalues. Indeed, the truncated spectral representation generally requires a high number of eigen-

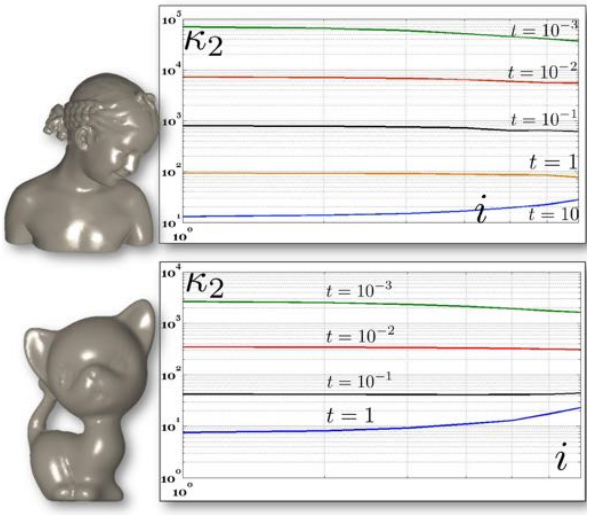

Figure 10: *Conditioning number κ_2 (y-axis) of the matrices $\{(t\mathbf{L}+\theta_i\mathbf{B})\}_{i=1}^{r}$, $r=7$, for the diffusion distance at scale t; the indices of the poles $\{\theta_i\}_{i=1}^{7}$ are reported on the x-axis.*

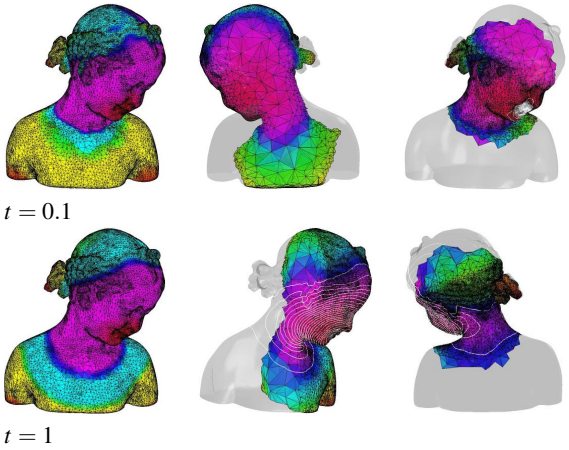

Figure 11: *Volumetric heat kernel ($r=7$). The colors begin with red, pass through yellow, green, cyan, blue, and magenta, and return to red. At scale $t=1$, the level-sets correspond to iso-values uniformly sampled in the range of the solution restricted to the volume boundary.*

pairs and does not achieve the approximation accuracy of our approach, which remains lower than 8.9×10^{-6} for all the scales. According to [VBCG10], there are no theoretical guarantees on the approximation accuracy of the heat kernel provided by multi-resolution prolongation operators. Furthermore, a low-resolution sampling of the input surface might affect the resulting accuracy. Our tests (Table 1, Fig. 5) show that the proposed computation is competitive with respect to the computational cost of prolongation operators and provides a higher approximation accuracy.

We also compare the accuracy of the diffusion distances computed with (i) the proposed approach; (ii) the spectral representation of the heat kernel $\mathbf{K}_t$, with k eigenpairs; (iii) the Euler backward method; and (iv) the power method. We recall that the power method applies the approximation $\mathbf{K}_t \approx (\mathbf{I}-\frac{t}{m}\tilde{\mathbf{L}})^m$, where m is chosen in such a way that t/m is sufficiently small, and the Euler backward method [CDR00,DMSB99] solves the heat equation through the iterative scheme $(t\tilde{\mathbf{L}}+\mathbf{I})\mathbf{F}_{k+1}(t)=\mathbf{F}_k(t)$, $\mathbf{F}_0=\mathbf{f}$. For all the scales (Fig. 6), the accuracy of the Padé-Chebyshev method is higher than the truncated approximation with k eigenpairs, $k=1,\ldots,10^3$, the Euler backward method, and the power method. Reducing the scale, the accuracy of the Padé-Chebyshev remains almost unchanged while the other methods are affected by a larger discrepancy and tend to have an analogous behavior ($t=10^{-4}$). Finally, the Euler backward method generally over-smooths the solution, which converges to a constant as $k\to+\infty$, and the selection of the power m is guided by heuristics.

