

# Volumetric Heat Kernel: Padé-Chebyshev Approximation, Convergence, and Computation

Giuseppe Patané, CNR-IMATI - Italy  
patane@ge.imati.cnr.it

---

## Abstract

This paper proposes an accurate and computationally efficient solver of the heat equation  $(\partial_t + \Delta)F(\cdot, t) = 0$ ,  $F(\cdot, 0) = f$ , on a volumetric domain, through the  $(r, r)$ -degree Padé-Chebyshev rational approximation of the exponential representation  $F(\cdot, t) = \exp(-t\Delta)f$  of the solution. To this end, the heat diffusion problem is converted to a set of  $r$  differential equations, which involve only the Laplace-Beltrami operator, and whose solution converges to  $F(\cdot, t)$ , as  $r \rightarrow +\infty$ . The discrete heat equation is equivalent to  $r$  sparse, symmetric linear systems and is independent of the volume discretization as a tetrahedral mesh or a regular grid, the evaluation of the Laplacian spectrum, and the selection of a subset of eigenpairs. Our approach has a super-linear computational cost, is free of user-defined parameters, and has an approximation accuracy lower than  $10^{-r}$ . Finally, we propose a simple criterion to select the time value that provides the best compromise between approximation accuracy and smoothness of the solution.

*Keywords:* Volumetric heat kernel, diffusion geometry, Laplacian spectrum, shape analysis

---

## 1. Introduction

The heat kernel plays a central role in several applications, such as surface [3, 16] and image [34, 40] smoothing, shape segmentation [12] and comparison [5, 6, 14, 25, 26, 35]. Furthermore, the wavelet operator [17], the geodesic [11] and diffusion [4, 10, 27] distances have been recently rewritten in terms of the heat kernel. Among its main properties, we mention the intrinsic and multi-scale encoding of the input shape, the invariance to isometries, the shape-awareness, the robustness to noise and tessellation.

In several applications, volumetric representations and descriptors are more suited than a two-dimensional manifold to model the shape invariance under rigid and elastic transformations. Furthermore, tetrahedral meshes are efficiently generated from surfaces [2, 33] and are a standard volumetric representation for the discretization of differential equations. Due to the high computational cost for the solution to the heat equation, previous work has been mainly focused on the diffusion kernel and distance on surfaces rather than on volumes.

Given the complexity of volumetric computation, several alternatives to the heat kernel were proposed in the literature. FEM discretizations [1] of the heat equation tessellate the volume with a voxel grid or cuboid voxels [30] and apply a 6-neighborhood stencil [23,

29] or a geometry-driven approximation field [22, 36]. These approximations provide a low accuracy of the solution in a neighbor of the volume boundary, which is generally represented as a triangle mesh. Even though multi-resolution prolongation operators [39] and Chebyshev polynomials [27, 28] can be extended to volumes, they have not been applied to the computation of the volumetric heat kernel or to the selection of the optimal time value. Additionally, the multi-resolution simplification of the input volume is time-consuming and the selection of the volume resolution with respect to the expected approximation accuracy is generally guided by heuristics. Further approaches extend the solution to the heat equation computed on the input surface to its interior through barycentric coordinates or a non-linear approximation, as done for the Laplacian [31, 32] and harmonic [21, 24] maps. Note that these methods do not intend to approximate the heat kernel quantitatively, but provide alternative approaches that qualitatively behave like the heat kernel on volumes.

*Overview and contribution.* We propose an accurate and computationally efficient solver of the heat equation  $(\partial_t + \Delta)F(\cdot, t) = 0$ ,  $F(\cdot, 0) = f$ , on a closed and connected manifold  $\mathcal{M}$  of  $\mathbb{R}^3$ , such that its boundary  $\partial\mathcal{M}$  is a smooth and closed two-dimensional manifold. We also introduce a simple criterion to select the time value

