Multi-Scale Feature Spaces for Shape Processing and Analysis

Giuseppe Patanè and Bianca Falcidieno
Consiglio Nazionale delle Ricerche
Istituto di Matematica Applicata e Tecnologie Informatiche
Genova, Italy
{patane,falcidieno}@ge.imati.cnr.it

***Abstract*—In digital geometry processing and shape modeling, the Laplace-Beltrami and the heat diffusion operator, together with the corresponding Laplacian eigenmaps, harmonic and geometry-aware functions, have been used in several applications, which range from surface parameterization, deformation, and compression to segmentation, clustering, and comparison. Using the linear FEM approximation of the Laplace-Beltrami operator, we derive a discrete heat kernel that is linear, stable to an irregular sampling density of the input surface, and scale covariant. With respect to previous work, this last property makes the kernel particularly suitable for shape analysis and comparison; in fact, local and global changes of the surface correspond to a re-scaling of the time parameter without affecting its spectral component. Finally, we study the scale spaces that are induced by the proposed heat kernel and exploited to provide a multi-scale approximation of scalar functions defined on 3D shapes.**

***Keywords*-scale-space methods; heat kernel; Laplacian matrix; spectral analysis; signal and function smoothing; critical points; shape analysis**

1. Introduction

In digital geometry processing and shape modeling, the Laplace-Beltrami and the heat diffusion operator, together with the corresponding Laplacian eigenmaps, harmonic and geometry-aware functions, have been used in several contexts. For instance, the eigenvectors of the graph Laplacian are exploited to project the input signal into the frequency [23], [50] or a lower dimensional space [2], [22]; to smooth surfaces [8], [19], [20], [27], [29], [48] and signals [33], [45]; to compress 3D shapes [18], [42]; to process meshes [31], [32] and graphs [9], [21]. The Laplacian eigenvectors [34], [35], [37], [38], [46] are also used for parameterizing surfaces homeomorphic to the sphere [15] or with an arbitrary genus [51], [52]. In the frequency space, mesh segmentation [24], [49], shape correspondence [16] and comparison [34], [37], [38], [39], [40] have been successfully addressed. Finally, mesh Laplacian operators, whose stability and convergence have been studied in [12], [17], [47] and [34], [36], [37], play a central role in the definition of differential coordinates for surface deformation [41], [43] and quadrilateral remeshing [2], [10], [11], [30].

The Laplace-Beltrami operator is strictly related to the heat diffusion equation, which provides an embedding of a given function in a hierarchy of smoothed approximations. The heat kernel and the associated diffusion metric have been exploited to approximate the gradient of maps defined on triangulated surfaces or point sets [25]. Finally, the heat kernel has important applications such as shape segmentation [14], [7] and matching [3], [28] with diffusion distances [6], [22], multi-scale [44] and isometry-invariant [3], [28] signatures.

The main limitation of the current discretizations of the heat kernel k_t proposed by previous work is its dependence on the scale of $\mathcal{M}$; i.e., rescaling $\mathcal{M}$ to $\alpha\mathcal{M}$, $\alpha \in \mathbb{R}^+$, k_t becomes $\tilde{k}_t = \alpha^2 k_{\alpha^2 t}$. This means that global and local changes to the surface provide different kernels, which must be "normalized" before their use as shape signatures. To partially overcome this drawback, in [26] the heat kernel signature is sampled logarithmically in time, scaled, and derived; then, the descriptor is defined by the magnitude of the Fourier Transform coefficients. In this case, the normalization steps are neither unique nor intrinsically defined by the input shape. Alternatively, in [14] the eigenvalues are normalized by the first non-null eigenvalue, thus guaranteeing that $\tilde{k}_t = \alpha^2 k_t$ but without removing the scale term α^2.

Overview and contributions

Combining the weak formulation of the heat equation with the linear FEM approximation of the Laplace-Beltrami operator, we derive a discretization of the heat kernel that is linear, *intrinsically* invariant to surface scalings, and stable to irregular sampling densities. As main feature with respect to previous work, the proposed kernel is *scale covariant*; i.e., local and global changes of the surface correspond to a re-scaling of the time parameter without affecting the spectral term of the kernel. Then, the *scale invariance* is achieved by normalizing the eigenvalues by the first non-null eigenvalue, thus avoiding *a-posteriori* changes to the kernel itself or to the surface. Due to these two properties, the kernel is particularly suitable for shape comparison and the proposed approach improves the invariance and stability of the corresponding shape signatures. To assess these aims, the underlying idea is to compute the solution of the heat equation using the inner product induced by the mass matrix of the linear FEM discretization. In fact, this scalar product is adapted to the sampling density of $\mathcal{M}$ through the distribution of the areas of its triangles.

Combining the heat kernel operator with the Laplacian spectral properties, we define a *feature map* $\Phi_t : \mathcal{F}(\mathcal{M}) \rightarrow \mathcal{F}(\mathcal{M})$, where $\mathcal{F}(\mathcal{M})$ is the space of piecewise linear scalar functions defined on a triangulated surface $\mathcal{M}$. The real parameter t and the map Φ_t induce a multi-scale hierarchy of approximations $\{\Phi_t(f)\}_t$ of an arbitrary scalar function $f : \mathcal{M} \rightarrow \mathbb{R}$. Through Φ_t, the local noise of f (e.g., noisy level sets and/or critical points of f with low persistence) is removed with small time values, and the global behavior of f is enhanced at larger scales. In this way, the global structure of f is separated from local details, which are preserved or discarded according to the target accuracy. Constraining the approximation to preserve a set of feature values of the input map, we also define a *multi-scale* and *feature-driven approximation*. To further characterize Φ_t, we will discuss its linearity, the estimation of the approximation error, and the computational aspects behind the induced approximations. Finally, recent results on the robustness and usefulness of the proposed heat kernel for shape retrieval are presented in [4].

The paper is organized as follows. Section II defines a discrete heat kernel on trianglulated surfaces. In Section III, we introduce the feature spaces and their main properties. Future work is outlined in Section IV.

2. Discrete heat kernel induced by the linear FEM weights

Assuming that $h : \mathcal{N} \subseteq \mathbb{R}^d \rightarrow \mathbb{R}$ is a scalar function defined on a compact manifold $\mathcal{N}$, the *scale-based representation* $H : \mathcal{N} \times \mathbb{R} \rightarrow \mathbb{R}$ *of* h, with $\lim_{t\rightarrow 0} H(\mathbf{x}, t) = h(\mathbf{x})$, $\mathbf{x} \in \mathcal{N}$, provides an embedding of h in a *hierarchy* of simplified and/or smoothed approximations. One way to accomplish the construction of this one-parameter family is to solve the linear diffusion problem (*heat equation*)

$$\begin{cases} \partial_t H(\mathbf{x},t) = -\frac{1}{2}\Delta H(\mathbf{x},t), \\ H(\mathbf{x},0) = h(\mathbf{x}), \end{cases} \qquad \mathbf{x} \in \mathcal{N}, \qquad t \in \mathbb{R}, \tag{1}$$

with Δ Laplace-Beltrami operator. Since $\mathcal{N}$ is compact, $H(\mathbf{x},t) := k_t(\mathbf{x},\cdot) \star h = \int_{\mathcal{N}} k_t(\mathbf{x},\mathbf{y})h(\mathbf{y})d\mathbf{y}$ is the scale-based representation of h, where $k_t(\cdot,\cdot)$ is the *heat kernel* and $\star$ is the convolution operator. Indicating with ϕ_i the i^{th} eigenfunction of Δ related to the eigenvalue λ_i (i.e., $\Delta\phi_i = \lambda_i\phi_i$), $k_t(\mathbf{x},\mathbf{y}) := \sum_{i=0}^{+\infty} \exp(-\frac{1}{2}\lambda_i t)\phi_i(\mathbf{x})\phi_i(\mathbf{y})$ is the spectral decomposition of k_t. Finally, the heat diffusion and the Laplace-Beltrami operator share the same eigenfunctions $\{\phi_i\}_{i=1}^{+\infty}$ and λ_i, $\rho_i := \exp(-\frac{1}{2}\lambda_i t)$ are the corresponding eigenvalues.

In the following, we introduce the weighted linear FEM discretization of the heat kernel (Section II-A), prove its invariance to scalings (Section II-B), discuss its properties and computational aspects (Section II-C).

2.1 Weak and discrete formulation of the heat equation

In the following, we derive the weak formulation of the heat equation through the Galerkin formulation and introduce its weighted linear FEM discretization.