4.3. Approximation stability

We discuss the stability of our approach with respect to shape changes or to the approximation of the spectral kernel. To this end, we consider the solution $\mathbf{K}_t\mathbf{e}_i$ to the heat diffusion process, whose initial condition takes value 1 at the anchor point $\mathbf{p}_i$ and 0 otherwise. On noisy and irregularly-sampled meshes (Figs. 7, 8) or point sets (Fig. 9), the level-sets of $\mathbf{K}_t\mathbf{e}_i$ are smooth, well-distributed around the anchor point $\mathbf{p}_i$, and remain almost unchanged and coherent with respect to the original shape. These results confirm the robustness of the Padé-Chebyshev approximation to surface discretization. In these examples, the maximum variation between the spectral distances on the complete surface and its representation with holes is lower than 10^{-3}.

According to [MVL03], the approximation of the matrix $\rho(\tilde{\mathbf{L}})$ might be numerically unstable if $\|\tilde{\mathbf{L}}\|_2$ is large. From the bound $\|\mathbf{B}^{-1}\mathbf{L}\|_2 \le \lambda_{\min}^{-1}(\mathbf{B})\lambda_{\max}(\mathbf{L})$, a well-conditioned mass matrix $\mathbf{B}$ guarantees that $\|\mathbf{B}^{-1}\mathbf{L}\|_2$ is bounded. Recalling that $\mathbf{X}^{\top}(\mathbf{B}+\beta_i\mathbf{L})\mathbf{X}=(\mathbf{I}+\beta_i\Lambda)$, $\{1+\beta_i\lambda_j\}_{j=1}^{n}$ are the eigenvalues of $(\mathbf{B}+\beta_i\mathbf{L})$ and its conditioning number is bounded by the constant $(1+\beta_{\max}\lambda_n)$, $\beta_{\max}:=\max_{i=1,\ldots,n}|\beta_i|$. Indeed, the coefficient matrices in Eq. (6) are well-conditioned, as also confirmed by our tests (Fig. 10). In any case, pre-conditioners and regularization techniques [GV89] can be applied to attenuate numerical instabilities. Finally, our computation applies to any discretization of the input domain; in particular, volumes (Fig. 11) that have been considered in few approaches (e.g., [ABBK11, LBB11, LBB12, RBBK10]) due to the high computational cost and low accuracy of the truncated approximation.