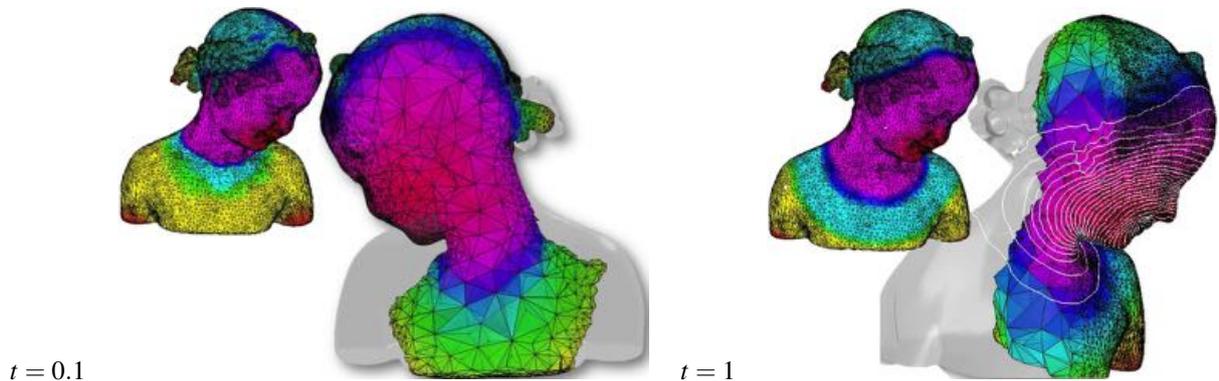


Figure 1: **Volumetric heat kernel.** Color map of the solution to the volumetric heat equation at two time values and computed with the Padé-Chebyshev approximation of degree  $r = 7$ ; the initial condition takes value 1 at a point of the lips and 0 at the other vertices of the tetrahedralization. The color map varies the hue component of the hue-saturation-value color model; the colors begin with red, pass through yellow, green, cyan, blue, and magenta, and return to red. At scale  $t = 1$ , the level-sets on the volume boundary correspond to iso-values uniformly sampled in the range of the solution restricted to the points of the volume boundary.

(or scale) that provides the best compromise between approximation accuracy and smoothness of the solution.

The idea behind our approach (Sect. 2) is to apply the  $(r, r)$ -degree Padé-Chebyshev rational approximation to the exponential representation  $F(\cdot, t) = \exp(-t\Delta)f$  of the solution to the heat equation. Then, the diffusion problem is converted to a set of  $r$  differential equations, which involve only the Laplace-Beltrami operator, and the resulting solution converges to  $F(\cdot, t)$ , as  $r \rightarrow +\infty$ . Through the proposed approach, the solution to the heat equation is approximated in a low-dimensional space generated by  $r+1$  functions, which are induced by the input volume, the initial condition  $f$ , and the selected time value. Furthermore, the approximation accuracy is lower than  $10^{-r}$  (e.g.,  $r = 5, 7$ ). In comparison, the Laplacian eigenfunctions only encode the domain geometry and it is difficult to select the number of eigenpairs necessary to achieve a given approximation of  $F(\cdot, t)$  with respect to  $t$  and  $f$ .

While a discretization of the heat kernel on a voxel grid is accurate enough for the evaluation of diffusion descriptors [23, 29], which are quantized and clustered in bags-of-features, we focus on the computation of the heat kernel on tetrahedral meshes (Fig. 1). Our discretization (Sect. 3) is equivalent to a set of  $r$  sparse, symmetric linear systems and is applied to any representation of the input domain and of the Laplace-Beltrami operator. Furthermore, it properly encodes the local and global features in the heat kernel and bypasses the computation of the Laplacian spectrum. For a given time value, the overall computational cost of the  $r$ -degree Padé-Chebyshev rational polynomial is  $O(rn)$ , where  $n$  is the number of volume vertices. Indeed, our approx-

imation is competitive with respect to multi-resolutive simplification/prolongation operators, the Euler backward method, and the truncated spectral approximation.

As main novelties with respect to previous work [27], we apply the Padé-Chebyshev approximation to the more complex case of the heat kernel on volumes, also addressing the convergence of the approximation scheme and the selection of the time value.

For our experiments (Sect. 4), we consider volumetric diffusion smoothing, which is typically applied to thin film evolution [19], to the analysis of multi-material volume grids and their interfaces [20], and to volumetric shape deformation [22]. Other possible applications, which are not addressed in this paper, include volume-based approximation and the evaluation of volumetric descriptors.

## 2. Volumetric heat equation

Let us consider the heat equation  $(\partial_t + \Delta)F(\cdot, t) = 0$ ,  $F(\cdot, 0) = f$ , on a closed, connected manifold  $\mathcal{M}$  of  $\mathbb{R}^3$ , with  $f: \mathcal{M} \rightarrow \mathbb{R}$  and  $\partial\mathcal{M}$  smooth, closed two-dimensional boundary of  $\mathcal{M}$ . Then, the solution  $F(\mathbf{p}, t) = K_t(\mathbf{p}, \cdot) \star f$  is the convolution between the heat kernel  $K_t(\mathbf{p}, \mathbf{q}) := (4\pi t)^{-3/2} \exp(-\|\mathbf{p} - \mathbf{q}\|_2^2 / 4t)$  and  $f$ .

