

An Interactive Analysis of Harmonic and Diffusion Equations on Discrete 3D Shapes

Giuseppe Patané^a, Michela Spagnuolo^a,

^aConsiglio Nazionale delle Ricerche
Istituto di Matematica Applicata e Tecnologie Informatiche
Via De Marini, 6, 16149 Genova, Italy
{patane,spagnuolo}@ge.imati.cnr.it

Abstract

Recent results in geometry processing have shown that shape segmentation, comparison, and analysis can be successfully addressed through the spectral properties of the Laplace-Beltrami operator, which is involved in the harmonic equation, the Laplacian eigenproblem, the heat diffusion equation, and the definition of spectral distances, such as the bi-harmonic, commute time, and diffusion distances. In this paper, we study the discretization and the main properties of the solutions to these equations on 3D surfaces and their applications to shape analysis. Among the main factors that influence their computation, as well as the corresponding distances, we focus our attention on the choice of different Laplacian matrices, initial boundary conditions, and input shapes. These degrees of freedom motivate our choice to address this study through the executable paper, which allows the user to perform a large set of experiments and select his/her own parameters. Finally, we represent these distances in a unified way and provide a simple procedure to generate new distances on 3D shapes.

Keywords: Spectral methods, harmonic and bi-harmonic functions, heat diffusion equation, bi-harmonic and diffusion distances.

1. Introduction

Three main classes of functions are associated to the Laplace-Beltrami operator $\Delta : C^2(\mathcal{N}) \rightarrow C^0(\mathcal{N})$, where $C^k(\mathcal{N})$ is the class of maps defined on the manifold \mathcal{N} and with order k of differentiability: the harmonic maps, the Laplacian eigenfunctions, and the solutions to the heat diffusion equation [57]. The spectral properties of this operator have also several applications in geometry processing and shape analysis, such as parameterization, remeshing, segmentation, and comparison. In this context, our work addresses the main discretizations and properties of the harmonic functions and the solutions to the heat diffusion equation. We also take into account the corresponding distances and provide a unified characterization based on the filtered spectral distances, thus providing a simple way to generate new distances on 3D shapes.

For the discretization of the heat diffusion kernel [51], in the space $\mathcal{F}(\mathcal{M})$ of piecewise linear scalar maps defined on a triangle mesh \mathcal{M} we consider the *weighted scalar product* $\langle \mathbf{f}, \mathbf{g} \rangle_B := \mathbf{f}^T B \mathbf{g}$, $\mathbf{f}, \mathbf{g} \in \mathcal{F}(\mathcal{M})$, where $\mathbf{f} := (f(\mathbf{p}_i))_{i=1}^n$ is the array of the f -values at the mesh vertices $\{\mathbf{p}_i\}_{i=1}^n$. This product is induced by the symmetric and positive-definite matrix B associated to the linear FEM discretization $\tilde{L} := B^{-1}L$ of the Laplace-Beltrami operator [55, 74]. Here, the stiffness matrix L is the Laplacian matrix with cotangent weights and the mass matrix B , or its lumped version, encodes the areas of the triangles, or of the Voronoi regions, of \mathcal{M} . The product $\langle \cdot, \cdot \rangle_B$ generalizes the Euclidean product (i.e., $B := I$); is intrinsic to the surface on which the scalar maps are defined; and is adapted to the

local sampling of \mathcal{M} through the variation of the triangle areas. Under this formulation, the generalized eigenproblem for (L, B) is $LX = BX\Gamma$ [25], where $X := [\mathbf{x}_1, \dots, \mathbf{x}_n]$, $\Gamma := \text{diag}(\lambda_i)_{i=1}^n$ are the eigenvalues' and eigenvectors' matrices.

In $\mathcal{F}(\mathcal{M})$, we derive the *wFEM heat kernel* $K_t := XD_t X^T B$ and an *intrinsic, multi-scale* scalar product $\langle \cdot, \cdot \rangle_t$, which makes K_t self-adjoint and induces the corresponding wFEM diffusion distances. If the mass matrix B is lumped to the positive diagonal matrix D , then the wFEM heat kernel $K_t := XD_t X^T B$ becomes equal to the discretization $K_t^* := XD_t X^T D$, which holds for Laplacians of type $L := D^{-1}W$ [9, 49, 60, 71, 75]. Here, W has the mask of the mesh adjacency matrix and the diagonal entries of D are the areas of the Voronoi regions associated to the vertices of \mathcal{M} . If B is equal to the identity matrix, then K_t reduces to the spectral representation of the heat diffusion kernel $\tilde{K}_t := XD_t X^T$ with cotangent weights [10, 14, 23, 76]. Using the mass matrix B instead of its lumped version D , or the identity matrix, allows us to accurately encode the geometry of the input surface through the area of its triangles instead of its Voronoi regions. In this way, the wFEM heat kernel and the corresponding FEM diffusion distances have a higher robustness against topological and scale changes, irregular sampling, and noise. Finally, this choice guarantees that the induced wFEM heat kernel is intrinsically scale-covariant and scale-invariant through a normalization of the Laplacian eigenvalues.

The Laplacian spectrum is also fundamental to define the diffusion [5, 13, 37], commute time, and biharmonic [61] distances. For shape analysis, these distances must be a metric

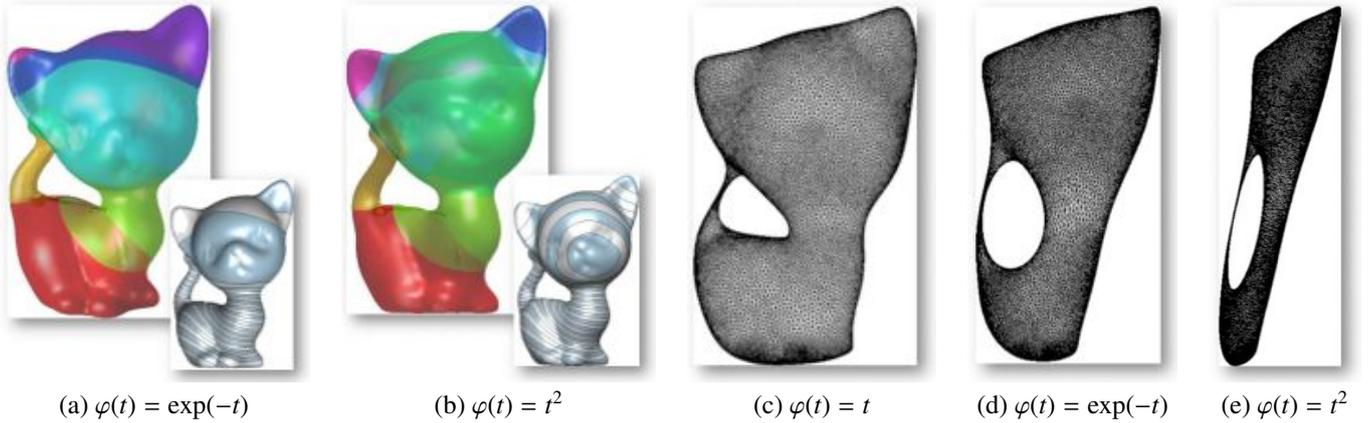


Figure 1: Level sets and surface decomposition into 0-genus patches induced by (a) the heat and (b) biharmonic shape signatures. (c-e) Shape embeddings with respect to different filter maps. In all these examples, we used the Laplacian eigenvectors related to the first $k = 500$ eigenvalues of smaller magnitude.

(i.e., nullity, symmetry, triangular inequality hold) and satisfy additional properties, such as local and global shape-awareness, isometry invariance, robustness to surface sampling, topological and geometric noise. Noting that these distances are related to each others by applying a different weighting to the Laplacian eigenvalues, they are represented in a unified way through the *filtered spectral distance*

$$d(\mathbf{p}_i, \mathbf{p}_j) := \left[\sum_{k=1}^n \frac{|x_k(\mathbf{p}_i) - x_k(\mathbf{p}_j)|^2}{\varphi(\lambda_k)} \right]^{1/2}, \quad \mathbf{p}_i, \mathbf{p}_j \in \mathcal{M}.$$

Here, $x_k(\mathbf{p}_i)$ is the i th component of the vector \mathbf{x}_k and $\varphi : \mathbb{R} \rightarrow \mathbb{R}$ is a strictly positive map, which is applied to the Laplacian eigenvalues. Specific filters provide well-known or new spectral distances on \mathcal{M} ; for instance, the linear $\varphi(s) := s$, quadratic $\varphi(s) := s^2$, or exponential $\varphi_i(s) := \exp(st)$ maps induce the commute time [53], bi-harmonic [61], or diffusion distances [9], respectively. In [6], the filter map $\varphi(s) := \exp(-is)$, $s \in [0, 2\pi]$, defines the wave kernel signature $\mathcal{S}(\mathbf{p}_i) := \sum_{j=1}^n \varphi(\lambda_j) x_j^2(\mathbf{p}_i)$ and the filter values $\{\varphi(\lambda_i)\}_{i=1}^n$ are the frequency responses induced by the Schroedinger equation. In this case, the corresponding solution $F(\cdot, t) := \sum_{j=1}^n \varphi(\lambda_j t) \langle f, \mathbf{x}_j \rangle_2 \mathbf{x}_j$ is associated to a complex filter, whose values belong to the interval $[-1, 1]$. In our discussion, we focus on the spectral distances, together with its relation with the generalized Green kernel, instead of kernel signatures; consider only positive maps that converge to zero, as $s \rightarrow +\infty$; and derive the discretization induced by a scalar product intrinsic to the input shape instead of the Euclidean product.

Starting from these filters, we can design new distances; the only constraint is that φ is a strictly positive map. The capability of locally or globally characterizing the input shape depends on the convergence of the filtered eigenvalues $(\varphi(\lambda_i))_{i=1}^n$ to zero. Furthermore, the condition $\lim_{s \rightarrow +\infty} \varphi(s) = 0$ guarantees a smooth spectral distance and a good approximation accuracy through the truncated approximation (c.f., Eq. (13)). Increasing the filter decay to zero, the effects of the eigenvalues of largest magnitude and of the corresponding eigenvectors on the filtered spectral distance are negligible with respect to the contribution of the lower eigenvalues. The resulting distance characterizes the global properties of the input shape, while

poorly identifying its local properties. Reducing the filter decay to zero, local shape features are better characterized. A trade-off between the measure of both local and global properties can be achieved by selecting as φ a convex combination between the filters associated to the diffusion and bi-harmonic distances, respectively.

As main applications of the filtered spectral distances, we focus on shape segmentation into 0-genus patches with respect to generalized auto diffusion functions and shape embeddings. For instance, selecting the filter φ we compute a family of segmentations into 0-genus patches (Fig. 1 (a,b)) and embed 3D shapes in abstract representations (Fig. 1 (c-d)). For a more detailed description of the main applications of the filtered spectral distances and related shape descriptors, we refer the reader to Sect. 2.

Through the executable paper, the reader will be able to compute the harmonic functions and the solutions to the heat diffusion equation, by selecting (i) a specific set of weights for the Laplacian matrix (e.g., cot, Voronoi-cot, linear FEM weights); (ii) the input shapes, which are represented as triangle meshes; and (iii) the initial boundary conditions. Among the main properties, we discuss their robustness to shape discretization and noise, smoothness, scale-invariance, dependence on the discretization of the Laplace-Beltrami operator, and computational stability. The resulting scalar functions and distances are visualized using a color coding of the function values, the evolution of the corresponding level sets and critical points. Finally, the filtered spectral distances are computed through the selection of the spectrum of different Laplacian matrices, filter maps, and source points. For our tests, we have selected 3D shapes that are (i) characterized by a different sampling density and geometric/topological complexity; (ii) modified through isometric and non-rigid almost inelastic deformations, global and local scaling; and (iii) affected by noise of different magnitude.

The paper is organized as follows. After a brief overview on previous work (Sect. 2), we introduce the discretization and properties of the harmonic map, the Laplacian eigenfunctions (Sect. 3), and the solutions to the heat diffusion equation (Sect. 4).

Then (Sect. 5), we discuss the FEM heat diffusion distances and propose a unified view on the definition of distances on 3D shapes induced by the Laplacian spectrum.

2. Previous work

We briefly review previous work on the harmonic maps, the Laplacian eigenfunctions, the heat diffusion kernel, and the corresponding distances.