The truncated spectra1 approximation of the diffusion dis-

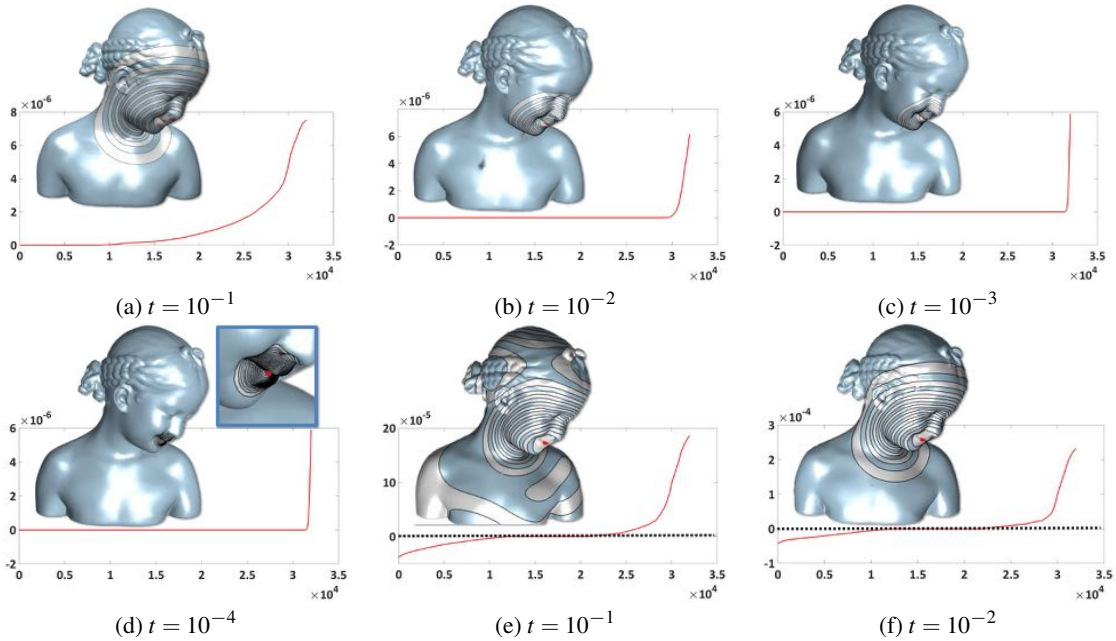

(a) $t = 10^{-1}$ (b) $t = 10^{-2}$ (c) $t = 10^{-3}$

(d) $t = 10^{-4}$ (e) $t = 10^{-1}$ (f) $t = 10^{-2}$

Figure 12: *(a-d) Robustness of the Padé-Chebyshev approximation and (e,f) sensitiveness of truncated spectral approximation to the Gibbs phenomenon. At all scales (a-d), the distance values (red curve) computed with the Padé-Chebyshev approximation are positive; at large scales (e,f), the truncated spectral approximation is already affected by the Gibbs phenomenon, as represented by the part of the plot below the zero line (black curve).*

Table 2: *Timings (in seconds) for the evaluation of the heat kernel on surfaces with n points ($K := 10^3$), approximated with $k = 500$ eigenpairs* (Eigs) *and the Padé-Chebyshev approximation* (Ch., $r = 5$). *Column '×' indicates the number of times the cost is reduced. Tests have been performed on a* 2.7 *GHz Intel Core i7 Processor, with* 8 *GB memory.*

Hand (Fig. 9a)							
n K	Ch.	$k =$ 50	×	$k =$ 100	×	$k =$ 500	×
100	5.0	4.9	0.9	9.2	1.8	49.2	9.7
300	54.9	46.0	0.8	90.4	1.6	423.2	7.6
600	230.5	170.3	0.7	334.1	1.4	1232.0	5.3
800	43.4	312.2	7.1	523.1	12.0	1345.1	30.9
Cat (Fig. 9b)							
200	22.0	21.2	0.9	39.2	1.7	198.1	9
500	159.2	119.2	0.7	231.2	1.4	1087.3	6.8
900	523.4	458.7	0.8	1138.3	2.1	4232.2	8.0
1M	753.2	957.2	1.2	1746.4	2.3	7897.3	10.4
Bimba (Fig. 9c)							
400	98.4	82.4	0.8	156.8	1.5	745.2	7.5
600	212.2	162.3	0.7	323.5	1.4	1189.2	5.3
700	319.3	265.5	0.8	489.6	1.5	1203.5	3.7
900	520.1	438.4	0.8	1076.8	2.1	4156.0	8.0

tance is generally affected by the Gibbs phenomenon; i.e., small negative distance values. This phenomenon is more evident at small cases, which induce diffusion distances that decrease fast to zero and that are largely affected by small negative values. In fact, at small scales the diffusion distances decrease fast to zero and the negative values are no more compensated by the Laplacian eigenvectors related to smaller eigenvalues, as they are not included in the approximation (Fig. 12(e,f)). For the Padé-Chebyshev approximation (Fig. 12(a-d)), the distance values are positive at all the scales; in fact, we approximate the filter map without selecting a sub-part of the Laplacian spectrum.