Our approach applies the Padé-Chebyshev rational approximation to the exponential representation  $F(\cdot, t) = \exp(-t\Delta)f$  of the solution to the heat equation. According to [15], on  $\mathbb{R}^+$  the best  $(r, r)$ -degree rational polynomial approximation of  $\exp(-x)$  with respect to the  $\mathcal{L}_\infty$  norm is  $c_{rr}(x) = \alpha_0 + \sum_{i=1}^r \alpha_i (x - \theta_i)^{-1}$ , with poles  $\{\theta_i\}_{i=1}^r$  and coefficients  $\{\alpha_i\}_{i=1}^r$ . These values are

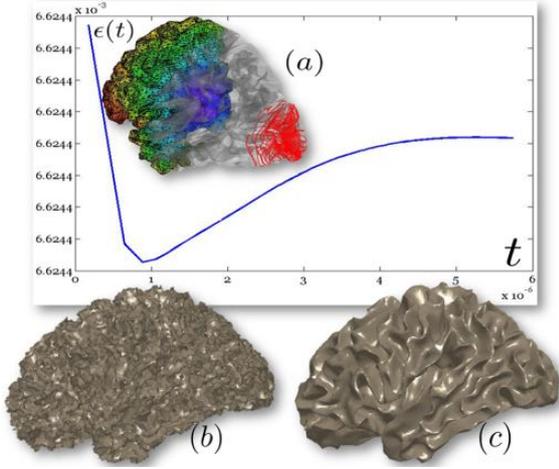


Figure 2: **Volumetric diffusion smoothing of noisy data.** (a) Input tetrahedral mesh and L-curve of the approximation accuracy (y-axis) versus the solution smoothness (x-axis). (b) Data set achieved by adding a Gaussian noise to the tetrahedra mesh (a). (c) Diffusion smoothing computed with the Padé-Chebyshev approximation of degree  $r = 7$ . The  $\ell_\infty$  error between the volumetric ground-truth (a) and the smooth approximation of (b) is lower than 1% for the Padé-Chebyshev method (c) and varies from 12% ( $k = 100$ ) up to 13% ( $k = 1K$ ) for the truncated spectral approximation.

precomputed for any degree through standard numerical routines; for more details, we refer the reader to [8]. Indicating with  $\text{id}(\cdot)$  the identity operator, the function

$$\begin{aligned} F(\cdot, t) &= \exp(-t\Delta)f \approx \alpha_0 f - \sum_{i=1}^r \alpha_i (\Delta + \theta_i \text{id})^{-1} f \\ &= \alpha_0 f + \sum_{i=1}^r \alpha_i g_i, \quad (t\Delta + \theta_i \text{id})g_i = -f, \end{aligned} \quad (1)$$

is approximated by a linear combination of the solutions to  $r$  equations induced by the Laplace-Beltrami operator. The resulting approximation of  $F(\cdot, t)$  belongs to the linear space  $\mathcal{H}$  generated by  $f$  and  $\{g_i\}_{i=1}^r$ , which depend on the input volume, the initial condition  $f$ , and the selected time value  $t$ . In comparison, the Laplacian eigenfunctions  $\{(\lambda_n, \phi_n)\}_{n=0}^{+\infty}$ ,  $\Delta\phi_n = \lambda_n\phi_n$ , encode only the domain geometry and it is difficult to select the number  $k$  of eigenpairs that are necessary to achieve an accurate approximation of  $F(\cdot, t)$  through the truncated spectral representation  $F(\cdot, t) \approx \sum_{n=1}^k \exp(-\lambda_n t) \langle f, \phi_n \rangle_2 \phi_n$ . Furthermore, a larger number of eigenpairs is necessary to accurately recover the solution at small time values.