Weak formulation of the heat equation

To convert the heat equation into a variational problem, we miltiply (1) with test functions $\varphi \in \mathcal{C}^2$ and integrate the resulting relation over $\mathcal{N}$; i.e.,

$$\int_{\mathcal{N}} \varphi\partial_t H d\sigma + \frac{1}{2}\int_{\mathcal{N}} \varphi\Delta H d\sigma = 0,$$

where $d\sigma$ is the surface element. Then, using the Green formula we get

$$\int_{\mathcal{N}} \varphi\partial_t H d\sigma + \frac{1}{2}\int_{\mathcal{N}} \nabla(H,\varphi) d\sigma = 0,$$

or equivalently

$$\int_{\mathcal{N}} \varphi\partial_t H d\sigma + \frac{1}{2}\int_{\mathcal{N}} \sum_{r,s} g^{r,s}(\partial_{x_r} H)(\partial_{x_s}\varphi) d\sigma = 0, \tag{2}$$

where g^{rs}'s are the entries of the first fundamental matrix. To compute the solution of the variational problem, we apply the Galerkin techniques. More precisely, we choose n linearly independent basis functions $\mathcal{B} := \{\varphi_i\}_{i=1}^n$, $\varphi_i : \mathcal{N} \rightarrow \mathbb{R}$ and consider the linear space $\mathcal{F}$ generated by $\mathcal{B}$. We now approximate the solution $H(\mathbf{x},t)$ of Equation (2) as a linear combination of the basis functions in $\mathcal{B}$; i.e.,

$$\tilde{H}(\mathbf{x},t) := \sum_{i=1}^{n} a_i(t)\varphi_i(\mathbf{x}), \qquad \mathbf{x} \in \mathcal{N}, \qquad t \in \mathbb{R}^+.$$

Then, for any value $t \in \mathbb{R}^+$ we compute the n coefficients $\mathbf{a}(t) := (a_i(t))_{i=1}^n$ by imposing that Equation (2) is satisfied by $\tilde{H}(\mathbf{x},t)$ for any test function φ_j, $j = 1,\ldots,n$; i.e.,

$$\int_{\mathcal{N}} \varphi_j\partial_t\tilde{H} d\sigma + \frac{1}{2}\int_{\mathcal{N}} \varphi_j\Delta\tilde{H} d\sigma = 0.$$

Indeed, for $j = 1,\ldots,n$ we get the relation

$$\begin{aligned} &\sum_{i=1}^{n} \partial_t a_i(t) \int_{\mathcal{N}} \varphi_i\varphi_j d\sigma + \\ &+ \frac{1}{2}\sum_{i=1}^{n} a_i(t) \int_{\mathcal{N}} \sum_{r,s} g^{rs}(\partial_{x_r}\varphi_i)(\partial_{x_s}\varphi_j) d\sigma = 0. \end{aligned} \tag{3}$$

Introducing the matrices $L := (L(i,j))_{i,j=1}^n$ and $B := (B(i,j))_{i,j=1}^n$, whose elements are

$$\begin{cases} L(i,j) := \int_{\mathcal{N}} \sum_{r,s} g^{rs}(\partial_{x_r}\varphi_i)(\partial_{x_s}\varphi_j) d\sigma, \\ B(i,j) := \int_{\mathcal{N}} \varphi_i\varphi_j d\sigma, \end{cases}$$

Equation (3) is rewritten as

$$\sum_{i=1}^{n} B(i,j)\partial_t a_i(t) + \frac{1}{2}\sum_{i=1}^{n} L(i,j)a_i(t) = 0,$$

$j = 1,\ldots,n$, and its matrix formulation is $B\partial_t\mathbf{a}(t) + \frac{1}{2}L\mathbf{a}(t) = \mathbf{0}$. An analogous relation can be derived for the boundary condition $H(\mathbf{x},0) = h(\mathbf{x})$, $\mathbf{x} \in \mathcal{N}$. Since B is the Gram matrix associated to $\mathcal{B}$, it is invertible and the previous system of equations becomes $\left[\partial_t + \frac{1}{2}B^{-1}L\right]\mathbf{a}(t) = \mathbf{0}$. Comparing this expression with Equation (1), it follows that the Laplace-Beltrami operator is discretized by the matrix $\tilde{L} := B^{-1}L$.

Weighted linear FEM discretization of the heat kernel

To define the discrete counterpart of Equation (1), let $\mathcal{M} := (M, T)$ be a triangulated surface that approximates $\mathcal{N}$. Here $M := \{\mathbf{p}_i,\ i = 1, \dots, n\}$ is a set of n vertices and T is an *abstract simplicial complex*, which contains the adjacency information. Choosing a common set of piecewise linear basis and test functions $\{\varphi_i\}_{i=1}^n$ such that $\varphi_i(\mathbf{p}_j) := \delta_{ij}$, the matrices L and B previously introduced become [38], [39], [46]

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

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

Here, $N(i) := \{j :\ (i,j) \text{ edge}\}$ is the 1-*star* of i; $|\ t_i\ |$ is the area of the triangle t_i; t_r and t_s are the triangles that share the edge (i,j); and α_{ij}, β_{ij} are the angles opposite to the edge (i,j).

According to the weak formulation of the heat equation previously discussed, the discrete heat problem is

$$\begin{cases} \partial_t F(\mathbf{p}, t) = -\frac{1}{2}\tilde{L}F(\mathbf{p}, t), & \mathbf{p} \in \mathcal{M}, \\ F(\mathbf{p}_i, 0) = f(\mathbf{p}_i), & i = 1, \dots, n, \end{cases} \tag{4}$$

where $\tilde{L} := B^{-1}L$ is the weighted Laplacian matrix associated to $\mathcal{M}$ and $F : \mathcal{M} \times \mathbb{R} \rightarrow \mathbb{R}$ is the unknown function, which is a $n \times 1$ vector for each value of the real parameter t. To compute the solution to (4), let us consider the generalized eigensystem $\{(\lambda_i, \mathbf{x}_i)\}_{i=1}^n$ of (L, B), which satisfies the relations $L\mathbf{x}_i = \lambda_i B\mathbf{x}_i$, $i = 1, \dots, n$. Since the Laplacian eigenvectors $\{\mathbf{x}_i\}_{i=1}^n$ form a basis of $\mathbb{R}^n$ and $(F(\mathbf{p}_i, t))_{i=1}^n \in \mathbb{R}^n$, for any $t \in \mathbb{R}$ we express the solution as $F(\cdot, t) = \sum_{i=1}^n \alpha_i(t)\mathbf{x}_i$, where $\alpha := (\alpha_i(t))_{i=1}^n$ is the unknown vector.

Using the aforementioned relations, the invertibility of the matrix B, and the linear independence of the Laplacian eigenfunctions, the following identities hold

$$\begin{aligned} & B\partial_t\left[\sum_{i=1}^n \alpha_i(t)\mathbf{x}_i\right] = -\frac{1}{2}L\left[\sum_{i=1}^n \alpha_i(t)\mathbf{x}_i\right] \longleftrightarrow \\ & \sum_{i=1}^n \alpha_i'(t)\mathbf{x}_i = -\frac{1}{2}\sum_{i=1}^n \lambda_i\alpha_i(t)\mathbf{x}_i \longleftrightarrow \\ & \sum_{i=1}^n \left[\alpha_i'(t) + \frac{1}{2}\lambda_i\alpha_i(t)\right]\mathbf{x}_i = \mathbf{0} \longleftrightarrow \\ & \alpha_i'(t) + \frac{1}{2}\lambda_i\alpha_i(t) = 0, \quad i = 1, \dots, n. \end{aligned}$$

Then, the unknown vector satisfies a system of differential equations of first order. Since the Laplacian eigenvectors are orthonormal with respect to $\langle\cdot,\cdot\rangle_B$, the boundary conditions $\{F(\mathbf{p}_i, 0) = f(\mathbf{p}_i)\}_{i=1}^n$ are equivalent to

$$\begin{aligned} & (F(\mathbf{p}_i, 0))_{i=1}^n = \mathbf{f} \leftrightarrow \sum_{i=1}^n \alpha_i(0)\mathbf{x}_i = \mathbf{f} \leftrightarrow \\ & \alpha_i(0) = \langle\mathbf{f}, \mathbf{x}_i\rangle_B, \quad i = 1, \dots, n. \end{aligned}$$

We conclude that the coefficients $\alpha(t) := (\alpha_i(t))_{i=1}^n$ that satisfy the system of differential equations

$$\begin{cases} \alpha_i'(t) = -\frac{1}{2}\lambda_i\alpha_i(t), \\ \alpha_i(0) = \langle\mathbf{f}, \mathbf{x}_i\rangle_B, \end{cases} \quad i = 1, \dots, n,$$

are $\alpha_i(t) = \exp\left(-\frac{1}{2}\lambda_i t\right)\langle\mathbf{f}, \mathbf{x}_i\rangle_B$, $i = 1, \dots, n$ and the scale-based representation of $f : \mathcal{M} \rightarrow \mathbb{R}$ is

$$F(\cdot, t) = \sum_{i=1}^n \exp\left(-\frac{1}{2}\lambda_i t\right)\langle\mathbf{f}, \mathbf{x}_i\rangle_B\mathbf{x}_i, \quad t \in \mathbb{R}, \tag{5}$$

or, in matrix form,

$$\begin{aligned} & F(\cdot, t) = XD_tX^TB\mathbf{f}, \qquad X := [\mathbf{x}_1, \dots, \mathbf{x}_n], \\ & D_t := \operatorname{diag}\left(\exp\left(-\tfrac{1}{2}\lambda_1 t\right), \dots, \exp\left(-\tfrac{1}{2}\lambda_n t\right)\right). \end{aligned} \tag{6}$$

We refer to $K_t := XD_tX^TB$ as the *weighted linear FEM heat kernel*; if $B := I$, then $K_t := XD_tX^T$ is the *linear FEM heat kernel*, which is commonly used by previous work. Figure 1 shows the multi-scale approximations of a noisy map, whose behavior at different scales is almost the same in terms of level sets. Their main difference is related to the number and location of the critical points (Figure 1(i-k)).