Harmonic maps and Laplacian eigenfunctions. Mesh Laplacian operators were initially developed for shape filtering [73]; spectral partitioning [33] and analysis [81]. Successively, Laplacian operators have been associated to a set of differential coordinates for surface deformation [66, 69] and quadrangulation with harmonic maps and Laplacian eigenfunctions [4, 19, 18, 38, 46]. The eigenvector corresponding to the smallest non-zero eigenvalue of the Laplacian graph (i.e., the *Fiedler vector*) has been applied to graph partitioning [1, 21, 43]; to sparse matrices for reducing their bandwidth [3]; to clustering [62] (§ 14) and dimensionality reduction [4]; to graph/mesh layout [17, 36] and image segmentation [63].

In digital geometry processing, low-pass filters [72] have been designed using the spectral properties of the uniform discrete Laplacian and encoding the local geometry of the input surface [15, 34, 52]. Further applications include implicit mesh fairing [15, 34, 79], mesh watermarking [48, 47], and geometry compression [33, 68]. Moreover, the spectral properties of the Laplacian matrix are at the basis of the definition of fairing functionals [35, 40], which optimize the triangles' shape and/or the surface smoothness [44].

The Laplacian eigenvectors were also used for embedding a surface of arbitrary genus into the plane [82, 84] and mapping a closed 0-genus surface into a spherical domain [26]. In the frequency space, shape segmentation [39, 80], correspondence [30, 31], and comparison [30, 55] have been successfully addressed. Recent applications of the Laplacian spectrum include shape segmentation and analysis through nodal domains [54] and constrained smoothing of scalar functions [50]. Finally, theoretical results on the sensitivity of the Laplacian spectrum against geometry changes, irregular sampling density, and connectivity variations were discussed in [20, 32, 78].

Heat diffusion maps and diffusion distances. In geometry processing and shape analysis, several problems have been addressed through the properties of the heat diffusion kernel on a 3D shape. The discrete heat diffusion kernel has been applied to shape segmentation [14] and comparison [9, 10, 16, 23, 49] through heat kernel shape descriptors and auto-diffusion maps [23]; dimensionality reduction [4, 77] with spectral embeddings; the computation of the gradient of discrete maps [76] and the multi-scale approximation of functions [51]. In [75], prolongation operators extend the values of the heat diffusion kernel computed on a low resolution representation of the input shape to higher resolutions, through the hierarchy associated to multiresolutive simplification algorithms.

Recently, the heat equation and the associated diffusion metric have been used to define multi-scale shape signatures [71], compare 3D shapes [41], and approximate the gradients of scalar functions defined on triangulated surfaces and point sets [76]. The heat diffusion kernel on manifolds also plays a central role in Machine Learning [13, 24, 27, 28, 37, 62, 64]; dimensionality reduction with spectral embeddings [4, 77]; data visualization [4, 28, 58], representation [11, 65, 83], and classification [45, 63, 70].

3. Laplacian matrix and harmonic maps

We briefly recall the discretization of the Laplace-Beltrami operator, the classification of the critical points for piecewise scalar functions defined on triangulated surfaces, and the main properties of the harmonic functions.

Laplacian matrix. Let us consider a triangle mesh $\mathcal{M} := (\mathcal{P}, T)$, which discretizes a manifold \mathcal{N} , where $\mathcal{P} := \{\mathbf{p}_i\}_{i=1}^n$ is the set of n vertices and T is the connectivity graph. The piecewise linear scalar function $f : \mathcal{M} \rightarrow \mathbb{R}$ is defined by linearly interpolating the values $\mathbf{f} := (f(\mathbf{p}_i))_{i=1}^n$ of f at the vertices using barycentric coordinates. Then, the Laplacian matrix is defined as $\tilde{L} := B^{-1}L$, where

$$L(i, j) := \begin{cases} w(i, j) := \frac{\cot \alpha_{ij} + \cot \beta_{ij}}{2} & j \in N(i), \\ -\sum_{k \in N(i)} w(i, k) & i = j, \\ 0 & \text{else.} \end{cases}$$

is the Laplacian matrix with cotangent weights; here, $N(i)$ is the set of the indices of the vertices incident to \mathbf{p}_i and α_{ij}, β_{ij} are the angles that are opposite to the edge (i, j) . The weight matrix B is the diagonal matrix whose entries are the areas of the Voronoi regions of the mesh vertices (*Voronoi-cot weights*) [15]

$$B := \text{diag}(a(i))_{i=1}^n, \quad a(i) := \frac{1}{3} \sum_{t \in N(i)} \text{area}(t).$$

Alternatively, B is the FEM mass matrix [56, 74] (*linear FEM weights*), which codes the variation of the triangle areas

$$B(i, j) := \begin{cases} \frac{|t_a| + |t_b|}{12} & j \in N(i), \\ \frac{\sum_{k \in T(i)} |t_k|}{6} & i = j, \\ 0 & \text{else,} \end{cases}$$

where $\text{area}(t)$ is the area of the triangle t , $T(i)$ is the set of triangles that have $\{i\}$ among their vertices, and t_a, t_b are the triangles that share the edge (i, j) . If the FEM mass matrix is lumped to the diagonal matrix, then the FEM Laplacian matrix becomes equal to the Voronoi-cot Laplacian matrix. If $B := I$, then we get the Laplacian matrix with cot weights [52].

If the input surface is bordered (e.g., Fig. 13), then we symmetrize \tilde{L} as $(\tilde{L} + \tilde{L}^T)/2$; since our approach works mainly on matrices, this symmetrization is enough to guarantee that the computation of the filtered spectral distances is not affected by surface holes and local shape artifacts.

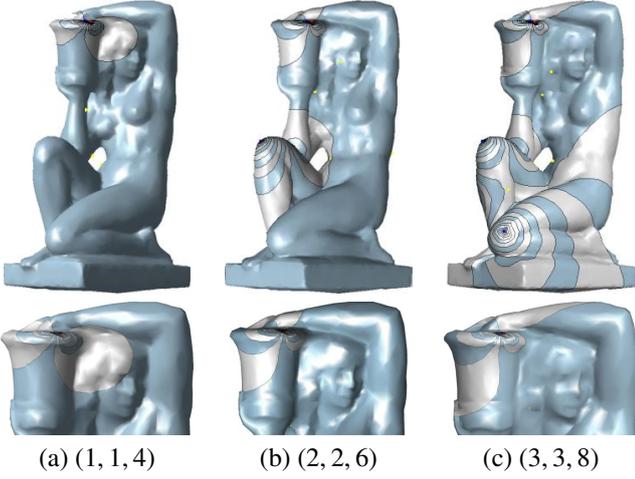


Figure 2: Level sets and critical points (m, M, s) of harmonic maps with (a) two, (b) four, and (c) six Dirichlet boundary conditions, which are shown in blue and red. Blue, red, and green points indicate the m minima, M maxima, and s saddles of each map, respectively. The insertion of new boundary constraints locally affects the resulting harmonic map. See also Fig. 3.

Critical points. In the discrete case, the behavior of a piecewise linear scalar function $f : \mathcal{M} \rightarrow \mathbb{R}$ is conveyed by the corresponding level sets $f^{-1}(\alpha)$ and the critical points of f , at which the number of connected components of the level sets changes. Let us assume that $f : \mathcal{M} \rightarrow \mathbb{R}$ is *general*; i.e., $f(\mathbf{p}_i) \neq f(\mathbf{p}_j)$, for each edge (i, j) . Then, the critical points of $f : \mathcal{M} \rightarrow \mathbb{R}$ are computed by analyzing the distribution of the f -values on the neighborhood of each vertex [2]. Under the assumption that f is a regular map, the *Euler formula*

$$\chi(\mathcal{M}) = m - s + M, \quad g = \frac{1}{2}(2 - \chi(\mathcal{M})), \quad (1)$$

gives the link between the critical points of (\mathcal{M}, f) , the Euler characteristic $\chi(\mathcal{M})$ of \mathcal{M} [2, 42], and the genus g of \mathcal{M} . Here, m and M are the number of minima and maxima, respectively; the $s := \sum_{\mathbf{p}_i \text{ saddle}} m_i$ saddle points of f are counted with their multiplicity m_i .

Harmonic maps. The *harmonic function* $h : \mathcal{N} \rightarrow \mathbb{R}$ is the solution of the Laplace equation $\Delta h = 0$ with Dirichlet boundary conditions $h|_S = h_0$, $S \subset \mathcal{N}$, where h_0 is the initial condition. The *locality property* guarantees that if \mathbf{p} and \mathbf{q} are two distinct points, then $\Delta h(\mathbf{p})$ is not affected by the value of h at \mathbf{q} . According to the *maximum principle* [57], a harmonic map has no local extrema other than at constrained vertices. In the case that all constrained minima are assigned the same global minimum value and all constrained maxima are assigned the same global maximum value, all the constraints will be extrema in the resulting field.

The harmonic equation is approximated at the vertices of \mathcal{M} as the homogeneous linear system $L\mathbf{f} = \mathbf{0}$, with Dirichlet boundary conditions $\mathcal{B} := \{f(\mathbf{p}_i) = a_i\}_{i \in \mathcal{I}}$, $\mathcal{I} \subseteq \{1, \dots, n\}$. According to the Euler formula (1), the number of critical points of a harmonic map depends on the number of Dirichlet boundary conditions, which determine the maxima and minima of the resulting harmonic map. In particular, a harmonic function with one

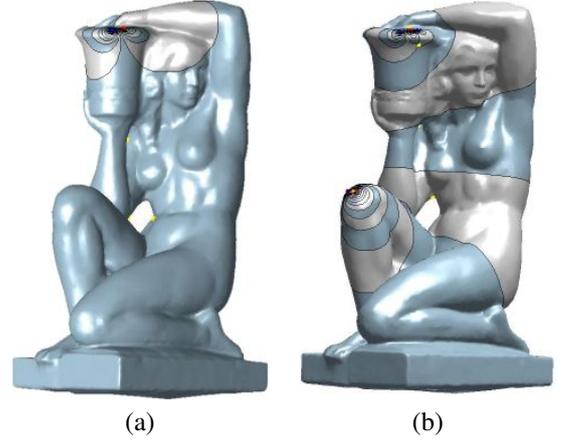


Figure 3: With reference to Fig. 2(a,b), level sets and critical points of harmonic maps with (a) two (b) four Dirichlet boundary conditions on a high-resolution (26K vertices) representation of the input model (5K vertices) in Fig. 2.

maximum and one minimum has a minimal number of $2g$ saddle points, which are located on the topological handles of \mathcal{M} .

Fig. 2 shows a family of harmonic functions achieved by increasing the number of Dirichlet boundary conditions. In (a), we have selected two boundary conditions, which identify the maximum and minimum of f ; in (b), we have added two new boundary conditions, while maintaining the previous ones. According to the locality property, the resulting harmonic function and its level sets remain unchanged in a neighborhood of the Dirichlet points related to the previous step. An analogous remark applies to (b,c), where in (c) we added two constraints to the set of Dirichlet conditions used in (b). While the position of the extrema is determined *a-priori* by the Dirichlet boundary conditions, only the number of saddles and not their locations on the input surface are determined through the Euler formula. Fig. 3 shows an analogous example on a high resolution model.