4.4. Computational cost

Approximating an arbitrary filter map with a rational polynomial of degree r, the evaluation of the corresponding spectral distance between two points is reduced to solve r sparse, symmetric, linear systems (c.f., Eqs. (5), (6)), whose coefficient matrices have the same structure and sparsity of the connectivity matrix of the input triangle mesh or of the k-nearest neighbor graph for a point set. Applying an iterative and sparse linear solver (e.g., Gauss-Seidel method, conjugate gradient) [GV89] (Ch. 10), the computational cost for the evaluation of any Laplacian spectral kernel and distance between two points is $\mathcal{O}(r\tau(n))$, where $\mathcal{O}(\tau(n))$ is the computational cost of the selected solver. Here, the function $\tau(n)$, which depends on the number n of shape samples and the sparsity of the coefficient matrix, varies from $\tau(n) = n$ to $\tau(n) = n \log n$, where $\mathcal{O}(n \log n)$ is the average computational cost of the iterative solvers of sparse linear systems.

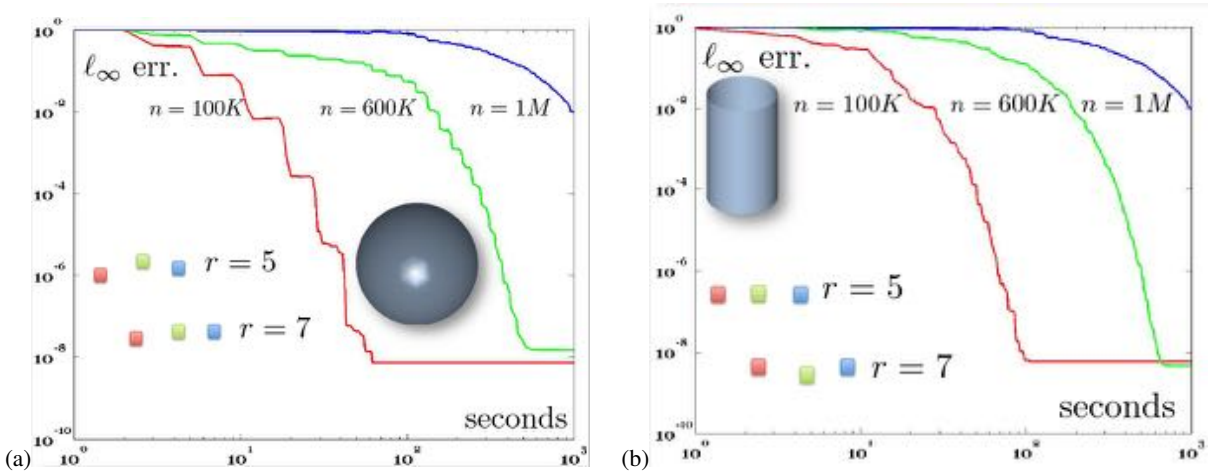

Figure 13: *Trade-off between accuracy (y-axis) and time (x-axis) for the Padé-Chebyshev of degree r ($r = 5, 7$, black boxes) and truncated approximation ($k = 50$ eigenpairs) on the (a) sphere and (b) cylinder. Tests have been performed on a 2.7 GHz Intel Core i7 Processor, with 8 GB memory.*