2.2 Scale invariance of the discrete heat kernel

We now verify that the discretization of the heat kernel previously introduced is scale covariant and intrinsically independent of uniform rescalings. This means that rescaling $\mathcal{M}$ to $\alpha\mathcal{M}$ changes the kernel $k_t(\mathcal{M})$ to $k_{t/\alpha^2}(\alpha\mathcal{M})$; i.e., only the time component is rescaled. Roughly speaking, this property is guaranteed by the mass matrix B, which changes according to the rescaling of the input surface and compensates the variation of the corresponding Laplacian eigenvalues. Finally, we normalize the eigenvalues to make the kernel scale invariant.

According to the continuous case previously introduced, $\{(\exp(-\frac{1}{2}\lambda_i t), \mathbf{x}_i)\}_{i=1}^n$ is the eigensystem of K_t; in fact, from equation (6) it follows that

$$\begin{aligned} K_t\mathbf{x}_i &= \sum_{j=1}^n \exp\left(-\frac{1}{2}\lambda_j t\right)\langle\mathbf{x}_i, \mathbf{x}_j\rangle_B\mathbf{x}_j \\ &= \exp\left(-\frac{1}{2}\lambda_i t\right)\mathbf{x}_i, \qquad i = 1, \dots, n. \end{aligned}$$

Rescaling $\mathcal{M}$ by a factor α, the matrix L is the same, the mass matrix B becomes $\alpha^2 B$, and the eigensystem $\{(\lambda_i, \mathbf{x}_i)\}_{i=1}^n$ changes into $\{(\frac{\lambda_i}{\alpha^2}, \frac{\mathbf{x}_i}{\alpha})\}_{i=1}^n$. Replacing the matrices B, D_t, and X in $K_t(\mathcal{M}) := XD_tX^TB$ with $\alpha^2 B$, D_{t/α^2}, and X/α, we have

$$K_t(\alpha\mathcal{M}) = \left(\frac{X}{\alpha}\right)D_{t/\alpha^2}\left(\frac{X^T}{\alpha}\right)\alpha^2 B = K_{t/\alpha^2}(\mathcal{M}),$$

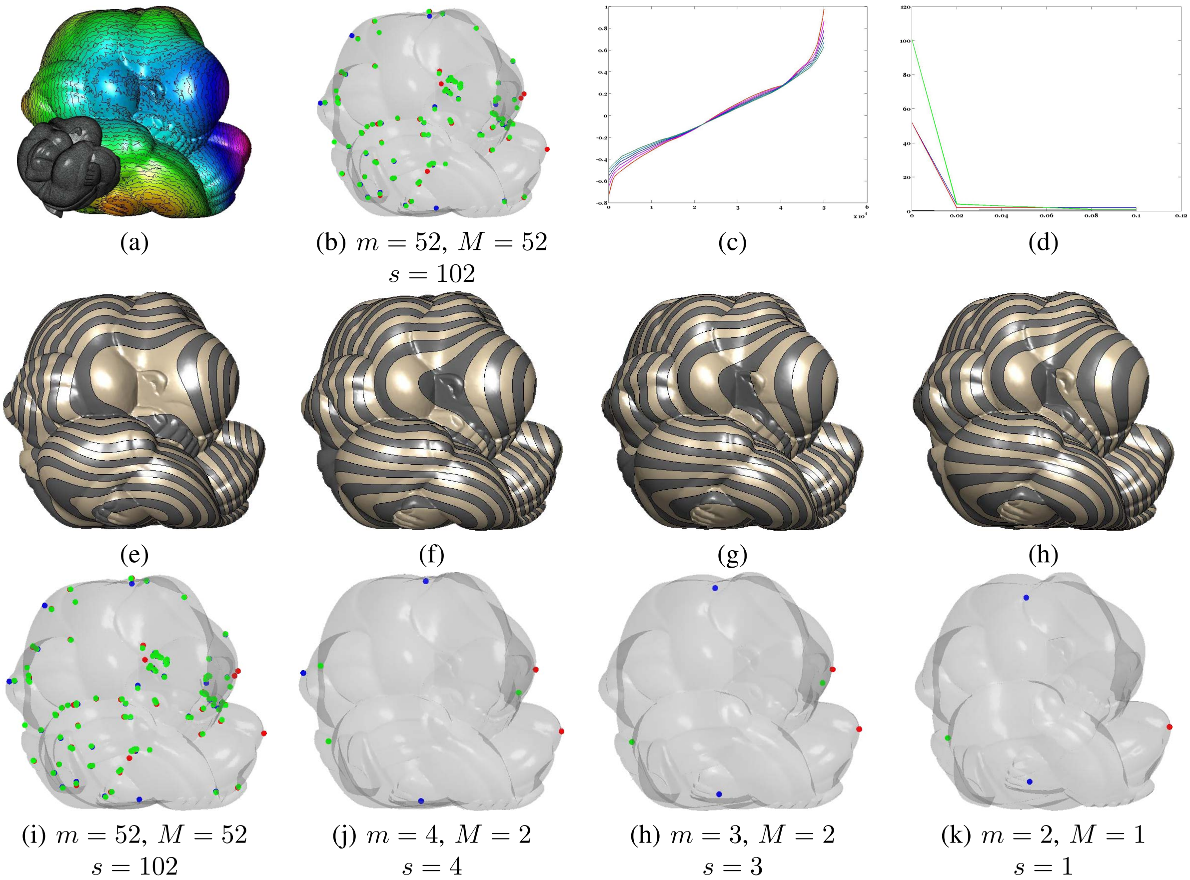

Fig. 1. (a) Level sets and (b) critical points of a noisy map f. Variation (y-axis) of the (c) approximating map Φ_t and (d) critical points at different scales t (x-axis). (e-h) Level sets and (i-k) critical points of $\Phi_t(f)$ at each scale. The M maxima, m minima, and s saddles are shown in red, blue, and green. See also Figure 3.

for any $\alpha \in \mathbb{R}\backslash 0$. It follows that the discrete heat kernel induced by the linear FEM discretization of the Laplace-Beltrami operator is *scale covariant*. Finally, the entries of $K_t := (K_t(i,j))_{i,j=1}^n$ are

$$K_t(i,j) := \sum_{a,b=1}^{n} \exp\left(-\frac{1}{2}\lambda_b t\right) \mathbf{x}_b^{(i)} \mathbf{x}_b^{(a)} B(a,j),$$

and the heat kernel signature is

$$K_t(i,i) := \sum_{a,b=1}^{n} \exp\left(-\frac{1}{2}\lambda_b t\right) \mathbf{x}_b^{(i)} \mathbf{x}_b^{(a)} B(a,i),$$

for $i = 1,\dots,n$. Assuming that $\lambda_2 \neq 0$ and re-defining K_t as $\tilde{K}_t := X\tilde{D}_t X^T B$, where $\tilde{D}_t$ is the $n \times n$ diagonal matrix with entries $\exp\left(-\frac{1}{2}\frac{\lambda_i}{\lambda_2}t\right)\delta_{ij}$, $i,j = 1,\dots,n$, we get that $\tilde{K}_t$ is *scale invariant*.

2.3 Remarks on the weighted linear FEM heat kernel

Before discussing the computation and possible generalizations of the weighted linear FEM heat kernel, we briefly recall that the heat kernel is commonly discretized by the $n \times n$ matrix whose entries are $k_t(\mathbf{p}_r,\mathbf{p}_s) := \sum_{i=1}^n \exp(-\frac{1}{2}\lambda_i t)\mathbf{x}_i^{(r)}\mathbf{x}_i^{(s)}$, where $\mathbf{x}_i^{(r)}$ is the r^{th} component of the Laplacian eigenfunction $\mathbf{x}_i$ [44]. Using the relation between the convolution operator and Equation (1), the heat approximation of f at time t is usually computed as the vector $(F(\mathbf{p}_i,t))_{i=1}^n$, $t \in \mathbb{R}$, whose entries are

$$F(\mathbf{p}_i,t) := \sum_{j=1}^{n} k_t(\mathbf{p}_i,\mathbf{p}_j)a(j)f(\mathbf{p}_j), \quad i = 1,\dots,n, \tag{7}$$

where $a(j)$ is the Voronoi area of the vertex $\mathbf{p}_j$. Recent results on the robustness and usefulness of the proposed kernel for shape retrieval are presented in [4]. Here, among the compared feature detection and description methods the heat kernel descriptors provide the highest overall repeatability in most of the transformation classes. Among these descriptors and with respect to Equation (7), our discretization of the heat kernel provides the higher robustness against topological and local scale changes, irregular sampling densities, and noise.