Executable paper.¹ For the experiments presented in the paper, the user can select the input surface among 30 triangle meshes, which have been stored in *off* format as *shape-i.off*, $i = 1, \dots, 30$. For the tests in Figs. 2, 3, the reader can run the experiment *Harmonic_function.exp.xml*, which asks to select the *off* file of the input shape. The default file is *3torus.off*; the acquarius models at low (Fig. 2) and high resolution (Fig. 3) are stored in *acquarius-5Kv.off* and *acquarius-26Kv.off*, respectively. The routine computes three harmonic functions with two, four, and six Dirichlet boundary conditions and displays the corresponding M maxima, m minima, and s saddles. The first two initial conditions are placed at the vertices with minimum and maximum x -coordinates; the corresponding f -values are 0 and 1. For four and six initial conditions, we also consider the vertices of minimum and maximum y - and z -coordinates. The input shape, the colormap, and the level sets of each harmonic function are stored in *harmonic-map-i.wrl* and displayed

¹<https://collage.elsevier.com/collage/template?doi=10.0000/1352475983201>

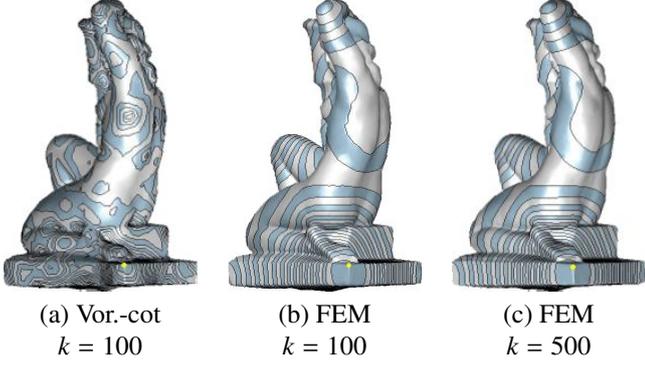


Figure 4: Level sets of the basis function $K_i e_j$ centered at the yellow point; here, we have selected the (a) Voronoi-cot and (b,c) the FEM weights, with the same time parameter t and number k of Laplacian eigenpairs.

in *harmonic-map-i.png*, $i = 1, 2, 3$. The colormap varies the hue component of the hue-saturation-value color model; the colors begin with red, pass through yellow, green, cyan, blue, magenta, and return to red. Finally, the array of the values of each harmonic map at the mesh vertices is stored in *harmonic-field-i.txt*, $i = 1, 2, 3$. For the visualization of the *wrl* file, the user can use *View3DScene*² or analogous visualization tools.

4. wFEM Heat kernel

We briefly address the discretization of the heat diffusion kernel with respect to linear finite elements [51] (Sect. 4.1), its properties (Sect. 4.2), and its computation (Sect. 4.3). Then (Sect. 4.4), we discuss the main features of this discretization with respect to previous work. Among them, we focus on the multi-scale property; the stability under shape perturbations (e.g., sampling, connectivity, topological noise); and the invariance to isometries and inelastic deformations.

4.1. Linear FEM discretization

The *scale-based representation* $H : \mathcal{N} \times \mathbb{R}^+ \rightarrow \mathbb{R}$ of the map $h : \mathcal{N} \rightarrow \mathbb{R}$ is the solution to the *heat diffusion equation*

$$\partial_t H(\mathbf{p}, t) = -\Delta H(\mathbf{p}, t), \quad H(\mathbf{p}, 0) = h(\mathbf{p}), \quad (2)$$

$(\mathbf{p}, t) \in \mathcal{N} \times \mathbb{R}^+$ and it can be written through the convolution operator \star as $H(\mathbf{p}, t) := k_t(\mathbf{p}, \cdot) \star h = \int_{\mathcal{N}} k_t(\mathbf{p}, \mathbf{q}) h(\mathbf{q}) d\mathbf{q}$, where $k_t(\mathbf{p}, \mathbf{q}) := \sum_{i=1}^{+\infty} \exp(-\lambda_i t) \varphi_i(\mathbf{p}) \varphi_i(\mathbf{q})$ is the *heat diffusion kernel*.

To derive the discretization to Eq. (2), we replace the space $\mathcal{L}_2(\mathcal{N})$ of square integrable functions defined on \mathcal{N} with the finite dimensional subspace

$$\mathcal{F}(\mathcal{M}) := \{f : \mathcal{M} \rightarrow \mathbb{R}, f \text{ PL}, \mathbf{f} := (f(\mathbf{p}_i))_{i=1}^n \in \mathbb{R}^n\},$$

of piecewise linear (PL, for short) scalar functions defined on the triangle mesh \mathcal{M} . In $\mathcal{F}(\mathcal{M})$, the piecewise linear scalar function $F(\cdot, t) : \mathcal{M} \rightarrow \mathbb{R}$ is uniquely identified by the F -values $\mathbf{F}(t) := (F(\mathbf{p}_i, t))_{i=1}^n$ at the vertices of the input triangle mesh. In

$\mathcal{F}(\mathcal{M})$, let ψ_i be the map that takes value 1 at \mathbf{p}_i and 0 at the other vertices. To re-write the solution $\mathbf{F}(t) = \sum_{i=1}^n F(\mathbf{p}_i, t) \psi_i$ to the heat diffusion equation in terms of the basis of $\mathcal{F}(\mathcal{M})$, let us multiply both sides of the equation $\partial_t F(\cdot, t) + \Delta F(\cdot, t) = 0$ with the test functions in $\mathcal{F}(\mathcal{M})$; integrating this equation over \mathcal{M} , we get that

$$\begin{aligned} 0 &= \langle \partial_t F(\cdot, t) + \Delta F(\cdot, t), \psi_j \rangle_2 \\ &= \sum_{i=1}^k \partial_t F(\mathbf{p}_i, t) \langle \psi_i, \psi_j \rangle_2 + \sum_{i=1}^k F(\mathbf{p}_i, t) \langle \Delta \psi_i, \psi_j \rangle_2 \\ &= \sum_{i=1}^k \partial_t F(\mathbf{p}_i, t) B(i, j) + \sum_{i=1}^k F(\mathbf{p}_i, t) L(i, j). \end{aligned} \quad (3)$$

for $j = 1, \dots, n$. In matrix form, Eq. (3) is rewritten as

$$B \partial_t \mathbf{F}(t) + L \mathbf{F}(t) = \mathbf{0}, \quad \mathbf{F}(0) = \mathbf{f}.$$

Let us now consider the generalized eigensystem

$$\{(\lambda_i, \mathbf{x}_i)\}_{i=1}^n, \quad L \mathbf{x}_i = \lambda_i B \mathbf{x}_i, \quad \mathbf{x}_i^T B \mathbf{x}_j = \delta_{ij},$$

of the couple (L, B) , which defines the linear FEM discretization $\tilde{L} := B^{-1}L$ of the Laplace-Beltrami operator (Sect. 3), or in matrix form

$$LX = BX\Gamma, \quad \Gamma := \text{diag}(\lambda_i)_{i=1}^n, \quad X^T BX = I, \quad (4)$$

where Γ is the diagonal matrix of the Laplacian eigenvalues and $X := [\mathbf{x}_1, \dots, \mathbf{x}_n]$ is the matrix of the eigenvectors. Since the unknown function $F : \mathcal{M} \times \mathbb{R} \rightarrow \mathbb{R}$ is a $n \times 1$ vector for each value of the parameter t , the solution $\mathbf{F}(t) := \sum_{i=1}^n \alpha_i(t) \mathbf{x}_i$ to the discrete heat diffusion equation is expressed as a linear combination of the eigensystem of (L, B) , where $\alpha(t) := (\alpha_i(t))_{i=1}^n$ is the unknown vector. Since B is positive-definite, the space $\mathcal{F}(\mathcal{M})$ is endowed with the *weighted inner product*

$$\langle \mathbf{f}, \mathbf{g} \rangle_B = \mathbf{f}^T B \mathbf{g}, \quad \mathbf{f} := (f(\mathbf{p}_i))_{i=1}^n, \quad \mathbf{g} := (g(\mathbf{p}_i))_{i=1}^n, \quad (5)$$

which is intrinsic to the geometry underlying \mathcal{M} and adapted to its discretization. Using the invertibility of the matrix B , the linear independence of the Laplacian eigenfunctions, and the identity $\mathbf{f} = \sum_{i=1}^n \langle \mathbf{f}, \mathbf{x}_i \rangle_B \mathbf{x}_i$, each component $\alpha_i(t)$ satisfies the differential equation

$$\alpha_i'(t) + \lambda_i \alpha_i(t) = 0, \quad \alpha_i(0) = \langle \mathbf{f}, \mathbf{x}_i \rangle_B, \quad i = 1, \dots, n.$$

Indeed, the *scale-based representation* of $f : \mathcal{M} \rightarrow \mathbb{R}$ is

$$F(\cdot, t) = \sum_{i=1}^n \exp(-\lambda_i t) \langle \mathbf{f}, \mathbf{x}_i \rangle_B \mathbf{x}_i, \quad t \in \mathbb{R}, \quad (6)$$

which is re-written in matrix form as $F(\cdot, t) = K_t \mathbf{f}$, where

$$K_t := X D_t X^T B, \quad D_t := \text{diag}(\exp(-\lambda_i t))_{i=1}^n, \quad (7)$$

is the *weighted linear FEM (wFEM) heat kernel matrix*.

²<http://castle-engine.sourceforge.net/view3dscene.php>

Table 1: Definition and properties of different discretizations of the heat kernel: sparsity, positive definiteness, and symmetry. The full \bullet and empty \circ circle means that the corresponding property is or is not satisfied, respectively.

Method	Matrix K_t	Sp.	Pos. Def.	Sym.
Std. HK	XD_tX^T	\circ	\bullet	\bullet
Vor.-cot HK	XD_tX^TD	\circ	\bullet	\circ
wFEM HK	XD_tX^TB	\circ	\bullet	\circ

4.2. Properties of the wFEM heat kernel and comparison with previous work

Lumping the mass matrix B , we get the diagonal matrix $D := \text{diag}(d(i))_{i=1}^n$, whose entries $d(i) := 1/3 \sum_{t \in N(i)} |t|$ are the areas of the Voronoi regions of \mathcal{M} and $|t|$ is the area of the triangle t . Then, K_t in Eq. (7) becomes equal to the *Voronoi-cot* (or *lumped FEM*) heat kernel [9, 49, 60, 71, 75]

$$K_t^* := XD_tX^TD, \quad LX = DXT,$$

Choosing $B := I$ in Eq. (7), we get the *linear FEM heat kernel matrix* $\tilde{K}_t := XD_tX^T$. Using the mass matrix B instead of its lumped version D allows us to accurately encode the geometry of the input surface through the area of its triangles instead of its Voronoi regions.

Comparing the kernels \tilde{K}_t and K_t , we notice that $K_t = \tilde{K}_t B$ and that the wFEM discretization explicitly involves the mass matrix B . In this case, \tilde{K}_t is scale-dependent; i.e., rescaling \mathcal{M} to $\alpha\mathcal{M}$, $\alpha > 0$, the corresponding kernels satisfy the relation $\tilde{K}_t(\alpha\mathcal{M}) = \alpha^{-2}\tilde{K}_{\alpha^{-2}t}(\mathcal{M})$. It follows that the geometric and temporal components of $\tilde{K}_t(\alpha\mathcal{M})$ are affected by the surface rescaling and the corresponding heat kernel shape descriptors cannot be directly used for matching. On the contrary, K_t and K_t^* are intrinsically scale-covariant (i.e., without shape or kernel normalization) and scale-invariant through a normalization of the Laplacian eigenvalues. In [7, 8], these properties and the robustness of the wFEM heat kernel against shape transformation, sampling, and noise have been verified by testing the matching performances of the wFEM heat kernel descriptors on the SHREC'10 data set. Finally, Table 1 summarizes the properties of different discretizations of the heat kernel.