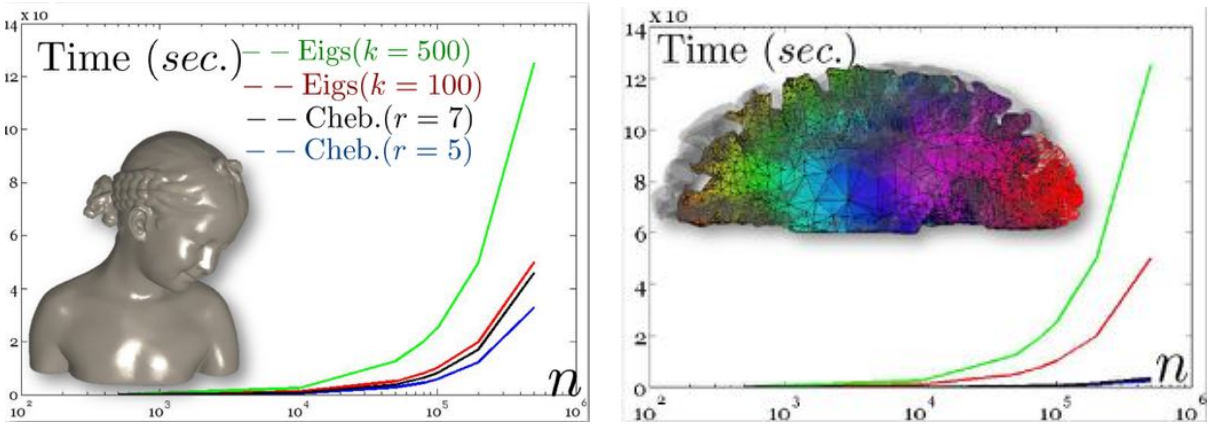

Figure 14: *Timings (in seconds) for the evaluation of the heat kernels on a domain with n points, approximated with $k = 100, 500$ eigenpairs* (Eigs) *and the Padé-Chebyshev approximation ($r = 7$). Tests have been performed on a 2.7 GHz Intel Core i7 Processor, with 8 GB memory.*

Since the truncated spectral approximation and the multi-resolution prolongation operators are independent of the selected filter map, the computed eigenpairs and the multi-resolution structures can be used to evaluate several spectral distances on the same domain. Our computation of the one-to-all distances $\{d(\mathbf{p}_i, \mathbf{p}_j)\}_{j=1}^{n}$ takes $\mathcal{O}(rn\tau(n))$ time; in fact, we solve the linear system (6) with n different right-hand vectors $(\mathbf{e}_i - \mathbf{e}_j)$, $j = 1, \ldots, n$. Computing a fixed number k of eigenpairs in $\mathcal{O}(kn^2)$ time, the one-to-all distance is evaluated in constant time; indeed, the resulting computational cost is competitive with respect to the truncated spectral approximation with $k(n) \geq r\tau(n)/n$ Laplacian eigenpairs. In the average case, $\tau(n) \approx n \log n$ and $k(n) \geq k_n$, $k_n = r \log n$.

For a surface sampled with $n = 10^4, 10^5, 10^6$ points and a degree $r = 5$, the number of eigenpairs is $k_n = 46, 58, 69$; in particular, this growth of k_n with respect to n is slow, as a matter of the logarithm in k_n. Fixing the number of Laplacian eigenpairs makes the truncated spectral approximation of the one-to-all distances faster than ours but generally provides a lower approximation accuracy (Sect. 4.2). Slowly-increasing filters and small scales for the diffusion distances also require the computation of a large number of Laplacian eigenpairs, thus reducing the gap between the computational cost of the proposed approximation of the one-to-all distances and previous work. An analogous discussion applies to prolongation operators, which still compute the truncated spec-

tral approximation on a lower resolution of the input shape. Furthermore, previous work has not addressed methods for the selection of the proper number of eigenpairs with respect to the target approximation accuracy, which cannot be estimated without computing the whole Laplacian spectrum.

Results in Figs. 13, 14 and Table 2 confirm that the diffusion distances at small scales generally require a number of eigenpairs that is much higher than the estimated value k_n. This case makes our computation of the one-to-all distance competitive with respect to its truncated approximation and useful to evaluate the distances among seed points, which is typical in the evaluation of shape descriptors [OFCD02] and bags-of-features [BB11, BBOG11]. Here, the number of seeds is much lower than the number of samples and the higher accuracy of our computation improves the discrimination capabilities of descriptors based on spectral distances.

5. Conclusions and future work

This paper has presented a computation of spectral distances and kernels through the solution of a set of sparse, symmetric, well-conditioned linear systems. Our approach is independent of user-defined or heuristic parameters, the computation of the Laplacian spectrum, multi-resolution prolongation operators, and the discretization of both the Laplace-Beltrami operator and the input domain. As main future work, we foresee a deeper analysis of the constraints that the filter map must satisfy in order to define spectral distances and kernels for shape analysis.

Acknowledgements We acknowledge the precious comments of the anonymous Reviewers, which helped us to improve the paper content and presentation. This work has been partially supported by the *Project TEDIG*, MIUR-SIIT-Regione Liguria, the *FAS Project I-REUMA*, Regione Liguria, and the *FP7 Project IQmulus*. Shapes are courtesy of AIM@SHAPE Repository and tet-meshes were generated by the TETGEN software (http://wias-berlin.de/software/tetgen/).