Numerical computation and iterative discretization of the heat kernel

The decay of the *filter factor* $\sigma_i := \exp(-\frac{1}{2}\lambda_i t)$ in (5) as λ_i decreases and the computational bottleneck of evaluating the whole Laplacian spectrum suggest us to consider only a part of the Laplacian spectrum. To this end, the sum in (5) is truncated by considering only the contribution related to the first k eigenvalues and

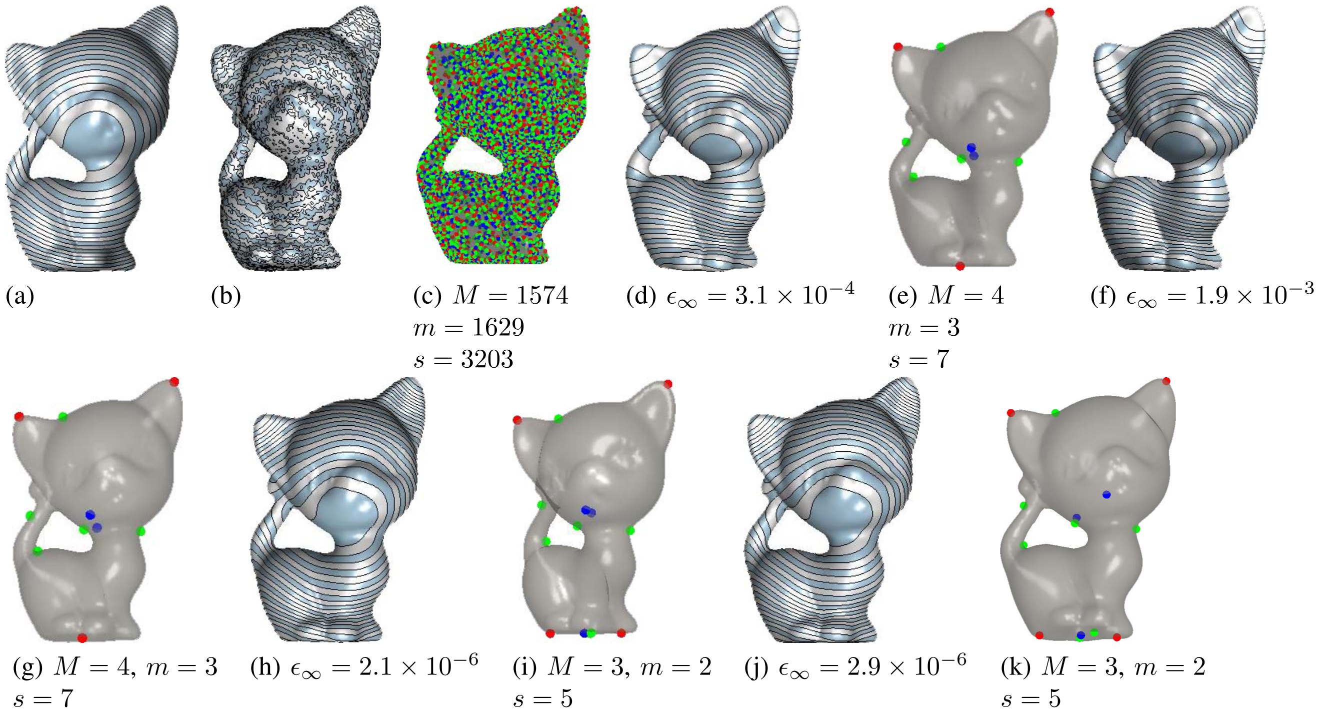

Fig. 2. (a) Input map and (b,c) its noisy perturbation f. Level sets and critical points of the multi-scale approximation of f at two scales using the first (d,e), (f,g) 10 and (h,i), (j,k) 100 Laplacian eigenvalues. The value ϵ_∞ refers to the discrepancy between the input map and its approximation. A lager number of eigenfunctions provides a higher approximation accuracy and a number of critical points of the same order. See also Figure 7.

eigenvectors; i.e.,

$$\begin{aligned} F_k(\cdot, t) &= \sum_{i=1}^{k} \exp\left(-\frac{1}{2}\lambda_i t\right) \langle \mathbf{f}, \mathbf{x}_i \rangle_B \mathbf{x}_i \\ &= X_k D_t^{(k)} X_k^T B \mathbf{f}, \qquad t \in \mathbb{R}, \end{aligned}$$

where the $n \times k$ full matrix $X_k := [\mathbf{x}_1, \ldots, \mathbf{x}_k]$ has the first k eigenfunctions as columns and $D_t^{(k)} := \mathrm{diag}(\exp(-\frac{1}{2}\lambda_1 t), \ldots, \exp(-\frac{1}{2}\lambda_k t)) \in \mathrm{Gl}_k(\mathbb{R})$ 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 the least-squares approximation of f in the linear space generated by the first k eigenfunctions and with respect to the norm $\|\cdot\|_B$. Note that both the parameters k and t define the hierarchy of approximations. In fact, reducing the number of basis functions results in a smoothing of the input map, a larger approximation error, and a further simplification of its critical points, with more emphasis on those with low persistence values (Figure 2(d-g), (h-k)). The approximation error is measured as $\epsilon_\infty := \frac{\|\mathbf{f} - \Phi_t(\mathbf{f})\|_\infty}{\|\mathbf{f}\|_\infty}$.

In our implementation, the multi-scale hierarchy is generated by varying the parameter t on an uniform sampling of the interval $[0, \lambda_k^{-1}]$; generally, from five to ten scales $\{t_i := \frac{i}{10}\lambda_k^{-1}\}_{i=1}^{10}$ are enough to provide a set of approximations that highlight the global and local behavior of the input map (Section III-C). As shown in [46], the first k Laplacian eigenvalues and eigenvectors are computed in superlinear time and the computational cost for the evalutation of the scale-based representation at time t and for any f is $O(kn)$.

Alternatively, the partial derivative $\partial_t F(\mathbf{p}, \cdot)$ is approximated by the incremental ratio

$$\partial_t F(\mathbf{p}, \cdot) \approx \frac{F(\mathbf{p}, t + dt) - F(\mathbf{p}, t)}{dt}, \quad dt \to 0.$$

Indeed, the heat equation becomes

$$\begin{cases} \frac{F(\mathbf{p}_i, t+dt) - F(\mathbf{p}_i, t)}{dt} = -\frac{1}{2} B^{-1} L F(\mathbf{p}_i, t), \\ F(\mathbf{p}_i, 0) = f(\mathbf{p}_i), \end{cases} \quad i = 1, \ldots, n,$$

which is equivalent to the sparse linear system

$$BF(\cdot, t + dt) = \left[-\frac{1}{2} dt\, L + B\right] F(\cdot, t). \qquad (8)$$

In this case, the mass matrix still provides better results than the identity matrix. However, the smoothing effect is not as good as the one provided by the schemes previously introduced. In fact, the quality of the approximation of the partial derivative with respect to t decreases while increasing t (Figure 3(d,e)). Furthermore, the smoothing effect of the projection on the Laplacian eigenfunctions is not exploited in (8). This approach also resembles the one discussed in [8], where the implicit integration of the diffusion equation is used to fairing a noisy 3D surface $\mathcal{M}$. In this case, the new approximation $\mathcal{M}^{(k+1)}$ of $\mathcal{M}^{(k)}$ is computed by solving a sparse linear system whose coefficient matrix involves the Laplacian matrix of $\mathcal{M}^{(k)}$ and is updated at each iteration k. On the contrary, we iteratively smooth the input signal f without recomputing the eigensystem in (5) or the coefficient matrix of the corresponding linear system in (8).

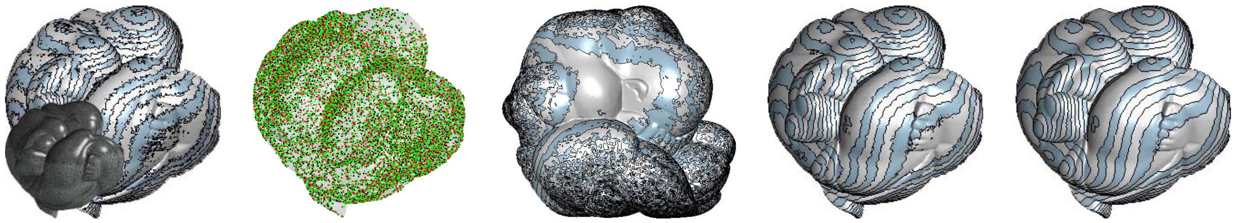

(a) (b) (c) (d) $t := 0.1$ (e) $t := 0.2$

Fig. 3. (a) Level sets and (b) critical points of a noisy map f on an irregularly-sampled surface $\mathcal{M}$. Map achieved by applying (c) the discretization (7) and (d,e) the iterative kernel approximation $\Phi_t(f)$ in (8) with two different scales. Comparing (d) and (e) shows that a higher value of t provides a smoother result. See also Figure 6.