4.3. Computation of the wFEM heat kernel

Once the Laplacian eigensystem has been computed in super-linear time [74], $F(\cdot, t)$ is evaluated in $O(n)$ -time. However, using the whole spectrum is computationally unfeasible in terms of computational cost and storage of full $n \times n$ matrices. Since the exponential decay of the *filter factor* $\sigma_i := \exp(-\lambda_i t)$ increases with λ_i , we consider only a part of the Laplacian spectrum. To this end, the sum in (6) is truncated by considering only the contribution related to the first k eigenpairs; i.e.,

$$F_k(\cdot, t) = \sum_{i=1}^k \exp(-\lambda_i t) \langle \mathbf{f}, \mathbf{x}_i \rangle_B \mathbf{x}_i = X_k D_t^{(k)} X_k^T B \mathbf{f}, \quad t \in \mathbb{R}, \quad (8)$$

where $X_k := [\mathbf{x}_1, \dots, \mathbf{x}_k]$ is the $n \times k$ matrix of Laplacian eigenfunctions and $D_t^{(k)} := \text{diag}(\exp(-\lambda_i t))_{i=1}^k$ is the diagonal matrix with the filter factors. If $t := 0$, then $F_k(\cdot, 0) = \sum_{i=1}^k \langle \mathbf{f}, \mathbf{x}_i \rangle_B \mathbf{x}_i$ is

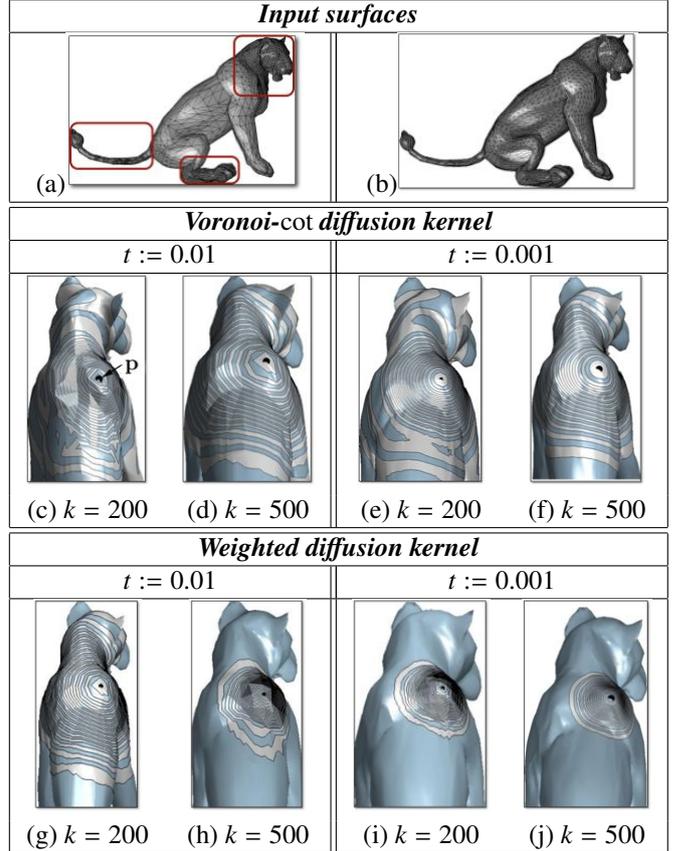


Figure 5: Approximation $F_k(\cdot, t)$ (c.f., Eq. (8)) of the basis function $K_t \mathbf{e}_i$ generated using k Laplacian eigenfunction, the (c-f) Voronoi-cot and (g-j) FEM weights on a 3D shape with a (a) coarse and (b) fine sampling.

the least-squares approximation of f in the linear space generated by the first k eigenfunctions and with respect to $\|\cdot\|_B$.

We notice that both the parameters k and t define the hierarchy of approximations. In fact, reducing the number of Laplacian eigenfunctions results in a smoothing of the input map and a further simplification of its critical points, with more emphasis on those with a low persistence value. Furthermore, the eigenpairs related to the low frequencies code the global structure of the input shape [54, 55]; on the contrary, the highest frequencies code noise and local details. Finally, the choice of the number k of Laplacian eigenpairs is up-to the user and its effects on the final approximation accuracy cannot be *a priori* estimated.

4.4. Examples and discussions

Exploiting the isomorphism between $\mathcal{F}(\mathcal{M})$ and \mathbb{R}^n , we interpret the canonical basis $\mathcal{E} := \{\mathbf{e}_i\}_{i=1}^n$ of \mathbb{R}^n as the set of trivial maps that take value one at a given vertex of \mathcal{M} and zero otherwise. Then, $\mathcal{B} := \{K_t \mathbf{e}_i\}_{i=1}^n$ is a counterpart of \mathcal{E} in $\mathcal{F}(\mathcal{M})$, where each piecewise linear map $K_t \mathbf{e}_i$ is achieved by applying the diffusion process to the function that takes value one at the *anchor* \mathbf{p}_i and zero otherwise. For shape analysis and comparison, as anchors of a 3D shape we can select the maxima and minima of the Laplacian eigenfunctions related to the smallest eigenvalues [59] or of the auto-diffusion maps [23], without/with simplification based on persistence homology [85]. The map

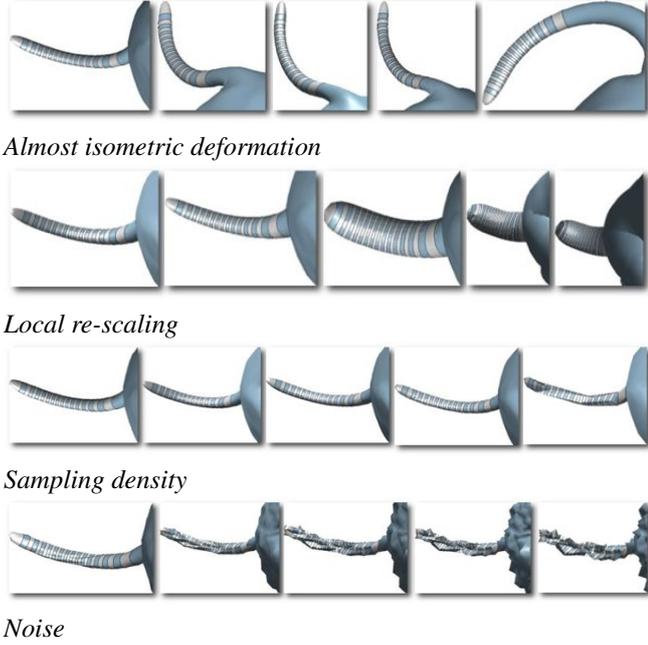


Figure 6: Robustness of the computation of the basis function induced by the wFEM heat kernel and centered at the spike of the tail. The transformation strength increases from left to right.

$K_t \mathbf{e}_i$ summarizes the shape distribution in a neighbor of \mathbf{p}_i , in a way similar to the local point signatures [12] and geodesic-based descriptors [22, 29]. This basis is a multi-scale version of the geometry-aware functions [67]; however, the definition and properties of \mathcal{B} are driven by the geometry of \mathcal{M} and not uniquely by its connectivity. It is also intrinsic to the input shape, multi-scale, and localized, according to the values of the time parameter.

Increasing t , the map $K_t \mathbf{e}_i$ changes from the Dirichlet function ($t := 0$) to a constant map. In fact, recalling that the first eigenpair is $(0, \mathbf{1})$ and using the relation

$$K_t \mathbf{e}_i = \langle \mathbf{1}, \mathbf{e}_i \rangle_B \mathbf{1} + \sum_{i=2}^n \exp(-\lambda_i t) \langle \mathbf{e}_i, \mathbf{x}_i \rangle_B \mathbf{x}_i \rightarrow_{t \rightarrow +\infty} (\mathbf{e}_i^T B \mathbf{1}) \mathbf{1},$$

we get that, as $t \rightarrow +\infty$, $K_t \mathbf{e}_i$ converges to the constant function $K_{+\infty} \mathbf{e}_i = \sum_{j=1}^n B(i, j)$. To better characterize these asymptotic properties and analyze different discretizations of the heat diffusion kernel, we compare the basis function

$$K_t^D \mathbf{e}_i := X D_t X^T D \mathbf{e}_i, \quad K_t \mathbf{e}_i := X D_t X^T B \mathbf{e}_i,$$

which are induced by the Voronoi-cot and wFEM heat kernel, respectively.

The linear FEM Laplacian eigensystem (Fig. 4(b,c)), together with the induced B -scalar product (c.f., Eq. (5)) in $\mathcal{F}(\mathcal{M})$, provides smooth level sets that are well-distributed around the anchor point \mathbf{p}_i . On the contrary (Fig. 4(a)), irregularly-sampled data badly affect the smoothness of the standard discretization of the diffusion kernel, thus showing the higher robustness and stability of the proposed discretization with respect to the data sampling.

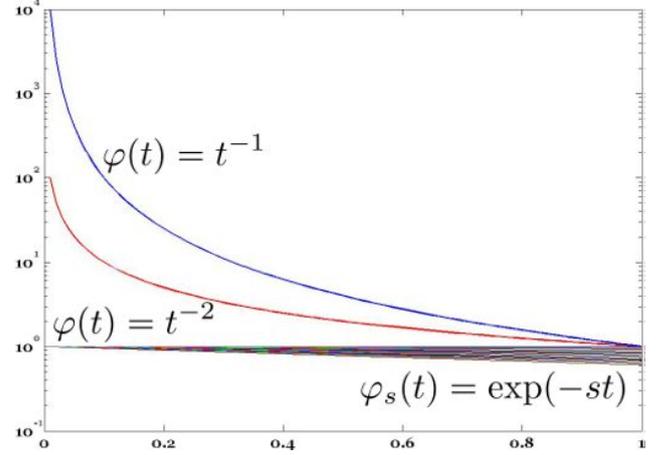


Figure 7: Filter functions that induce the diffusion, commute time, and bi-harmonic distances; on the y-axis, we have reported the log-scale.

The behavior of the level sets (Fig. 5(g-j)) and their uniform distribution around the anchor point confirm that the wFEM heat kernel is not affected by a different sampling, or noise, of the input surface. As a matter of the uneven sampling of \mathcal{M} and the wide variation of the areas of its triangles, the quality of the approximation provided by the Voronoi-cot weights is lower (Fig. 5(c-f)). Irregularly-sampled patches on \mathcal{M} generally affect the smoothness of $K_t^D \mathbf{e}_i$ at smaller scales; increasing t improves the distribution and smoothness of the level sets. We also notice that the smoothness of the weighted diffusion kernel is preserved among all the scales. Decreasing the time parameter t reduces the support size of $F_k(\cdot, t)$ (c.f., Eq. (8)). A higher resolution of \mathcal{M} improves the quality of the level sets of the solution, which are uniformly distributed around the anchor point \mathbf{p} . On the contrary, the smoothness of the solution computed with respect to the weighted scalar product is guaranteed through all the scales in spite of the discretization of \mathcal{M} . Finally, Fig. 6 shows the behavior of the basis function $K_t \mathbf{e}_i$ on almost isometric shapes.

Executable paper.³ For the tests in Figs. 4, 5, 6, the reader can run the experiment *Diffusion.basis.exp.xml*, which computes the approximation $F_k(\cdot, t)$ in Eq. (8) associated to a predefined anchor point \mathbf{p}_i . As input, the user must select the *off* file of the input shape; a discretization of the Laplace-Beltrami operator among three possible choices (i.e., cotangent, Voronoi-cot, FEM weights); the value t of the time parameter; and the k number of Laplacian eigenfunctions. The default experiment computes the basis function on the 3-torus with 100 linear FEM Laplacian eigenpairs and $t = 0.5$. The level sets and colormap of the basis function are displayed in *diffusion-basis.png* and stored in *diffusion-basis.wrl*. Finally, the diffusion basis function is stored in *diffusion-basis.txt*.

³<https://collage.elsevier.com/collage/template?doi=10.0000/1352476045780>

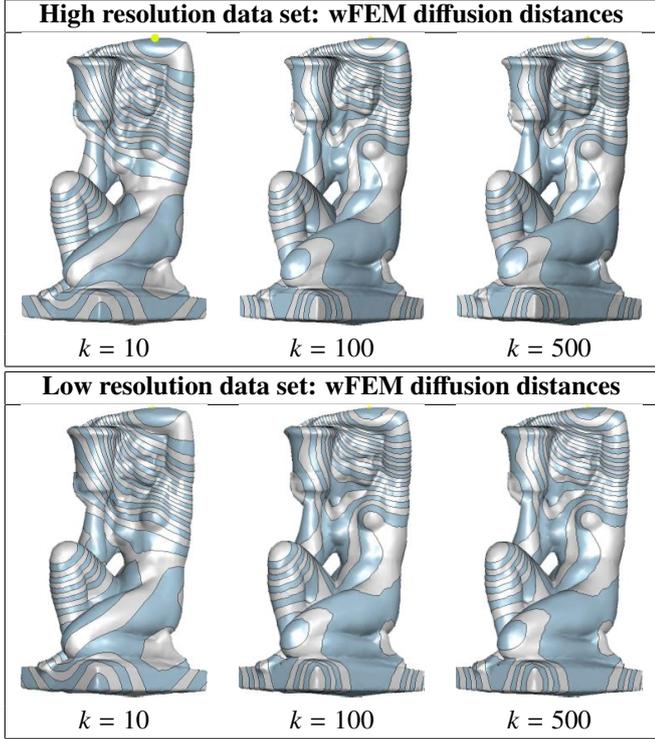


Figure 8: Level sets of the wFEM heat diffusion distances on the same shape with a high/low (first/second row) resolution and a different number k of Laplacian eigenpairs. For all the examples, $t = 0.5$ and the anchor point \mathbf{p}_i (yellow dot) is placed on the left arm.