References

[ABBK11] AFLALO Y., BRONSTEIN A. M., BRONSTEIN M. M., KIMMEL R.: Deformable shape retrieval by learning diffusion kernels. In *Scale space and Variational methods in computer Vision* (2011), pp. 689–700. 2, 7, 9

[ASC11] AUBRY M., SCHLICKEWEI U., CREMERS D.: The wave kernel signature: a quantum mechanical approach to shape analysis. In *IEEE Computer Vision Workshops* (2011), pp. 1626–1633. 1, 2

[AW11] ALEXA M., WARDETZKY M.: Discrete Laplacians on general polygonal meshes. *ACM Trans. on Graphics 30*, 4 (2011). 4

[BB11] BRONSTEIN M., BRONSTEIN A.: Shape recognition with spectral distances. *IEEE Trans. on Pattern Analysis and Machine Intelligence 33*, 5 (2011), 1065 –1071. 1, 2, 12

[BBG94] BERARD P., BESSON G., GALLOT S.: Embedding Riemannian manifolds by their heat kernel. *Geometric and Functional Analysis GAFA 4*, 4 (1994), 373–398. 1

[BBK*10] BRONSTEIN A., BRONSTEIN M., KIMMEL R., MAHMOUDI M., SAPIRO G.: A Gromov-Hausdorff framework with diffusion geometry for topologically-robust non-rigid shape matching. *Intern. Journal of Computer Vision 2-3* (2010), 266–286. 1

[BBOG11] BRONSTEIN A. M., BRONSTEIN M. M., OVSJANIKOV M., GUIBAS L. J.: Shape Google: geometric words and expressions for invariant shape retrieval. *ACM Trans. on Graphics 30*, 1 (2011). 1, 2, 12

[BN03] BELKIN M., NIYOGI P.: Laplacian eigenmaps for dimensionality reduction and data representation. *Neural Computations 15*, 6 (2003), 1373–1396. 1, 4

[BN08] BELKIN M., NIYOGI P.: Towards a theoretical foundation for Laplacian-based manifold methods. *Journal of Computer System Sciences 74*, 8 (2008), 1289–1308. 4

[BSW08] BELKIN M., SUN J., WANG Y.: Discrete Laplace operator on meshed surfaces. In *Proceedings of the Twenty-fourth Annual Symposium on Computational Geometry* (2008), pp. 278–287. 4

[CDR00] CLARENZ U., DIEWALD U., RUMPF M.: Anisotropic geometric diffusion in surface processing. In *IEEE Visualization* (2000), pp. 397–405. 9

[CL06] COIFMAN R. R., LAFON S.: Diffusion maps. *Applied and Computational Harmonic Analysis 21*, 1 (2006), 5–30. 1, 4

[CPS13] CRANE K., PINKALL U., SCHRÖDER P.: Robust fairing via conformal curvature flow. *ACM Transactions on Graphics 32*, 4 (2013), 61. 1

[CRV84] CARPENTER A., RUTTAN A., VARGA R.: Extended numerical computations on the "1/9" conjecture in rational approximation theory. In *Rational Approximation and Interpolation*, vol. 1105 of *Lecture Notes in Mathematics*. Springer, 1984, pp. 383–411. 6

[CWW13] CRANE K., WEISCHEDEL C., WARDETZKY M.: Geodesics in heat: a new approach to computing distance based on heat flow. *ACM Siggraph 32*, 14 (2013). 1, 2, 6

[dGGV08] DE GOES F., GOLDENSTEIN S., VELHO L.: A hierarchical segmentation of articulated bodies. *Computer Graphics Forum 27*, 5 (2008), 1349–1356. 1

[DMSB99] DESBRUN M., MEYER M., SCHRÖDER P., BARR A. H.: Implicit fairing of irregular meshes using diffusion and curvature flow. In *ACM Siggraph* (1999), pp. 317–324. 1, 4, 9

[FPS05] FOUSS F., PIROTTE A., SAERENS M.: A novel way of computing similarities between nodes of a graph, with application to collaborative recommendation. In *IEEE/WIC/ACM Intern. Conference on Web Intelligence* (2005), pp. 550–556. 1, 7