Generalized scale-based approximations

Note that we can also consider the generalized discrete problem

$$\begin{cases} \partial_t F(\mathbf{p},t) = -\frac{1}{2}\varphi(L)F(\mathbf{p},t), & \mathbf{p}\in\mathcal{M}, \\ F(\mathbf{p}_i,0) = f(\mathbf{p}_i), & i=1,\dots,n, \end{cases}$$

where $\varphi:\mathbb{R}\rightarrow\mathbb{R}$ is a *transfer function*. Solving this problem as done before shows that the *generalized scale-based approximation* of the function f is $F(\cdot,t) = \sum_{i=1}^n \exp\left(-\frac{1}{2}\varphi(\lambda_i)t\right)\langle\mathbf{f},\mathbf{x}_i\rangle_B\mathbf{x}_i$. This expression differs from (5) only for the spectral coefficients. Common choices of φ are the polynomials of degree strictly lower than n, which improve the convergence to zero of the filter factors in (5). For more details on the choice of the filter factor and related parameters, we refer the reader to [19].

3. Scale-spaces induced by the heat kernel

In the following, we introduce a *feature map* $\Phi_t:\mathcal{F}(\mathcal{M})\rightarrow\mathcal{F}(\mathcal{M})$ on the space of piecewise linear (PL, for short) maps defined on a triangulated surface $\mathcal{M}$; i.e., $\mathcal{F}(\mathcal{M}) := \{f:\mathcal{M}\rightarrow\mathbb{R},\ f \text{ PL function}\}$. To characterize Φ_t, we study its main properties such as linearity, approximation error, numerical computation, and stability against noise (Section III-A). Then, the canonic basis functions (Section III-B) are used to define a feature-driven and multi-scale approximation of discrete signals (Section III-C).

3.1 Feature map: definition and properties

To introduce a linear structure in the feature space $\mathcal{F}(\mathcal{M})$, we consider the following operations: $\forall f,g\in\mathcal{F}(\mathcal{M})$, $(f+g)(\mathbf{p}_i) := f(\mathbf{p}_i)+g(\mathbf{p}_i)$ and $(\alpha f)(\mathbf{p}_i) := \alpha f(\mathbf{p}_i)$, $\forall\alpha\in\mathbb{R}$, $i=1,\dots,n$. Note that $\mathcal{F}(\mathcal{M})$ is isomorphic to $\mathbb{R}^n$; i.e., any vector $\mathbf{f} := (f_i)_{i=1}^n\in\mathbb{R}^n$ is associated to a unique map f on $\mathcal{M}$ such that the value $f(\mathbf{p}_i)$ is equal to f_i and *viceversa*.

In $\mathcal{F}(\mathcal{M})$, we define the linear *feature map* associated to the weighted linear FEM heat kernel as

$$\begin{array}{llll} \Phi_t: & \mathcal{F}(\mathcal{M}) & \longrightarrow & \mathcal{F}(\mathcal{M}) \\ & \mathbf{f} & \mapsto & \Phi_t(\mathbf{f}) := XD_tX^TB\mathbf{f}. \end{array}$$

It follows that $\Phi_t(\mathbf{f})$ is the discrete solution to the heat equation with initial condition $\Phi_0(\mathbf{f}) := \mathbf{f}$ and the multi-scale approximation of f is given by $\Phi_t(\mathbf{f}) = \sum_{i=1}^n \exp(-\frac{1}{2}\lambda_i t)\langle\mathbf{f},\mathbf{x}_i\rangle_B\mathbf{x}_i$. Indeed, the behavior of $\Phi_t(\mathbf{f})$ is mainly controlled by the first eigenmaps, which code the global structure of f. The contribution of the eigenfunctions related to the largest eigenvalues, which locate local features and noise, is smoothed by the corresponding filter factors that decrease to zero as the eigenvalue magnitude grows.

We now verify that the feature map is *self-adjoint* with respect to the scalar product induced by B; in fact,

$$\begin{aligned} \langle\mathbf{f},\Phi_t(\mathbf{g})\rangle_B &= \mathbf{f}^TB\Phi_t(\mathbf{g}) \\ &= \mathbf{f}^TBXD_tX^TB\mathbf{g} \\ &= \langle\mathbf{f}^TBXD_tX^T,\mathbf{g}\rangle_B \\ &= \langle\Phi_t(\mathbf{f}),\mathbf{g}\rangle_B, \qquad \forall\mathbf{f},\mathbf{g}\in\mathcal{F}(\mathcal{M}). \end{aligned}$$

The self-adjointness of Φ_t is equivalent to the symmetry of the linear FEM kernel XD_tX^T for $B := I$. Since the matrix XD_tX^TB is invertible, the inverse functional $\Phi_t^{-1}:\mathcal{F}(\mathcal{M})\rightarrow\mathcal{F}(\mathcal{M})$ is defined as $\Phi_t^{-1}(\mathbf{f}) = B^{-1}X^{-1}D_t^{-1}X^{-T}\mathbf{f}$, $\mathbf{f}\in\mathcal{F}(M)$.

Linearity of Φ_t with respect to f

Let $f,g:\mathcal{M}\rightarrow\mathbb{R}$ be two scalar functions defined on $\mathcal{M}$, $\Phi_t(\mathbf{f})$ and $\Phi_t(\mathbf{g})$ their corresponding scale-based representations. Then, using the linearity of the scalar product we get that the scale-based representation $\Phi_t(\alpha\mathbf{f}+\beta\mathbf{g})$ of $\alpha f+\beta g$, $\alpha,\beta\in\mathbb{R}$, is a linear combination $\alpha\Phi_t(\mathbf{f})+\beta\Phi_t(\mathbf{g})$ of the corresponding approximations; in fact,

$$\begin{aligned} \Phi_t(\alpha\mathbf{f}+\beta\mathbf{g}) &= (XD_tX^TB)(\alpha\mathbf{f}+\beta\mathbf{g}) \\ &= \alpha\Phi_t(\mathbf{f})+\beta\Phi_t(\mathbf{g}), \qquad t\in\mathbb{R}. \end{aligned}$$

Estimation of the error between f and $\Phi_t(f)$

We now estimate the discrepancy between the input scalar function $\mathbf{f}\in\mathcal{F}(\mathcal{M})$ and the corresponding smooth approximation $\Phi_t(\mathbf{f})\in\mathcal{F}(\mathcal{M})$. Rewriting the f-values as $\mathbf{f} = \sum_{i=1}^n\langle\mathbf{f},\mathbf{x}_i\rangle_B\mathbf{x}_i$, the approximation error between $\mathbf{f}$ and $\Phi_t(\mathbf{f})$ is estimated as

$$\begin{aligned} \|\mathbf{f}-\Phi_t(\mathbf{f})\|_B &= \left\|\sum_{i=1}^n\left[1-\exp\left(-\frac{1}{2}\lambda_i t\right)\right]\langle\mathbf{f},\mathbf{x}_i\rangle_B\mathbf{x}_i\right\|_B \\ &\leq \sum_{i=1}^n\left|1-\exp\left(-\frac{1}{2}\lambda_i t\right)\right|\,|\langle\mathbf{f},\mathbf{x}_i\rangle_B| \\ &\leq 2\sum_{i=1}^n|\langle\mathbf{f},\mathbf{x}_i\rangle_B|, \qquad t\in\mathbb{R}, \end{aligned}$$