5. Filtered spectral distances

Analyzing the spectral representation of the commute time, bi-harmonic, and heat diffusion distances, we show that they can be represented in a unified way through the definition of the filtered spectral distance (Sect. 5.1). Then, we focus our discussion on the linear FEM heat diffusion distances (Sect. 5.2).

5.1. Filtered spectral distances

Using the generalized eigensystem $\{(\lambda_i, \mathbf{x}_i)\}_{i=1}^n$ of the couple (L, B) and a positive *filter map* $\varphi: \mathbb{R} \rightarrow \mathbb{R}$, we introduce the matrix operator

$$L_\varphi: \mathbb{R}^n \rightarrow \mathbb{R}^n$$

$$\mathbf{f} \mapsto L_\varphi \mathbf{f} := \sum_{i=1}^n \varphi(\lambda_i) \langle \mathbf{f}, \mathbf{x}_i \rangle_B \mathbf{x}_i. \quad (9)$$

For any $\mathbf{f} \in \mathbb{R}^n$, $L_\varphi \mathbf{f}$ is a filtered version of the input map f , where the *filter coefficients* $(\varphi(\lambda_i))_{i=1}^n$ in Eq. (9) are the values of the filter map at the Laplacian eigenvalues. Since $\{\mathbf{x}_i\}_{i=1}^n$ is an orthonormal basis with respect to the scalar product induced by the mass matrix B , any function f is uniquely represented as $\mathbf{f} = \sum_{i=1}^n \langle \mathbf{f}, \mathbf{x}_i \rangle_B \mathbf{x}_i$ and $L\mathbf{f} = \sum_{i=1}^n \lambda_i \langle \mathbf{f}, \mathbf{x}_i \rangle_B \mathbf{x}_i$ is the spectral representation of $L\mathbf{f}$. According to these relations, choosing $\varphi(t) := 1$ or $\varphi(t) := t$ the (discrete) operator L_φ is equal to the identity operator or to the Laplacian matrix. Indeed, the filter map acts on the Laplacian eigenvalues in such a way that specific aspects of f are enhanced, smoothed, or removed.

Given the operator L_φ in Eq. (9) and the input function $f: \mathcal{M} \rightarrow \mathbb{R}$, let us introduce the *generalized inhomogeneous*

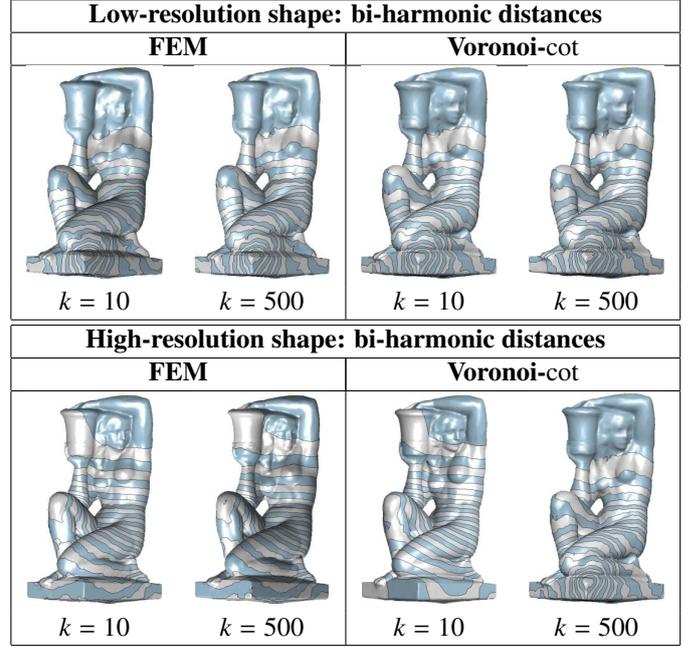


Figure 9: Variation of the bi-harmonic distance on a surface at different resolution, with respect to different Laplacian matrices and number of eigenpairs.

problem $L_\varphi \mathbf{u} = \mathbf{f}$. This equation reduces to the harmonic equation by selecting $\varphi(t) := t$, $f := 0$, and L_φ allows us to introduce the generalized Green kernel, together with the corresponding distances on \mathcal{M} . Using the relations

$$\mathbf{u} := \sum_{i=1}^n \alpha_i \mathbf{x}_i, \quad \mathbf{f} = \sum_{i=1}^n \langle \mathbf{f}, \mathbf{x}_i \rangle_B \mathbf{x}_i,$$

the generalized inhomogeneous problem is rewritten as

$$\sum_{i=1}^n \alpha_i \varphi(\lambda_i) \mathbf{x}_i = \sum_{i=1}^n \langle \mathbf{f}, \mathbf{x}_i \rangle_B \mathbf{x}_i.$$

Computing the B -scalar product of both sides of this equation with the Laplacian eigenvectors, we get that each coefficient is $\alpha_i = \frac{\langle \mathbf{f}, \mathbf{x}_i \rangle_B}{\varphi(\lambda_i)}$, $i = 1, \dots, n$. Then, the spectral representation of the solution to the generalized inhomogeneous problem is

$$\mathbf{u} = \sum_{k=1}^n \frac{\langle \mathbf{f}, \mathbf{x}_k \rangle_B}{\varphi(\lambda_k)} \mathbf{x}_k = G\mathbf{B}\mathbf{f},$$

where the *generalized Green kernel matrix* is defined as

$$G := X\Gamma^\dagger X^T, \quad \Gamma^\dagger := \text{diag}(\varphi^{-1}(\lambda_i))_{i=1}^n. \quad (10)$$

Assuming that $\varphi: \mathbb{R} \rightarrow \mathbb{R}$ is a strictly positive function, let us now introduce the *filtered spectral distance*

$$d(\mathbf{p}_i, \mathbf{p}_j) := \left[\sum_{k=1}^n \frac{|x_k(\mathbf{p}_i) - x_k(\mathbf{p}_j)|^2}{\varphi(\lambda_k)} \right]^{1/2}, \quad \mathbf{p}_i, \mathbf{p}_j \in \mathcal{M}, \quad (11)$$

where $x_k: \mathcal{M} \rightarrow \mathbb{R}$ is the piecewise linear map associated to the Laplacian eigenvector $\mathbf{x}_k := (x_k^{(i)})_{i=1}^n$; i.e., $x_k(\mathbf{p}_i) = x_k^{(i)}$.

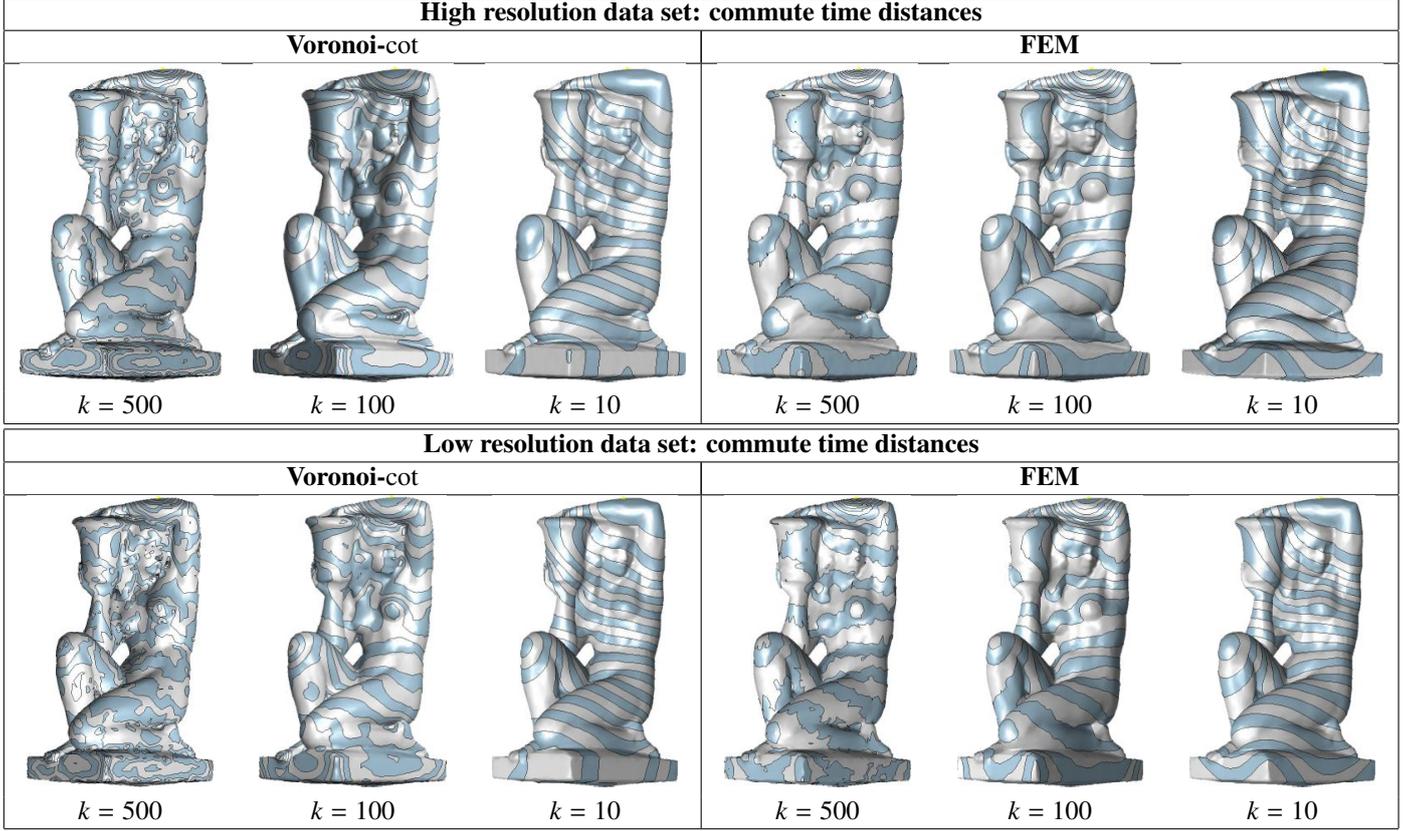


Figure 10: Level sets of the commute time distance on the same shape with a high/low (first/second row) resolution and a different number k of Laplacian eigenpairs.

According to the representation (11), diffusion, commute time, and bi-harmonic distances are achieved by selecting the following filters (Fig. 7)

$$\varphi(s) := \begin{cases} \exp(-st) & \text{diffusion distance,} \\ s & \text{commute time distance,} \\ s^2 & \text{bi-harmonic distance.} \end{cases} \quad (12)$$

Our interest on this formulation is motivated by its generalization of spectral distances and the possibility of adapting the filter function φ to the distribution of the Laplacian eigenvalues. Through specific choices of φ , local and global details of \mathcal{M} at different scales are identified in a way similar to the selection of the exponential and bi-harmonic weights.