[GBAL09] GEBAL K., BÆRENTZEN J. A., AANÆS H., LARSEN R.: Shape analysis using the auto diffusion function. *Computer Graphics Forum 28*, 5 (2009), 1405–1413. 1

[GK06] GINE E., KOLTCHINSKII V.: Empirical graph Laplacian approximation of Laplace-Beltrami operators: large sample results. *IMS Lecture Notes Monograph Series 51* (2006), 238. 1

[GV89] GOLUB G., VANLOAN G.: *Matrix Computations*. John Hopkins University Press, 2nd Edition, 1989. 2, 5, 6, 9, 10

[Hoe68] HOERMANDER L.: The spectral function of an elliptic operator. *Acta Mathematica 121*, 1 (1968), 193–218. 3

[HQ12] HOU T., QIN H.: Continuous and discrete Mexican hat wavelet transforms on manifolds. *Graphical Models 74*, 4 (2012), 221–232. 2

[HVG11] HAMMOND D. K., VANDERGHEYNST P., GRIBONVAL R.: Wavelets on graphs via spectral graph theory. *Applied and Computational Harmonic Analysis 30*, 2 (2011), 129 – 150. 1, 2, 3, 6, 7

[KR05] KIM B., ROSSIGNAC J.: Geofilter: Geometric selection of mesh filter parameters. *Computer Graphics Forum 24*, 3 (2005), 295–302. 1

[LBB11] LITMAN R., BRONSTEIN A., BRONSTEIN M.: Diffusion-geometric maximally stable component detection in deformable shapes. *Computers & Graphics 35*, 3 (2011), 549–560. 9

[LBB12] LITMAN R., BRONSTEIN A., BRONSTEIN M.: Stable volumetric features in deformable shapes. *Computers & Graphics 36*, 5 (2012), 569 – 576. 9

[LKC06] LAFON S., KELLER Y., COIFMAN R. R.: Data fusion and multicue data matching by diffusion maps. *IEEE Trans. on Pattern Analysis Machine Intelligence 28*, 11 (2006), 1784–1797. 1

[LPG12] LIU Y., PRABHAKARAN B., GUO X.: Point-based manifold harmonics. *IEEE Trans. on Visualization and Computer Graphics 18*, 10 (2012), 1693 –1703. 4

[LRF10] LIPMAN Y., RUSTAMOV R. M., FUNKHOUSER T. A.: Biharmonic distance. *ACM Trans. on Graphics 29*, 3 (2010). 1, 2, 5

[LS96] LEHOUCQ R., SORENSEN D. C.: Deflation techniques for an implicitly re-started Arnoldi iteration. *SIAM Journal of Matrix Analysis and Applications 17* (1996), 789–821. 5, 7

[LSW09] LUO C., SAFA I., WANG Y.: Approximating gradients for meshes and point clouds via diffusion metric. *Computer Graphics Forum 28* (2009), 1497–1508(12). 1

[Mem09] MEMOLI F.: Spectral Gromov-Wasserstein distances for shape matching. In *Workshop on Non-Rigid Shape Analysis and Deformable Image Alignment* (2009), pp. 256–263. 1

[MVL03] MOLER C., VAN LOAN C.: Nineteen dubious ways to compute the exponential of a matrix, twenty-five years later. *SIAM Review 45*, 1 (2003), 3–49. 9

[OFCD02] OSADA R., FUNKHOUSER T., CHAZELLE B., DOBKIN D.: Shape distributions. *ACM Trans. on Graphics 21*, 4 (2002), 807–832. 2, 12

[OMMG10] OVSJANIKOV M., MÉRIGOT Q., MÉMOLI F., GUIBAS L.: One point isometric matching with the heat kernel. *ACM Symposium on Discrete Algorithms* (2010), 650–663. 1

[OSV12] ORECCHIA L., SACHDEVA S., VISHNOI N. K.: Approximating the exponential, the Lanczos method and an $\tilde{o}(m)$-time spectral algorithm for balanced separator. In *Proc. of the 44th Symposium on Theory of Computing Conference* (2012), pp. 1141–1160. 6