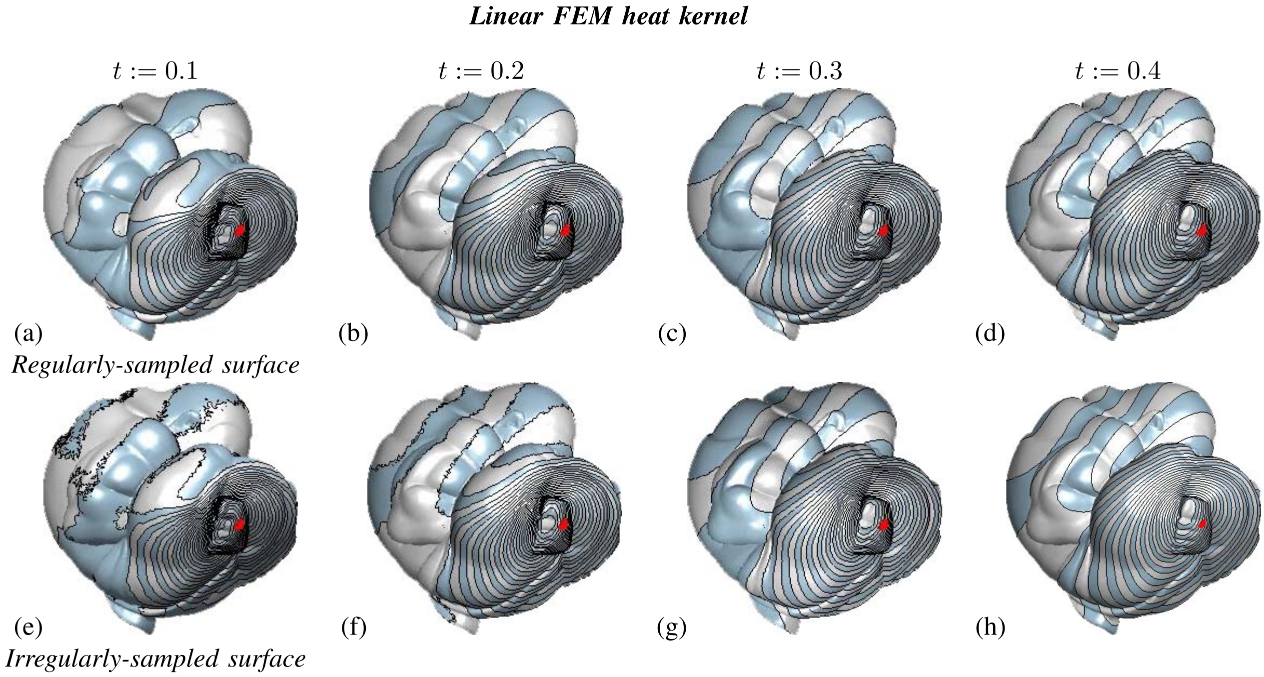

Fig. 4. Multi-scale canonic basis function $\Phi_t(\mathbf{e}_i)$ computed with respect to the linear FEM kernel and centered at a point $\mathbf{p}_i$ on the elbow of the (a-d) regularly- and (e-h) irregularly-sampled surfaces in Figure 1(a) and 3(a), respectively. See also Figure 5.

where we have used the relations $\|\mathbf{x}_i\|_B = 1$, $i = 1, \ldots, n$. Finally, the *energy* of Φ_t with respect to the Euclidean norm is bounded as follows

$$\begin{aligned}\|\Phi_t(\mathbf{f})\|_2 &\leq \|XD_tX^TB\mathbf{f}\|_2\\ &\leq \lambda^2_{\max}(X)\lambda_{\max}(B)\|\mathbf{f}\|_2.\end{aligned}$$

Using the norm induced by B, we have that

$$\begin{aligned}\|\Phi_t(\mathbf{f})\|_B &\leq \sum_{i=1}^{n}\exp\left(-\frac{1}{2}\lambda_i t\right)|\langle\mathbf{f},\mathbf{x}_i\rangle_B|\,\|\mathbf{x}_i\|_B\\ &\leq \sum_{i=1}^{n}|\langle\mathbf{f},\mathbf{x}_i\rangle_B|.\end{aligned}$$

Stability of Φ_t against noise

Perturbing the f-values as $\tilde{\mathbf{f}} := \mathbf{f} + \mathbf{e}$, the difference between the corresponding multi-scale approximations $\Phi_t(\mathbf{f})$, $\Phi_t(\tilde{\mathbf{f}})$ is bounded as follows

$$\begin{aligned}\|\Phi_t(\mathbf{f}) - \Phi_t(\tilde{\mathbf{f}})\|_B &= \|\Phi_t(\mathbf{e})\|_B\\ &\leq \left\|\sum_{i=1}^{n}\exp\left(-\frac{1}{2}\lambda_i t\right)\langle\mathbf{e},\mathbf{x}_i\rangle_B\mathbf{x}_i\right\|_B\\ &\leq \|\mathbf{e}\|_B\sum_{i=1}^{n}\exp\left(-\frac{1}{2}\lambda_i t\right)\\ &\leq n\|\mathbf{e}\|_B, \quad t\in\mathbb{R}.\end{aligned}$$

If we consider the l_2 norm, then this inequality becomes

$$\begin{aligned}\|\Phi_t(\mathbf{f}) - \Phi_t(\tilde{\mathbf{f}})\|_2 &= \|XD_tX^TB\mathbf{e}\|_2\\ &\leq \|X\|_2^2\|B\|_2\|\mathbf{e}\|_2\\ &\leq \lambda^2_{\max}(X)\lambda_{\max}(B)\|\mathbf{e}\|_2.\end{aligned}$$

Note that the previous bounds are independent of the time step and proportional to the perturbation vector $\mathbf{e}$.

Stop criteria

The parameter t provides an intrinsic notion of scale for the approximation and smoothing of the input scalar function f. Increasing t, local shape features together with noise are filtered out until the global structure of f appears and is preserved through diffusion. For the diffusion process, we consider two stop conditions. The first criterion is to terminate the iteration when the l_∞-error $\|\Phi_{t_i}(\mathbf{f}) - \Phi_{t_{i+1}}(\mathbf{f})\|_\infty$ between two consecutive approximations is lower than a given threshold. As second criterion, we consider the variation of the critical points of $\Phi_t(\mathbf{f})$ and stop the iteration when two consecutive approximations share the same set of critical points. For piecewise linear maps defined on triangulated surfaces, the critical points are computed according to [1]. Increasing t forces $\Phi_t(\mathbf{f})$ to converge to a constant function on $\mathcal{M}$; in fact $\lim_{t\to+\infty}\Phi_t(\mathbf{f}) = \mathbf{1}$.

For the tests of this paper, we have applied the first stop criterion and our experiments, which deal with noisy maps defined on irregularly-sampled surfaces, confirm that the smoothness of $\Phi_t(\mathbf{f})$ is effectively improved by the mass matrix. This property is strictly related to the higher accuracy of the FEM methods with respect to other discretizations (e.g., constant [5], cotangent [35], normalized cotangent weights [8]) of the Laplace-Beltrami operator, which has been investigated in the pioneering work [34], [39].

3.2 *Canonic basis in $\mathcal{F}(\mathcal{M})$*

Even though the function space $\mathcal{F}(\mathcal{M})$ is isomorphic to $\mathbb{R}^n$, only specific choices of the vector $\mathbf{f}\in\mathbb{R}^n$ identify scalar functions on $\mathcal{M}$ with interesting properties (e.g., smoothness, low number of critical points).

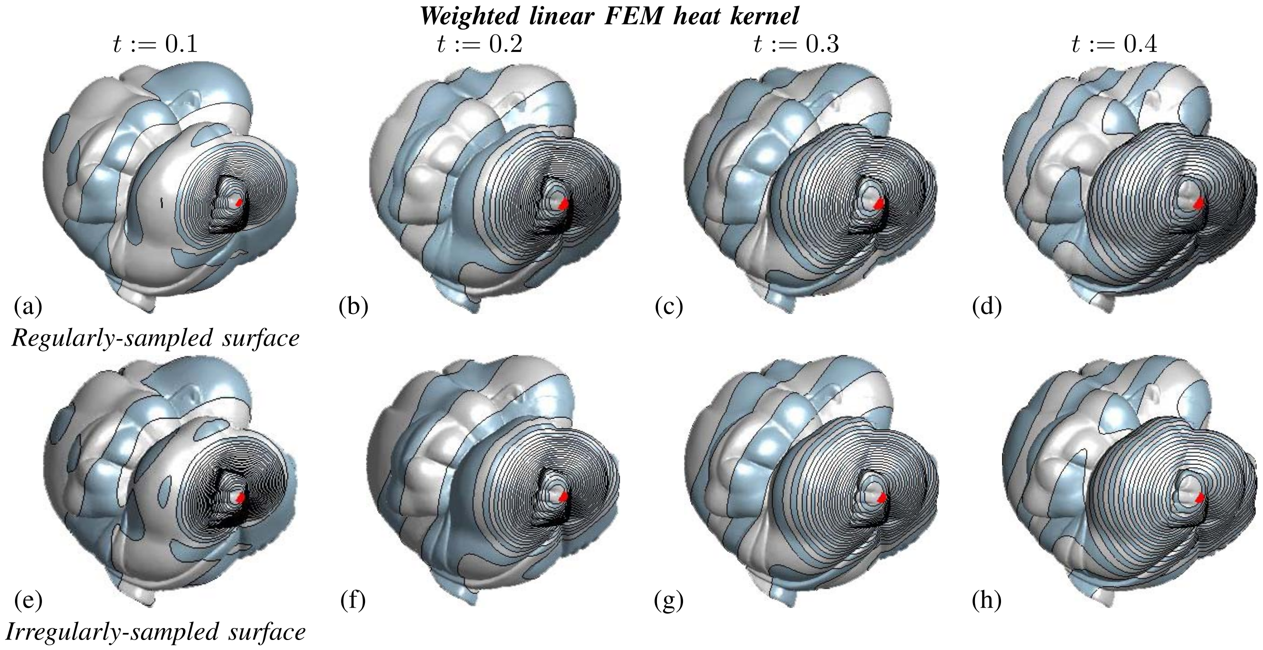

Fig. 5. With reference to Figure 4, the example shows the canonic basis functions associated to the weighted linear FEM heat kernel, which provides smoother results at every scale, on both (a-d) the regularly- and (e-h) irregularly-sampled surface.