Firstly, we show that $d(\cdot, \cdot)$ is a metric on \mathcal{M} . The positivity and symmetry of this distance are a direct consequence of the definition. For the nullity property, we notice that

$$\begin{aligned} d(\mathbf{p}_i, \mathbf{p}_j) = 0 & \iff x_k(\mathbf{p}_i) = x_k(\mathbf{p}_j), \quad \forall k, \\ (\mathbf{e}_i - \mathbf{e}_j)^T \mathbf{x}_k = 0, \forall k, & \iff i = j. \end{aligned}$$

The triangular inequality is verified as follows

$$\begin{aligned} d^2(\mathbf{p}_i, \mathbf{p}_j) &= \sum_{k=1}^n \frac{|x_k(\mathbf{p}_i) - x_k(\mathbf{p}_j)|^2}{\varphi(\lambda_k)} \\ &\leq \sum_{k=1}^n \frac{[|x_k(\mathbf{p}_i) - x_k(\mathbf{r})|^2 + |x_k(\mathbf{r}) - x_k(\mathbf{p}_j)|^2]}{\varphi(\lambda_k)} \\ &= d^2(\mathbf{p}_i, \mathbf{r}) + d^2(\mathbf{r}, \mathbf{p}_j), \quad \mathbf{p}_i, \mathbf{p}_j, \mathbf{r} \in \mathcal{N}. \end{aligned}$$

Rewriting the generalized Green kernel matrix as

$$G(\mathbf{p}_i, \mathbf{p}_j) := \sum_{k=1}^n \frac{x_k(\mathbf{p}_i)x_k(\mathbf{p}_j)}{\varphi(\lambda_k)}, \quad i, j = 1, \dots, n$$

The distance (11) is expressed in terms of G as

$$\begin{aligned} d^2(\mathbf{p}_i, \mathbf{p}_j) &:= \sum_{k=1}^n \frac{|x_k(\mathbf{p}_i) - x_k(\mathbf{p}_j)|^2}{\varphi(\lambda_k)} \\ &= \sum_{k=1}^n \frac{|x_k(\mathbf{p}_i)|^2}{\varphi(\lambda_k)} - 2 \sum_{k=1}^n \frac{x_k(\mathbf{p}_i)x_k(\mathbf{p}_j)}{\varphi(\lambda_k)} + \sum_{k=1}^n \frac{|x_k(\mathbf{p}_j)|^2}{\varphi(\lambda_k)} \\ &=_{(10)} G(\mathbf{p}_i, \mathbf{p}_i) - 2G(\mathbf{p}_i, \mathbf{p}_j) + G(\mathbf{p}_j, \mathbf{p}_j). \end{aligned}$$

Considering the embedding

$$\mathcal{E} : \mathbb{R}^3 \rightarrow \mathbb{R}^n, \quad \mathbf{p}_i \mapsto \sum_{k=1}^n \frac{x_k(\mathbf{p}_i)}{\varphi^{1/2}(\lambda_k)} \mathbf{x}_k,$$

of the vertices of \mathcal{M} to \mathbb{R}^n , the filtered spectral distance between \mathbf{p}_i and \mathbf{p}_j is equal to the B -norm of the embedded points $\mathcal{E}(\mathbf{p}_i), \mathcal{E}(\mathbf{p}_j)$; i.e., $d(\mathbf{p}_i, \mathbf{p}_j) = \|\mathcal{E}(\mathbf{p}_i) - \mathcal{E}(\mathbf{p}_j)\|_B$. Finally, the *generalized spectral function* is defined as $\mathbf{p}_i \in \mathcal{M} \mapsto G(\mathbf{p}_i, \mathbf{p}_i)$; if φ is the exponential map, then we get the auto diffusion functions introduced in [23] for shape analysis.

Computation of the filtered spectral distance. To make the computation of the generalized filtered distance feasible and tractable

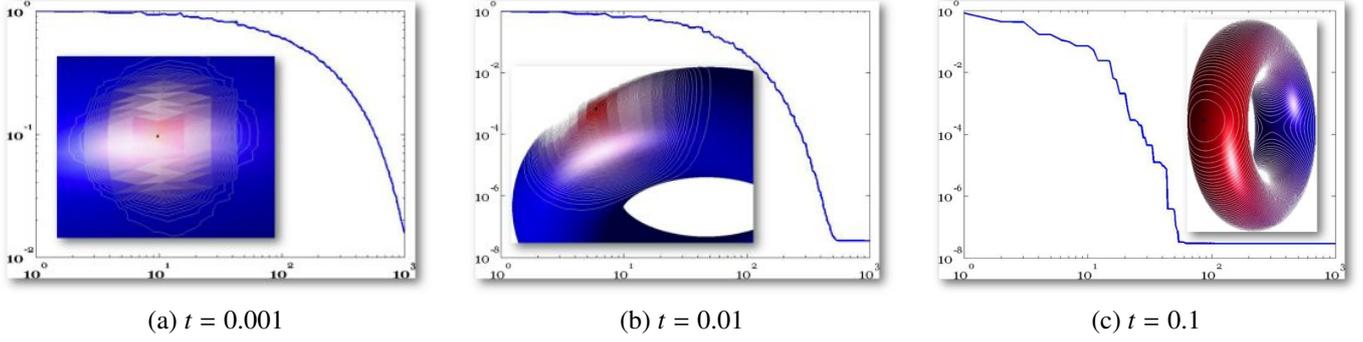


Figure 12: L^∞ error (y-axis) between the exact and approximated representation (9) of the function $L_\varphi \mathbf{e}_i$, computed using a different number k (x-axis) of eigenfunctions and (a-c) values of the time parameter t .

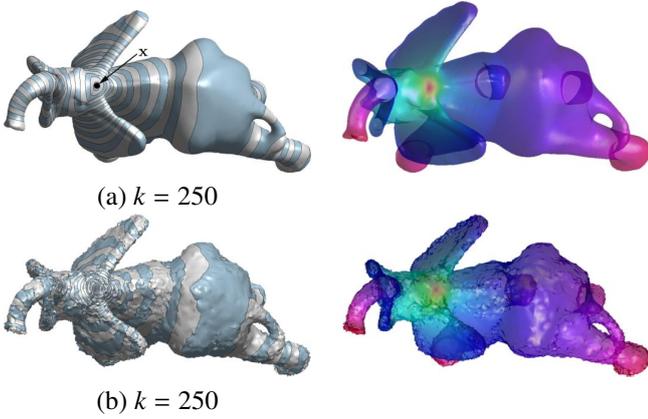


Figure 11: Level sets of $F_k(\cdot, t)$ on a (a, left) smooth and (b, left) noisy surface. (a-b, right) Behavior of the corresponding heat diffusion distances.

in real applications, the sum (11) is truncated by considering only the contribution related to the first k eigenpairs; i.e.,

$$d(\mathbf{p}_i, \mathbf{p}_j) := \left[\sum_{l=1}^k \frac{|x_l(\mathbf{p}_i) - x_l(\mathbf{p}_j)|^2}{\varphi(\lambda_l)} \right]^{1/2}, \quad \mathbf{p}_i, \mathbf{p}_j \in \mathcal{M}. \quad (13)$$

Indeed, the discretization of the Laplace-Beltrami operator, the filter φ , the number k of the eigenpairs, and the points $\mathbf{p}_i, \mathbf{p}_j$ will affect the value $d(\mathbf{p}_i, \mathbf{p}_j)$.

Selection of the filter map. Comparing the results on the smooth and irregularly sampled data highlights that the diffusion (Fig. 8), bi-harmonic (Fig. 9), and commute time (Fig. 10) distances have an analogous behavior, in terms of the shape and distribution of the corresponding level sets and colormaps. A similar discussion applies to the stability of the diffusion function and distance with respect to a change in the number of the eigenfunctions used in Eq. (13), noisy and irregularly-sampled data (Fig. 11). All these results confirm that the generalized eigen-system provides good results in terms of smoothness and distribution of the level sets with respect to uneven sampling densities, which are not affected by different resolutions of local areas of the input shape. Furthermore, increasing the number of generalized eigenpairs improves the quality of the approximation; on the contrary, a larger number of standard eigenvectors

results in a noisy map without improvements on the approximation accuracy and smoothness. In Fig. 12, we compare the L^∞ error ϵ_∞ (y-axis) between the analytical heat diffusion kernel on the torus with different discretizations of the Laplace-Beltrami operator and number k (x-axis) of eigenpairs. For small values of t , increasing k reduces the approximation error ϵ_∞ until it becomes almost constant and close to zero. In fact, in this case the behavior of the kernel is mainly influenced by the Laplacian eigenvectors related to smaller eigenvalues. Fig. 13 shows the robustness of the biharmonic distance from a source (black) point, which has been computed using the linear FEM mass matrix as weight, with respect to tiny and locally missing triangles, noise, holes.

Starting from the filter functions in Eq. (12), we can design new distances; the main constraint is that φ is strictly positive. The condition $\lim_{s \rightarrow +\infty} \varphi(s) = 0$ guarantees the smoothness of the resulting distance and a good approximation accuracy with the truncated approximation (13). On the one hand, the effects of the Laplacian eigenvalues of larger magnitude are enhanced by selecting as φ an increasing function such as a polynomial map $\varphi(s) := s^k, k \geq 1$. On the other hand, $\varphi(s) := \left[\prod_{i \in \mathcal{I}} (s - \lambda_i) \right]^{-1}$ suppresses the contribution of the eigenvectors related to a specific set $\{\lambda_k\}_{k \in \mathcal{I}}, \mathcal{I} \subseteq \{1, \dots, n\}$, of eigenvalues. The filter map $\varphi : \mathbb{R} \rightarrow \mathbb{R}$ must be chosen in such a way that the effects of the Laplacian eigenvalues and eigenfunctions are balanced in order to measure both local and global properties of the input shape. A simple way to generate new filters is to compute a convex combination of the maps in Eq. (12). For instance (Fig. 14), $\varphi_t(s) := (1 - \lambda) \exp(-st) + \lambda s^2, 0 \leq \lambda \leq 1$, is associated to a set of spectral distances that are closer to the diffusion or the bi-harmonic distance as much as λ is closer to 0 or 1, respectively.

Executable paper.⁴ For the tests in Figs. 8, 9, 10, the reader can run the experiment *SpectralDistances.exp.xml*, which computes the approximation $d(\mathbf{p}_i, \mathbf{p}_j), j = 1, \dots, n$, in Eq. (13) between a pre-defined anchor point \mathbf{p}_i and all the points of the input triangle mesh. The user must select the *off* file of the input shape; one spectral distance among three possible choices (i.e., diffusion, bi-harmonic, commute time distances); the value t of

⁴<https://collage.elsevier.com/collage/template?doi=10.0000/1352476136186>

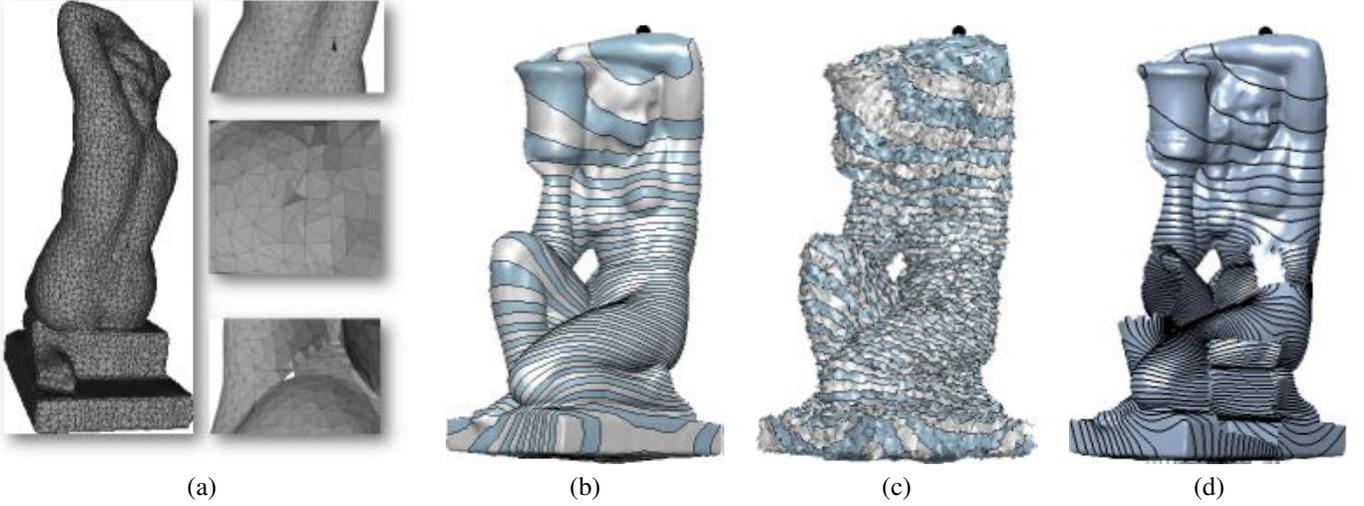


Figure 13: Robustness of the biharmonic distance from a source (black) point, which has been computed using the linear FEM mass matrix as weight, with respect to (b) tiny and missing triangles, (c) noise, (d) holes of an irregularly sampled surface (a) with local shape artifacts.

the time parameter for the diffusion distance; and the k number of Laplacian eigenpairs. The default experiment computes the bi-harmonic distance on the 3-torus with 100 linear FEM Laplacian eigenpairs. Then, the colormap and level sets of the computed distance are stored in *filtered-spectral-distance.wrl* and displayed in *filtered-spectral-distance.png*. The array of the values of the spectral distance is stored in *filtered-spectral-distance.txt*.