[Pat13] PATANÈ G.: wFEM heat kernel: discretization and applications to shape analysis and retrieval. *Computer Aided Geometric Design 30*, 3 (2013), 276–295. 1, 2, 6

[Pat14] PATANÈ G.: Laplacian spectral distances and kernels on 3D shapes. *Pattern Recognition Letters 47* (2014), 102–110. 1

[PP93] PINKALL U., POLTHIER K.: Computing discrete minimal surfaces and their conjugates. *Experimental Mathematics 2*, 1 (1993), 15–36. 4

[PS13] PATANÈ G., SPAGNUOLO M.: Heat diffusion kernel and distance on surface meshes and point sets. *Computers & Graphics 37*, 6 (2013), 676 – 686. 1

[RBBK10] RAVIV D., BRONSTEIN M. M., BRONSTEIN A. M., KIMMEL R.: Volumetric heat kernel signatures. In *Proceedings of the ACM Workshop on 3D Object Retrieval* (2010), 3DOR '10, pp. 39–44. 9

[RBG*09] REUTER M., BIASOTTI S., GIORGI D., PATANÈ G., SPAGNUOLO M.: Discrete Laplace-Beltrami operators for shape analysis and segmentation. *Computer & Graphics 33*, 3 (2009), 381–390. 7

[RS13] RAMANI K., SINHA A.: Multiscale kernels using random walks. *Computer Graphics Forum 33*, 1 (2013), 164–177. 1, 7

[Rus11] RUSTAMOV R. M.: Multiscale biharmonic kernels. *Computer Graphics Forum 30*, 5 (2011), 1521–1531. 2

[RWP06] REUTER M., WOLTER F.-E., PEINECKE N.: Laplace-Beltrami spectra as Shape-DNA of surfaces and solids. *Computer-Aided Design 38*, 4 (2006), 342–366. 4

[Sin06] SINGER A.: From graph to manifold Laplacian: the convergence rate. *Applied and Computational Harmonic Analysis 21*, 1 (2006), 128 – 134. 1

[SK03] SMOLA A. J., KONDOR R. I.: Kernels and regularization on graphs. In *Conference on Learning Theory* (2003), pp. 144–158. 1

[Sog88] SOGGE C. D.: Concerning the lp norm of spectral clusters for second-order elliptic operators on compact manifolds. *Journal of Functional Analysis 77*, 1 (1988), 123 – 138. 3

[SOG09] SUN J., OVSJANIKOV M., GUIBAS L. J.: A concise and provably informative multi-scale signature based on heat diffusion. *Computer Graphics Forum 28*, 5 (2009), 1383–1392. 1

[Sor92] SORENSEN D. C.: Implicit application of polynomial filters in a k-step arnoldi method. *SIAM Journal of Matrix Analysis and Applications 13*, 1 (1992), 357–385. 5, 7

[Tau95] TAUBIN G.: A signal processing approach to fair surface design. In *ACM Siggraph 1995* (1995), pp. 351–358. 1

[TZG96] TAUBIN G., ZHANG T., GOLUB G. H.: Optimal surface smoothing as filter design. In *4th European Conference on Computer Vision* (1996), pp. 283–292. 1

[Var90] VARGA R.: *Scientific computation on mathematical problems and conjectures*. SIAM, CBMS-NSF regional conference series in applied mathematics, 1990. 7

[VBCG10] VAXMAN A., BEN-CHEN M., GOTSMAN C.: A multi-resolution approach to heat kernels on discrete surfaces. *ACM Trans. on Graphics 29*, 4 (2010), 1–10. 1, 9

[VL08] VALLET B., LÉVY B.: Spectral geometry processing with manifold harmonics. *Computer Graphics Forum 27*, 2 (2008), 251–260. 1, 4, 7

[XHW10] XIAO B., HANCOCK E. R., WILSONB R.: Geometric characterization and clustering of graphs using heat kernel embeddings. *Image and Vision Computing 28*, 6 (2010), 1003 – 1021. 1

[ZF03] ZHANG H., FIUME E.: Butterworth filtering and implicit fairing of irregular meshes. In *Proc. of the Pacific Conference on Computer Graphics and Applications* (2003), p. 502. 1