*Convergence of the approximation.* Introducing the approximate solution  $F_r(\cdot, t) := \sum_{n=0}^{+\infty} c_{rr}(\lambda_n) \langle f, \phi_n \rangle_2 \phi_n$  to the (volumetric) heat equation induced by the  $r$ -degree Padé-Chebyshev polynomial  $c_{rr}$ , we show that the sequence  $(F_r(\cdot, t))_{r=0}^{+\infty}$  converges to  $F(\cdot, t)$ . First of all,

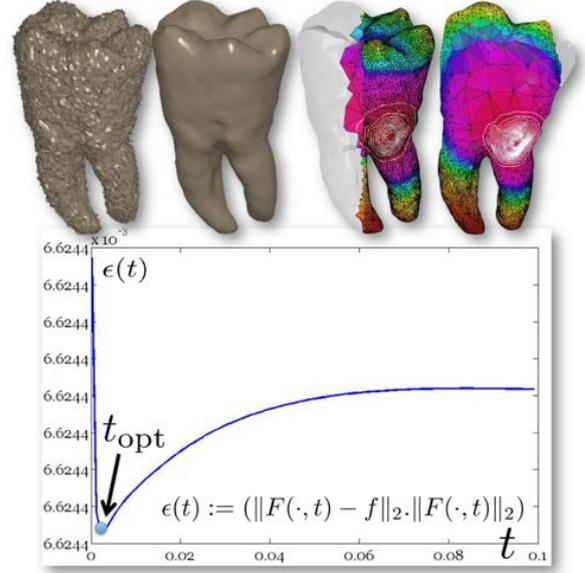


Figure 3: **L-curve and volumetric heat kernel.** Selection of the optimal scale and corresponding volumetric diffusion smoothing (upper part, right), Padé-Chebyshev approximation of degree  $r = 7$  on the noisy volumetric model of the teeth (upper part, left).

we notice that the approximation  $F_r(\cdot, t)$  is well-posed; in fact,  $\|c_{rr}\|_\infty \leq 1$  and  $\|F_r(\cdot, t)\|_2 \leq \|f\|_2$ . According to [38], the  $\mathcal{L}_\infty$  error between the exponential map and its rational polynomial approximation is bounded by the uniform rational Chebyshev constant  $\sigma_{rr}$ , which is independent of the evaluation point, and lower than  $10^{-r}$ . Applying the upper bound

$$\begin{aligned} \|F_r(\cdot, t) - F(\cdot, t)\|_2^2 &\leq \|c_{rr}(\cdot) - \exp(-\cdot)\|_\infty^2 \sum_{n=0}^{+\infty} |\langle f, \phi_n \rangle_2|^2 \\ &\leq \sigma_{rr}^2 \sum_{n=0}^{+\infty} |\langle f, \phi_n \rangle_2|^2 \leq 10^{-2r} \|f\|_2^2, \end{aligned}$$

we deduce that  $\lim_{r \rightarrow +\infty} F_r(\cdot, t) = F(\cdot, t)$ .

While the selection of a fixed number of eigenpairs does not allow us to estimate the resulting approximation accuracy, the projection of  $F(\cdot, t)$  on the linear space generated by  $\{f, \phi_1, \dots, \phi_r\}$  guarantees an accuracy lower than  $10^{-r}$ . Finally, this approximation is stable to a perturbation  $f + e$  of the initial condition; in fact, the variation of the corresponding solutions  $\tilde{F}_r(\cdot, t)$ ,  $F_r(\cdot, t)$  is bounded by the norm of the perturbation; i.e.,  $\|\tilde{F}_r(\cdot, t) - F_r(\cdot, t)\|_2 \leq \|c_{rr}\|_\infty \|e\|_2 \leq \|e\|_2$ .

*Optimal time value.* As optimal time value, we select the scale that provides a small residual error

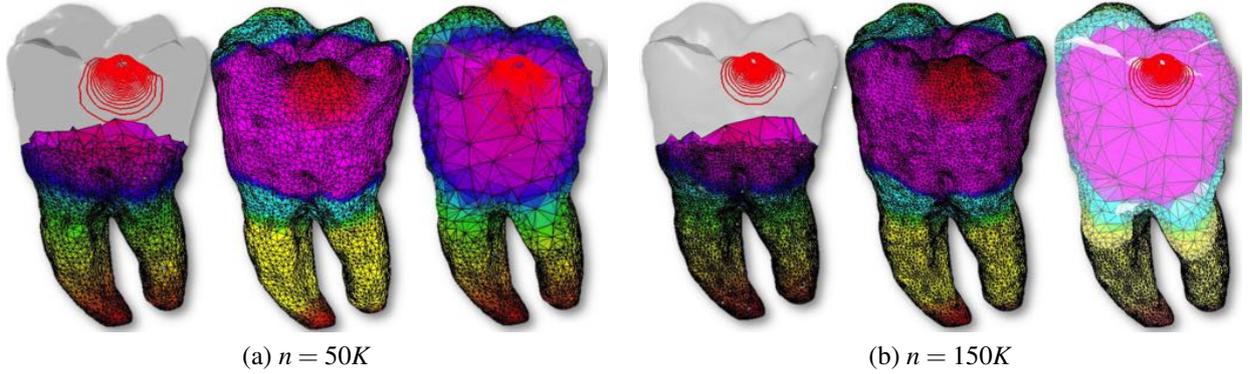


Figure 4: **Robustness to sampling.** Robustness of the Padé-Chebyshev approximation ( $r = 7$ ) of the volumetric heat kernel with respect to volume sampling;  $n$  is the number of vertices of the input tetrahedral mesh. The  $\ell_\infty$  discrepancy between the maps in (a) and (b) is lower than 1.4% ( $r = 5$ ).