5.2. wFEM Heat diffusion distances

For the wFEM heat kernel, the diffusion distances cannot be defined by decomposing K_t as $Y^T Y$ or $Y^T B Y$; in fact, K_t is no longer symmetric. However, the self-adjointness of K_t with respect to the scalar product $\langle \cdot, \cdot \rangle_B$ allows us to introduce a time-depending scalar product $\langle \cdot, \cdot \rangle_t$, which is induced by a symmetric, positive definite matrix S_t . To this end, the entries of the matrix S_t are rewritten as $S_t(i, j) := \langle \mathbf{h}_i^t, \mathbf{h}_j^t \rangle_B$.

In $\mathcal{F}(\mathcal{M})$, we introduce the time-depending scalar product

$$\langle \mathbf{f}, \mathbf{g} \rangle_t := \langle K_t \mathbf{f}, \mathbf{g} \rangle_B = \langle \mathbf{f}, K_t \mathbf{g} \rangle_B = \mathbf{f}^T B X D_t X^T B \mathbf{g}, \quad (14)$$

$\mathbf{f}, \mathbf{g} \in \mathcal{F}(\mathcal{M})$. Firstly, we show that $K_t := X D_t X^T B$ is self-adjoint with respect to $\langle \cdot, \cdot \rangle_B$ (i.e., $\langle K_t \mathbf{f}, \mathbf{g} \rangle_B = \langle \mathbf{f}, K_t \mathbf{g} \rangle_B$, $\mathbf{f}, \mathbf{g} \in \mathbb{R}^n$). In fact, for any $\mathbf{f}, \mathbf{g} \in \mathbb{R}^n$

$$\langle \mathbf{f}, \mathbf{g} \rangle_t = \langle K_t \mathbf{f}, \mathbf{g} \rangle_B = \mathbf{f}^T B X D_t X^T B \mathbf{g} = \mathbf{f}^T B K_t \mathbf{g} = \langle \mathbf{f}, K_t \mathbf{g} \rangle_B.$$

To verify that the scalar product in (14) is well-posed, we show that the matrix $S_t := B X D_t X^T B$ is symmetric and positive definite. From the identity $S_t \stackrel{(4)}{=} X^{-T} D_t X^{-1}$, we get that S_t is symmetric with strictly positive eigenvalues $(\exp(-\lambda_i t))_{i=1}^n$; indeed, S_t is also positive definite.

The self-adjointness of K_t with respect to the mass matrix B allows us to rewrite the entries of the matrix S_t in terms of the scalar product $\langle \cdot, \cdot \rangle_B$ as $S_t(i, j) = \langle \mathbf{h}_i^t, \mathbf{h}_j^t \rangle_B$, $\mathbf{h}_i^t = K_{t/2} \mathbf{e}_i$, where the wFEM diffusion maps $h_i^t : \mathcal{M} \rightarrow \mathbb{R}$ is induced by the wFEM

heat kernel K_t (c.f., Sect. 4.2). In fact,

$$\begin{aligned} S_t &= B X D_{t/2} D_{t/2} X^T B \\ &\stackrel{(4)}{=} B X D_{t/2} (X^T B X) D_{t/2} X^T B \\ &= Y^T B Y, \quad Y := X D_{t/2} X^T B = K_{t/2}, \end{aligned}$$

and therefore

$$S_t(i, j) = \mathbf{e}_i^T S_t \mathbf{e}_j = \langle \mathbf{h}_i^t, \mathbf{h}_j^t \rangle_B, \quad \mathbf{h}_i^t := K_{t/2} \mathbf{e}_i.$$

Indeed, the wFEM diffusion distances

$$\|\mathbf{h}_i^t - \mathbf{h}_j^t\|_B^2 = K_{t/2}(i, i) - 2K_{t/2}(i, j) + K_{t/2}(j, j) \quad (15)$$

are uniquely defined by the entries of the matrix K_t . Finally, we show that the wFEM diffusion distance of the Laplacian eigenfunctions $\{\mathbf{x}_i\}_{i=1}^n$ is $\|\mathbf{x}_i - \mathbf{x}_j\|_t^2 = \exp(-\lambda_i t) + \exp(-\lambda_j t)$, $i \neq j$. To this end, we notice that

$$\langle \mathbf{x}_i, \mathbf{x}_j \rangle_t = \langle K_t \mathbf{x}_i, \mathbf{x}_j \rangle_B = \langle \exp(-\lambda_i t) \mathbf{x}_i, \mathbf{x}_j \rangle_B = \exp(-\lambda_i t) \delta_{ij}.$$

It follows that the entry $S_t(i, j)$ can be interpreted as the scalar product between the (discrete) embedding $\mathbf{h}_i^t, \mathbf{h}_j^t$ of $\mathbf{p}_i, \mathbf{p}_j$ in $\mathcal{F}(\mathcal{M})$ induced by the wFEM heat diffusion matrix. Furthermore, the distance (15) is uniquely determined by the entries of the kernel matrix. The aforementioned properties apply to the wFEM heat kernel and are analogous to those that hold in the continuous case. This analogy confirms that the proposed discretization and the choice of the intrinsic scalar product $\langle \cdot, \cdot \rangle_B$ in $\mathcal{F}(\mathcal{M})$ maintain the main features of the heat diffusion kernel, together with a higher robustness to data discretization with respect to previous work.

6. Conclusions and future work

Through the executable paper, we have provided a set of experiments on the discretization and main properties of the

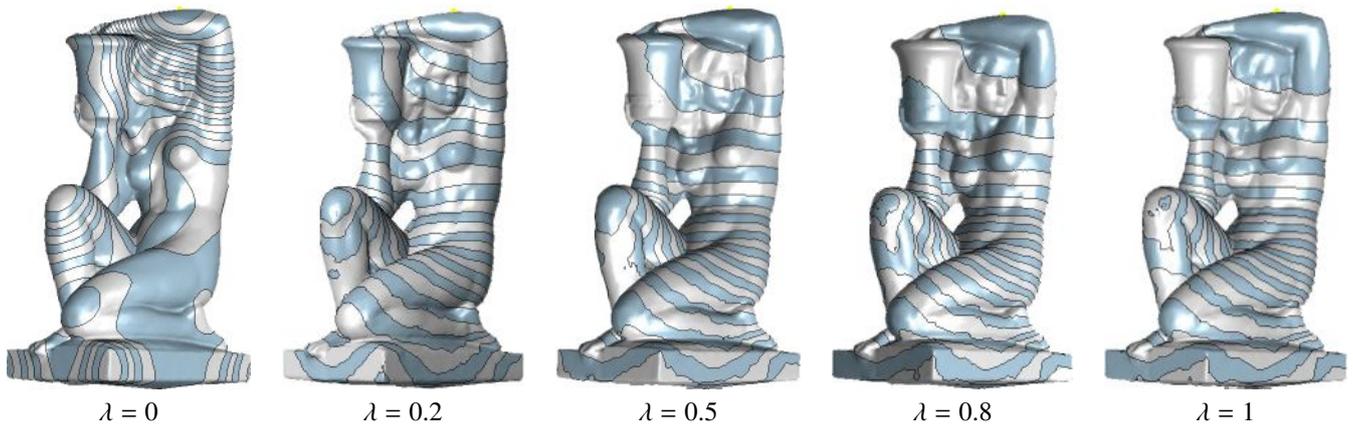


Figure 14: Level sets of the filtered spectral distance achieved as a convex combination of the filter maps associated to the diffusion ($\lambda = 0$, $t = 0.5$) and bi-harmonic ($\lambda = 1$) distances.

harmonic and bi-harmonic functions, the Laplacian eigenfunctions, and the solution to the heat diffusion equation, together with the corresponding bi-harmonic, commute time, and diffusion distances. This analysis has been based on the selection of different Laplacian weights, input shapes, transformations, initial boundary conditions, and source points. Analyzing the spectral representation of the aforementioned distances, we have shown that they can be represented in a unified way through a filtering of the Laplacian eigenpairs. Finally, adapting this filtering to the distribution of the Laplacian eigenvalues allow us to identify local and global details the input shape.

Acknowledgements

Special thanks are given to the anonymous Reviewers for their valuable comments, which helped us to improve the presentation and content of the paper, and to the Collage Support Team, in particular to Daniel Harezlak, for their help in the preparation of the experiments included in the paper. This work has been partially supported by the Research Project “*Methods and Techniques for the Development of Innovative Systems for Modeling and Analyzing Biomedical Data for Supporting Assisted Diagnosis*”, PO CRO Programme, European Social Funding Scheme, Regione Liguria, and the FP7 Marie Curie Initial Training Networks *MultiScaleHuman* - “Multi-scale Biological Modalities for Physiological Human Articulation”. 3D shapes are courtesy of the AIM@SHAPE Repository and the EPFL Computer Graphics and Geometry Laboratory.

References

- [1] C. J. Alpert, A. B. Kahng, and S.-Z. Yao. Spectral partitioning with multiple eigenvectors. *Discrete Applied Mathematics*, 90(1-3):3–26, 1999.
- [2] T. Banchoff. Critical points and curvature for embedded polyhedra. *Journal of Differential Geometry*, 1:245–256, 1967.
- [3] S. T. Barnard, A. Pothen, and H. D. Simon. A spectral algorithm for envelope reduction of sparse matrices. In *Proc. of the ACM Supercomputing*, pages 493–502, 1993.
- [4] M. Belkin and P. Niyogi. Laplacian eigenmaps for dimensionality reduction and data representation. *Neural Computations*, 15(6):1373–1396, 2003.
- [5] A. Bronstein, M. Bronstein, R. Kimmel, M. Mahmoudi, and G. Sapiro. A Gromov-Hausdorff framework with diffusion geometry for topologically-robust non-rigid shape matching. *International Journal of Computer Vision*, 2-3:266–286, 2010.
- [6] A. M. Bronstein. Spectral descriptors for deformable shapes. *CoRR*, abs/1110.5015, 2011.
- [7] A. M. Bronstein, M. M. Bronstein, B. Bustos, U. Castellani, M. Crisani, B. Falcidieno, L. J. Guibas, V. Murino I. Kokkinos, I. Isipiran, M. Ovsjanikov, G. Patanè, M. Spagnuolo, and J. Sun. SHREC 2010: robust feature detection and description benchmark. *Eurographics Workshop on 3D Object Retrieval*, 2010.
- [8] A. M. Bronstein, M. M. Bronstein, U. Castellani, B. Falcidieno, A. Fusiello, A. Godil, L.J. Guibas, I. Kokkinos, Z. Lian, M. Ovsjanikov, G. Patanè, M. Spagnuolo, and R. Toldo. SHREC 2010: robust large-scale shape retrieval benchmark. *Eurographics Workshop on 3D Object Retrieval*, 2010.
- [9] A. M. Bronstein, M. M. Bronstein, M. Ovsjanikov, and L. J. Guibas. Shape Google: geometric words and expressions for invariant shape retrieval. *ACM Transactions on Graphics*, 30(1), 2011.
- [10] M.M. Bronstein and I. Kokkinos. Scale-invariant heat kernel signatures for non-rigid shape recognition. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 1704–1711, 2010.
- [11] O. Chapelle, J. Weston, and B. Schölkopf. Cluster Kernels for Semi-Supervised Learning. In *Neural Information Processing Systems*, volume 15, pages 585–592, 2003.
- [12] C. S. Chua and R. Jarvis. Point signatures: A new representation for 3D object recognition. *International Journal of Computer Vision*, 25(1):63–85, 1997.
- [13] R. R. Coifman and S. Lafon. Diffusion maps. *Applied and Computational Harmonic Analysis*, 21(1):5–30, 2006.
- [14] F. de Goes, S. Goldenstein, and L. Velho. A hierarchical segmentation of articulated bodies. *Computer Graphics Forum*, 27(5):1349–1356, 2008.
- [15] M. Desbrun, M. Meyer, P. Schröder, and A. H. Barr. Implicit fairing of irregular meshes using diffusion and curvature flow. In *ACM Siggraph*, pages 317–324, 1999.
- [16] T. K. Dey, P. Ranjan, and Y. Wang. Convergence, stability, and discrete approximation of Laplace spectra. *ACM Symposium Discrete Algorithms*, pages 650–663, 2010.
- [17] J. Díaz, J. Petit, and M. Serna. A survey of graph layout problems. *ACM Computing Surveys*, 34(3):313–356, 2002.
- [18] S. Dong, P.-T. Bremer, M. Garland, V. Pascucci, and J. C. Hart. Spectral surface quadrangulation. *ACM Siggraph*, pages 1057–1066, 2006.
- [19] S. Dong, S. Kircher, and M. Garland. Harmonic functions for quadrilateral remeshing of arbitrary manifolds. *Computer Aided Geometric Design*, 22(5):392–423, 2005.
- [20] R. Dyer, R. H. Zhang, T. Moeller, and A. Clements. An investigation of the spectral robustness of mesh laplacians. 2007.
- [21] M. Fiedler. Algebraic connectivity of graphs. *Czechoslovak Mathematical Journal*, 23(98):298–305, 1973.