For instance, the canonic basis $\mathcal{B} := \{\mathbf{e}_i\}_{i=1}^n$ of $\mathbb{R}^n$, $(\mathbf{e}_i)_j := \delta_{ij}$, is not useful for shape analysis; in fact, the corresponding maps on $\mathcal{M}$ have abrupt variations from one to zero. Indeed, it is interesting to find a counterpart of this basis in $\mathcal{F}(\mathcal{M})$, which retains its main properties such as localization and linear independence, has a smooth behavior on $\mathcal{M}$, and is significative for signal approximation (Figure 4, 5).

Since the feature map Φ_t is bijective on the space $\mathcal{F}(\mathcal{M})$, any basis of $\mathbb{R}^n$ is mapped to a basis of $\mathcal{F}(\mathcal{M})$; in particular, as *canonic basis* of the feature space $\mathcal{F}(\mathcal{M})$ we refer to the set $\mathcal{C} := \{\Phi_t(\mathbf{e}_i) := XD_tX^TB\mathbf{e}_i\}_{i=1}^n$ corresponding to the (canonic) basis of $\mathbb{R}^n$. The function $\Phi_t(\mathbf{e}_i)$, $t \in \mathbb{R}$, is a multi-scale and smooth approximation of the map that attains the value one at the vertex $\mathbf{p}_i$ and zero at all the other vetices of $\mathcal{M}$. This basis will guide the feature-driven approximation of scalar functions, which is described in Section III-C.

If the surface $\mathcal{M}$ has a regular sampling density, then the canonic basis function $\Phi_t^I(\mathbf{e}_i) := XD_tX^T\mathbf{e}_i$ and $\Phi_t(\mathbf{e}_i) := XD_tX^TB\mathbf{e}_i$, which are provided by the linear and weighted linear FEM heat kernel, have a similar behavior at the same scales and in a neighbor of the point $\mathbf{p}_i$ (Figure 4(a-d), 5(a-d)). Irregularly-sampled patches on $\mathcal{M}$ generally affects the smoothness of $\Phi_t(\mathbf{e}_i)$ at smaller scales (Figure 4(e,f)); increasing t improves the smoothness of $\Phi_t(\mathbf{e}_i)$ in terms of regularity of the level sets and of a low number of critical points (Figure 4(g,h)). On the contrary, the smoothness of $\Phi_t(\mathbf{e}_i)$ is guaranteed through all the scales in spite of the discretization quality of $\mathcal{M}$ (Figure 5(e-h)). In all the examples, the canonic basis function $\Phi_t(\mathbf{e}_i)$ is distributed on the whole surface and concentrated mainly in a neighbor of its center $\mathbf{p}_i$. This is due to the fact that each basis $\Phi_t(\mathbf{e}_i)$ is a linear combination of the Laplacian eigenfunctions, which have a global support (i.e., they vanish on sets of null measure). Finally, the canonic basis provides a multi-scale alternative to the geometry-aware functions [42], which involves the linear FEM or any other discretization of the Laplace-Beltrami operator and is not restricted to constant weights.

To express a scalar function $f : \mathcal{M} \rightarrow \mathbb{R}$ as a linear combination of the canonic basis in $\mathcal{F}(\mathcal{M})$, let us identify f with the array $\mathbf{f} := (f(\mathbf{p}_i))_{i=1}^n$ and let $\alpha := (\alpha_i)_{i=1}^n$ be the unknown vector such that $\mathbf{f} = \sum_{i=1}^n \alpha_i \Phi_t(\mathbf{e}_i)$. Applying the linear operator Φ_t^{-1} to both members of the previous identity, we get that

$$\Phi_t^{-1}(\mathbf{f}) := \Phi_t^{-1}\left(\sum_{i=1}^n \alpha_i \Phi_t(\mathbf{e}_i)\right) = \sum_{i=1}^n \alpha_i \mathbf{e}_i = \alpha,$$

and the vector α is the solution to the sparse linear system $K_t\alpha = \mathbf{f}$. If $B := I$, then $\alpha = XD_t^{-1}X^T\mathbf{f}$ is evaluated in linear time as matrix-vector product.

3.3 Projection operator and feature-driven approximation of scalar functions

Given a map $f : \mathcal{M} \rightarrow \mathbb{R}$, the family $\{\Phi_t(f)\}_t$ is computed by treating all the f-values with the same degree of importance. In this context, it is interesting to adapt the approximation in such a way that specific values of f, which are related to its global behavior, are preserved at various scales. To this end, let us consider a set $\mathcal{A} \subseteq \{1, \dots, n\}$ of indices such that $\{f(\mathbf{p}_i)\}_{i\in\mathcal{A}}$ is the corresponding set of *feature values* of f; i.e., a set of values that characterize the behavior of f. These points $\{\mathbf{p}_i\}_{i\in\mathcal{A}}$ might be identified by using the simplification of the critical points of f [1], [13], [33] or a clustering

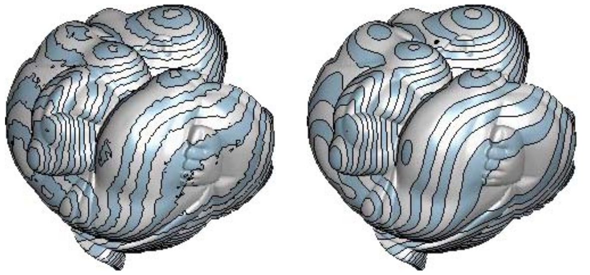

(a) (b)

Fig. 6. With reference to the map f shown in Figure 3(a), level sets of the projection $\Phi_t(pr_{\mathcal{A}}(f))$, where $\mathcal{A}$ is the set of (a) all the critical points of f and (b) those (i.e., the 42% of the input ones) preserved by the δ-simplifcation [33]. This last choice provides smoother results through a lower number of constraints.

of the f-values. In $\mathcal{F}(\mathcal{M})$, the *projection operator with respect to* $\mathcal{A}$ is defined as

$$
\begin{aligned}
pr_{\mathcal{A}}: \quad & \mathcal{F}(\mathcal{M}) \longrightarrow \mathcal{F}(\mathcal{M}) \\
& f \mapsto (pr_{\mathcal{A}}(f))(\mathbf{p}_i) := \begin{cases} f(\mathbf{p}_i) & i \in \mathcal{A}, \\ 0 & \text{else.} \end{cases}
\end{aligned}
$$

Then, the smooth approximation of f constrained to $\mathcal{A}$ is $\Phi_t(pr_{\mathcal{A}}(f)) \in \mathcal{F}(\mathcal{M})$. Once the set $\mathcal{A}$ has been chosen, we define the *global component of* f as

$$
f_{glob} := \Phi_t(pr_{\mathcal{A}}(f)) := \sum_{i \in \mathcal{A}} f(\mathbf{p}_i)\Phi_t(\mathbf{e}_i).
$$

Due to the linearity of Φ_t, the f_{glob}-values on $\mathcal{M}$ are computed as $\mathbf{f}_{glob} = K_t\tilde{\mathbf{f}}$, where the vector $\tilde{\mathbf{f}} \in \mathbb{R}$ has entry $\tilde{\mathbf{f}}^{(i)} := f(\mathbf{p}_i)$, if $i \in \mathcal{A}$, and zero otherwise. Indeed, it is not necessary to explicitly compute the canonic basis functions $\{\Phi_t(\mathbf{e}_i)\}_{i \in \mathcal{A}}$. Finally, the difference between f (or $\Phi_t(f)$) and f_{glob} is the corresponding *local component*. Note that the local component of $\Phi_t(f)$ is given by $\Phi_t(pr_{\mathcal{A}^C}(f))$; i.e., the image of the projection $pr_{\mathcal{A}^C}(f)$ of f through Φ_t using the complementary set $\mathcal{A}^C$.

For smoothing a noisy map f, as $\mathcal{A}$ we consider the set of critical points of f with a high persistence, which are computed according to [1], [13], [33]. In this case, we expect that f_{loc} codes the local feature or noise of f and f_{glob} identifies its global behavior. While traditional approaches to function approximation are mainly driven by a numerical error estimation, combining the critical points with the previous scheme provides a natural *feature-driven approximation* (Figure 6, 7). In fact, the critical points usually represent a relevant information about the phenomena coded by f. Using the critical points of f with a high persistence generally provides results that are smoother than using all the critical points (Figure 6) and affects neither the approximation accuracy ϵ_∞ nor the number of critical points (Figure 7).

4. Conclusions and future work

In this paper, a new weighted linear FEM discretization of the heat kernel has been exploited to define the feature spaces associated to 3D shapes, which have been characterized in the context of signal approximation. The proposed kernel is intrinsically scale invariant, stable to irregular sampling densities of the input surface, and suitable for shape comparison with heat kernel shape signatures and related descriptors. The usefulness of our results for shape retrieval through the heat kernel signature and related descriptors has been recently presented in [4]. As main future work, we plan to investigate the connection between the proposed approach, the diffusion wavelets [6], and the Tikhonov regularization.