- [22] R. Gal, A. Shamir, and D. Cohen-Or. Pose-oblivious shape signature. *IEEE Transactions on Visualization and Computer Graphics*, 13(2):261–271, 2007.
- [23] K. Gebal, J. A. Bærentzen, H. Aanæs, and R. Larsen. Shape analysis using the auto diffusion function. *Computer Graphics Forum*, 28(5):1405–1413, 2009.
- [24] E. Gine and V. Koltchinskii. Empirical graph Laplacian approximation of Laplace-Beltrami operators: large sample results. *IMS Lecture Notes Monograph Series*, 51:238, 2006.
- [25] G. Golub and G.F. VanLoan. *Matrix Computations*. John Hopkins University Press, 2nd Edition, 1989.
- [26] C. Gotsman, X. Gu, and A. Sheffer. Fundamentals of spherical parameterization for 3D meshes. In *ACM Siggraph 2003*, pages 358–363, 2003.
- [27] M. Hein. Uniform convergence of adaptive graph-based regularization. In *Conference on Learning Theory*, pages 50–64. Springer, 2006.
- [28] M. Hein, J.-Y. Audibert, and U. von Luxburg. From graphs to manifolds – weak and strong pointwise consistency of graph laplacians. In *Learning Theory*, volume 3559 of *Lecture Notes in Computer Science*, pages 470–485. Springer, 2005.
- [29] M. Hilaga, Y. Shinagawa, T. Kohmura, and T. L. Kunii. Topology matching for fully automatic similarity estimation of 3D shapes. In *ACM Siggraph*, pages 203–212, 2001.
- [30] V. Jain and H. Zhang. A spectral approach to shape-based retrieval of articulated 3D models. *Computer Aided Design*, 39:398–407, 2007.
- [31] V. Jain, H. Zhang, and O. van Kaick. Non-rigid spectral correspondence of triangle meshes. *International Journal on Shape Modeling*, 13(1):101–124, 2007.
- [32] K. Polthier, K. Hildebrandt and M. Wardetzky. On the convergence of metric and geometric properties of polyhedral surfaces. *Geometriae Dedicata*, pages 89–112, 2006.
- [33] Z. Karni and C. Gotsman. Spectral compression of mesh geometry. In *ACM Siggraph 2000*, pages 279–286, 2000.
- [34] B. Kim and J. Rossignac. Geofilter: Geometric selection of mesh filter parameters. *Computer Graphics Forum*, 24(3):295–302, 2005.
- [35] L. Kobbelt, S. Campagna, J. Vorsatz, and H.-P. Seidel. Interactive multi-resolution modeling on arbitrary meshes. In *ACM Siggraph 1998*, pages 105–114, 1998.
- [36] Y. Koren. On spectral graph drawing. pages 496–508. *Lecture Notes in Computer Science*, Vol. 2697, Springer Verlag., 2003.
- [37] S. Lafon, Y. Keller, and R. R. Coifman. Data fusion and multicue data matching by diffusion maps. *IEEE Transactions on Pattern Analysis Machine Intelligence*, 28(11):1784–1797, 2006.
- [38] B. Levy. Laplace-Beltrami eigenfunctions: towards an algorithm that understands geometry. In *Proc. of Shape Modeling and Applications*, page 13, 2006.
- [39] R. Liu and H. Zhang. Mesh segmentation via spectral embedding and contour analysis. *Eurographics 2007*, 26:385–394, 2007.
- [40] J.-L. Mallet. Discrete smooth interpolation. *ACM Transactions on Graphics*, 8(2):121–144, 1989.
- [41] F. Memoli. Spectral Gromov-Wasserstein distances for shape matching. In *Workshop on Non-Rigid Shape Analysis and Deformable Image Alignment*, pages 256–263, 2009.
- [42] J. Milnor. *Morse Theory*, volume 51 of *Annals of mathematics studies*. Princeton University Press, 1963.
- [43] B. Mohar and S. Poljak. Eigenvalues in combinatorial optimization. *Combinatorial and graph-theoretical problems in linear algebra*, 23(98):107–151, 1993.
- [44] A. Nealen, T. Igarashi, O. Sorkine, and M. Alexa. Laplacian mesh optimization. In *Proc. of Computer graphics and interactive techniques*, pages 381–389, 2006.
- [45] A. Y. Ng, M. I. Jordan, and Y. Weiss. On spectral clustering: Analysis and an algorithm. In *Advances in Neural Information Processing Systems 14*, pages 849–856. MIT Press, 2001.
- [46] X. Ni, M. Garland, and J. C. Hart. Fair morse functions for extracting the topological structure of a surface mesh. In *ACM Siggraph 2004*, pages 613–622, 2004.
- [47] R. Ohbuchi, A. Mukaiyama, and S. Takahashi. A frequency-domain approach to watermarking 3D shapes. *Computer Graphics Forum*, 21(3), 2002.
- [48] R. Ohbuchi, S. Takahashi, T. Miyazawa, and A. Mukaiyama. Watermarking 3D polygonal meshes in the mesh spectral domain. In *Graphics Interface 2001*, pages 9–17, 2001.
- [49] M. Ovsjanikov, Q. Mérigot, F. Mémoli, and L. Guibas. One point isometric matching with the heat kernel. *ACM Symposium on Discrete Algorithms*, pages 650–663, 2010.
- [50] G. Patanè and B. Falcidieno. Computing smooth approximations of scalar functions with constraints. *Computer & Graphics*, 33(3):399–413, 2009.
- [51] G. Patanè and B. Falcidieno. Multi-scale feature spaces for shape processing and analysis. In *Proc. of Shape Modeling International*, pages 113–123, 2010.
- [52] U. Pinkall and K. Polthier. Computing discrete minimal surfaces and their conjugates. *Experimental Mathematics*, 2(1):15–36, 1993.
- [53] M.M. Bronstein R. Litmana, A.M. Bronstein. Diffusion-geometric maximally stable component detection in deformable shapes. *Computers & Graphics*, 35(3):549–560, 2011.
- [54] M. Reuter, S. Biasotti, D. Giorgi, G. Patanè, and M. Spagnuolo. Discrete Laplace-Beltrami operators for shape analysis and segmentation. *Computer & Graphics*, 33(3):381–390, 2009.
- [55] M. Reuter, F.-E. Wolter, and N. Peinecke. Laplace-Beltrami spectra as Shape-DNA of surfaces and solids. *Computer-Aided Design*, 38(4):342–366, 2006.
- [56] M. Reuter, F.-E. Wolter, and N. Peinecke. Laplace-Beltrami spectra as Shape-DNA of surfaces and solids. *Computer-Aided Design*, 38(4):342–366, 2006.
- [57] S. Rosenberg. *The Laplacian on a Riemannian Manifold*. Cambridge University Press, 1997.
- [58] S. T. Roweis and L. K. Saul. Nonlinear dimensionality reduction by locally linear embedding. *Science*, 290:2323–2326, 2000.
- [59] M.R. Ruggeri, G. Patanè, M. Spagnuolo, and D. Saupe. Spectral-driven isometry-invariant matching of 3D shapes. *International Journal of Computer Vision*, 89(2-3):248–265, 2010.
- [60] R. M. Rustamov. Laplace-Beltrami eigenfunctions for deformation invariant shape representation. In *Proc. of the Symposium on Geometry processing*, pages 225–233, 2007.
- [61] R. M. Rustamov. Multiscale biharmonic kernels. *Computer Graphics Forum*, 30(5):1521–1531, 2011.
- [62] B. Schoelkopf and A. J. Smola. *Learning with Kernels*. The MIT Press, 2002.
- [63] J. Shi and J. Malik. Normalized cuts and image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 22(8):888–905, 2000.
- [64] A. Singer. From graph to manifold laplacian: the convergence rate. *Applied and Computational Harmonic Analysis*, 21(1):128–134, 2006.
- [65] A. J. Smola and R. I. Kondor. Kernels and regularization on graphs. In *Conference on Learning Theory*, pages 144–158, 2003.
- [66] O. Sorkine. Differential representations for mesh processing. *Computer Graphics Forum*, 25(4):789–807, 2006.
- [67] O. Sorkine, D. Cohen-Or, D. Irony, and S. Toledo. Geometry-aware bases for shape approximation. *IEEE Transactions on Visualization and Computer Graphics*, 11(2):171–180, 2005.
- [68] O. Sorkine, D. Cohen-Or, and S. Toledo. High-pass quantization for mesh encoding. In *Proc. of the Symposium on Geometry Processing*, pages 42–51, 2003.
- [69] O. Sorkine, Y. Lipman, D. Cohen-Or, M. Alexa, C. Roessl, and H.-P. Seidel. Laplacian surface editing. In *Proc. of the Symposium on Geometry Processing*, pages 179–188, 2004.
- [70] D. A. Spielman and S.-H. Teng. Spectral partitioning works: Planar graphs and finite element meshes. *Linear Algebra and its Applications*, 421:284–305, 2007.
- [71] J. Sun, M. Ovsjanikov, and L. J. Guibas. A concise and provably informative multi-scale signature based on heat diffusion. *Computer Graphics Forum*, 28(5):1383–1392, 2009.
- [72] G. Taubin. A signal processing approach to fair surface design. In *ACM Siggraph 1995*, pages 351–358, 1995.
- [73] G. Taubin. 3D Geometry Compression and Progressive Transmission. In *Eurographics Tutorials*, 1999.
- [74] B. Vallet and B. Lévy. Spectral geometry processing with manifold harmonics. *Comput. Graph. Forum*, 27(2):251–260, 2008.
- [75] A. Vaxman, M. Ben-Chen, and C. Gotsman. A multi-resolution approach to heat kernels on discrete surfaces. *ACM Transactions on Graphics*, 29(4):1–10, 2010.
- [76] Y. Wang. Approximating gradients for meshes and point clouds via dif-

- fusion metric. *Computer Graphics Forum*, 28:1497–1508(12), 2009.
- [77] B. Xiaoa, E. R. Hancock, and R.C. Wilsonb. Geometric characterization and clustering of graphs using heat kernel embeddings. *Image and Vision Computing*, 28(6):1003 – 1021, 2010.
- [78] G. Xu. Discrete Laplace-Beltrami operators and their convergence. *Computer Aided Geometric Design*, 8(21):398–407, 2007.
- [79] H. Zhang and E. Fiume. Butterworth filtering and implicit fairing of irregular meshes. In *Proc. of the Pacific Conference on Computer Graphics and Applications*, page 502, 2003.
- [80] H. Zhang and R. Liu. Mesh segmentation via recursive and visually salient spectral cuts. In *Proc. of Vision, Modeling, and Visualization*, pages 429–436, 2005.
- [81] H. Zhang, O. van Kaick, and R. Dyer. Spectral methods for mesh processing and analysis. In *Eurographics State-of-the-art Report*, pages 1–22, 2007.
- [82] K. Zhou, J. Synder, B. Guo, and H.-Y. Shum. Iso-charts: stretch-driven mesh parameterization using spectral analysis. In *Proc. of the Symposium on Geometry processing*, pages 45–54, 2004.
- [83] X. Zhu, Z. Ghahramani, and J. Lafferty. Semi-supervised learning using gaussian fields and harmonic functions. In *International Conference on Machine Learning*, pages 912–919, 2003.
- [84] G. Zigelman, R. Kimmel, and N. Kiryati. Texture mapping using surface flattening via multidimensional scaling. *IEEE Transactions on Visualization and Computer Graphics*, 8(2):198–207, 2002.
- [85] A. Zomorodian and G. Carlsson. Computing persistent homology. *Discrete Computational Geometry*, 33(2):249–274, 2005.