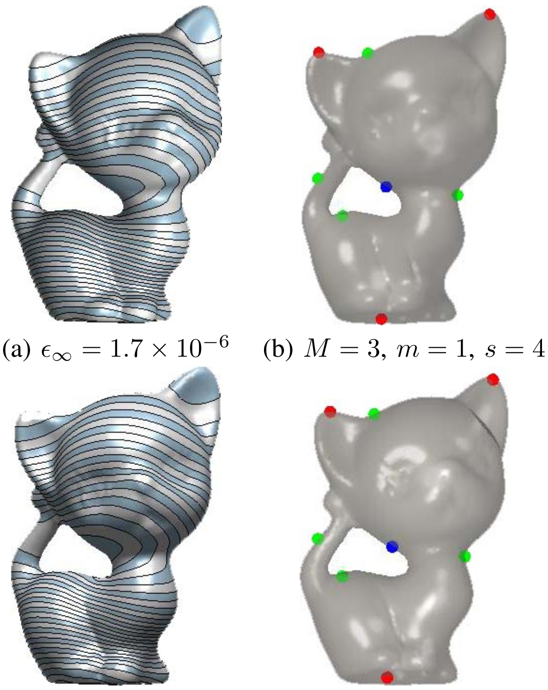

(a) $\epsilon_\infty = 1.7 \times 10^{-6}$ (b) $M = 3,\ m = 1,\ s = 4$

(c) $\epsilon_\infty = 1.9 \times 10^{-6}$ (d) $M = 3,\ m = 1,\ s = 4$

Fig. 7. With reference to Figure 2, level sets and critical points of feature-driven approximation with 100 Laplacian eigenfunctions and constrained to the f-values at (a,b) all (i.e., 6406) the critical points or (c,d) those (i.e., the 11% of the input ones) preserved by the simplification [33].

Acknowledgments

Special thanks are given to the anonymous reviewers for their comments. The activities of Giuseppe Patanè have been partially funded by the GNCS-INdAM "F. Severi", "Young Researcher Programme". This work has been partially supported by the FOCUS K3D Coordination Action. Models are courtesy of the AIM@SHAPE Repository.

References

[1] T. Banchoff. Critical points and curvature for embedded polyhedra. *Journal of Differential Geometry*, 1:245–256, 1967.

[2] M. Belkin and P. Niyogi. Laplacian eigenmaps for dimensionality reduction and data representation. *Neural Computations*, 15(6):1373–1396, 2003.

[3] 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*, To appear, 2009.

[4] 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 large-scale shape retrieval benchmark. *Eurographics Workshop on 3D Object Retrieval*, To appear, 2010.

[5] F. R. K. Chung. *Spectral graph theory*. American Mathematical Society, 1997.

[6] R. R. Coifman and S. Lafon. Diffusion maps. *Applied and Computational Harmonic Analysis*, 21(1):5–30, 2006.

[7] F. de Goes, S. Goldenstein, and L. Velho. A hierarchical segmentation of articulated bodies. *Computer Graphics Forum*, 27(5):1349–1356, 2008.

[8] M. Desbrun, M. Meyer, P. Schröder, and A. H. Barr. Implicit fairing of irregular meshes using diffusion and curvature flow. In *ACM Siggraph 1999*, pages 317–324, 1999.

[9] J. Díaz, J. Petit, and M. Serna. A survey of graph layout problems. *ACM Computing Surveys*, 34(3):313–356, 2002.

[10] S. Dong, P.-T. Bremer, M. Garland, V. Pascucci, and J. C. Hart. Spectral surface quadrangulation. *ACM Siggraph 2006*, pages 1057–1066, 2006.

[11] S. Dong, S. Kircher, and M. Garland. Harmonic functions for quadrilateral remeshing of arbitrary manifolds. *Computer Aided Geometric Design*, 22(5):392–423, 2005.

[12] R. Dyer, R. H. Zhang, T. Moeller, and A. Clements. An investigation of the spectral robustness of mesh laplacians. 2007.

[13] H. Edelsbrunner, D. Morozov, and V. Pascucci. Persistence-sensitive simplification functions on 2-manifolds. In *Proc. of the Symposium on Computational Geometry*, pages 127–134. ACM, 2006.

[14] K. Gebal, J. Andreas Bærentzen, H. Aanæs, and R. Larsen. Shape analysis using the auto diffusion function. *Computer Graphics Forum*, 28(5):1405–1413, 2009.

[15] C. Gotsman, X. Gu, and A. Sheffer. Fundamentals of spherical parameterization for 3D meshes. In *ACM Siggraph 2003*, pages 358–363, 2003.

[16] V. Jain and H. Zhang. A spectral approach to shape-based retrieval of articulated 3D models. *Computer Aided Design*, 39:398–407, 2007.

[17] K. Polthier K. Hildebrandt and M. Wardetzky. On the convergence of metric and geometric properties of polyhedral surfaces. *Geometria Dedicata*, pages 89–112, 2006.

[18] Z. Karni and C. Gotsman. Spectral compression of mesh geometry. In *ACM Siggraph 2000*, pages 279–286, 2000.

[19] B. Kim and J. Rossignac. Geofilter: Geometric selection of mesh filter parameters. *Computer Graphics Forum*, 24(3):295–302, 2005.

[20] 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.

[21] Y. Koren. On spectral graph drawing. volume 2697, pages 496–508. Lecture Notes in Computer Science, 2003.

[22] S. Lafon, Y. Keller, and R. R. Coifman. Data fusion and multicue data matching by diffusion maps. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 28:1784–1797, 2006.

[23] B. Levy. Laplace-Beltrami eigenfunctions: towards an algorithm that understands geometry. In *Proc. of the IEEE Shape Modeling and Applications*, page 13, 2006.

[24] R. Liu and H. Zhang. Mesh segmentation via spectral embedding and contour analysis. *Eurographics 2007*, 26:385–394, 2007.

[25] C. Luo, I. Safa, and Y. Wang. Approximating gradients for meshes and point clouds via diffusion metric. *Computer Graphics Forum*, 28(5):1497–1508, 2009.

[26] A. M. Bronstein M. M. Bronstein. Analysis of diffusion geometry methods for shape recognition. *IEEE Trans. Pattern Analysis and Machine Intelligence*, Submitted, 2010.

[27] J.-L. Mallet. Discrete smooth interpolation. *ACM Transactions on Graphics*, 8(2):121–144, 1989.

[28] F. Mémoli. Spectral Gromov-Wasserstein distances for shape matching. In *Workshop on Non-Rigid Shape Analysis and Deformable Image Alignment*, october 2009.

[29] A. Nealen, T. Igarashi, O. Sorkine, and M. Alexa. Laplacian mesh optimization. In *Proc. of Computer graphics and interactive techniques*, pages 381–389, 2006.

[30] 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.

[31] R. Ohbuchi, A. Mukaiyama, and S. Takahashi. A frequency-domain approach to watermarking 3D shapes. *Computer Graphics Forum*, 21(3), 2002.

[32] 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.

[33] G. Patanè and B. Falcidieno. Computing smooth approximations of scalar functions with constraints. *Computer & Graphics*, 33(3):399 – 413, 2009.

[34] N. Peinecke, F.-E. Wolter, and M. Reuter. Laplace spectra as fingerprints for image recognition. *Computer-Aided Design*, 39:460–476, 2007.

[35] U. Pinkall and K. Polthier. Computing discrete minimal surfaces and their conjugates. *Experimental Mathematics*, 2(1):15–36, 1993.

[36] 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.

[37] M. Reuter, F. Wolter, M. Shenton, and M. Niethammer. Laplace-beltrami eigenvalues and topological features of eigenfunctions for statistical shape analysis. *Computer Aided Design*, 41(10):739–755, 10 2009.

[38] M. Reuter, F.-E. Wolter, and N. Peinecke. Laplace-spectra as fingerprints for shape matching. In *Proc. of the Symp. on Solid and Physical Modeling*, pages 101–106. ACM, 2005.

[39] 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.

[40] R. M. Rustamov. Laplace-beltrami eigenfunctions for deformation invariant shape representation. In *Proc. of the Symposium on Geometry processing*, pages 225–233, 2007.

[41] O. Sorkine. Differential representations for mesh processing. *Computer Graphics Forum*, 25(4):789–807, 2006.

[42] 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.

[43] 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.

[44] 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.

[45] G. Taubin. A signal processing approach to fair surface design. In *ACM Siggraph 1995*, pages 351–358, 1995.

[46] B. Vallet and B. Levy. Manifold harmonics. *Computer Graphics Forum 27(2)*, 2008.

[47] G. Xu. Discrete Laplace-Beltrami operators and their convergence. *Computer Aided Geometric Design*, 8(21):398–407, 2007.

[48] 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.

[49] 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.

[50] 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.

[51] 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.

[52] 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.