$\|F(\cdot, t) - f\|_2$  and a low energy  $\|F(\cdot, t)\|_2$ , which controls the solution smoothness. Through the Laplacian spectrum  $\{(\lambda_n, \phi_n)\}_{n=0}^{+\infty}$ , the orthonormality of the Laplacian eigenfunctions, and the spectral representation  $f = \sum_{n=0}^{+\infty} \langle f, \phi_n \rangle_2 \phi_n$  of the initial condition, we rewrite these terms as

$$\begin{cases} \|F(\cdot, t) - f\|_2^2 = \sum_{n=0}^{+\infty} |1 - \exp(-2\lambda_n t)|^2 |\langle f, \phi_n \rangle_2|^2, \\ \|F(\cdot, t)\|_2^2 = \sum_{n=0}^{+\infty} \exp(-2\lambda_n t) |\langle f, \phi_n \rangle_2|^2; \end{cases}$$

indeed, the residual and penalty terms are increasing and decreasing maps with respect to  $t$ , respectively. If  $t$  tends to zero, then the residual becomes null and the energy converges to  $\|f\|_2$ . If  $t$  becomes large, then the residual tends to  $|\langle f, \phi_0 \rangle_2|$  and the solution norm converges to  $(\|f\|_2^2 - |\langle f, \phi_0 \rangle_2|^2)^{1/2}$ . According to these properties, the plot ( $L$ -curve) of  $\varepsilon(t) := (\|F(\cdot, t) - f\|_2, \|F(\cdot, t)\|_2)$  is  $L$ -shaped [18] and its minimum provides the optimal time value; i.e., the best compromise between approximation accuracy and smoothness. For the computation of the optimal time value (Figs. 2, 3), we apply the corner detection based on cubic B-splines approximation [18]; alternatives are the evaluation of the curvature of the graph of  $\varepsilon(t)$  or its adaptive pruning [18].

### 3. Computational aspects

Representing the input domain as a tetrahedral mesh  $\mathcal{M}$  with vertices  $\mathcal{P} := \{\mathbf{p}_i\}_{i=1}^n$ , any scalar function on  $\mathcal{M}$  is defined as  $g = \sum_{j=1}^n g_j \phi_j$ , where  $g_j := g(\mathbf{p}_j)$  and  $\phi_j$  is the piecewise linear basis function that takes value 1 at  $\mathbf{p}_j$  and 0 otherwise. To discretize the heat equation on  $\mathcal{M}$ , we recall that [22, 36]

$$\langle \Delta g, \phi_i \rangle_2 = \frac{1}{V_i} \sum_{j \in N(i)} \left( \frac{1}{6} \sum_{k=1}^n l_k \cot \alpha_k \right) (g_i - g_j), \quad (2)$$

where  $V_i$  is the tetrahedral volume at  $i$ ,  $N(i)$  is the set of vertices incident to  $i$ ,  $l_k$  is the length of the edge to which  $(i, j)$  is opposite. Then, the weak formulation  $\langle \nabla g, \nabla \phi_i \rangle_2 + \theta \langle g, \phi_i \rangle_2 = \langle f, \phi_i \rangle_2$  of Eq. (1) is rewritten as  $(\tilde{\mathbf{L}} + \theta \mathbf{I})\mathbf{g} = \mathbf{f}$ ,  $\mathbf{f} := (f(\mathbf{p}_i))_{i=1}^n$ . Here, the Laplacian matrix  $\tilde{\mathbf{L}} := \mathbf{B}^{-1}\mathbf{L}$  (c.f., Eq. (2)) is defined as the product between the diagonal matrix  $\mathbf{B}$ , which encodes the tetrahedral volume  $V_i$  at each vertex  $i$ , and  $\mathbf{L}$  is the Laplacian matrix with entries  $L(i, j) := w_{ij} := \frac{1}{6} \sum_{k=1}^n l_k \cot \alpha_k$  for each edge  $(i, j)$ ,  $L(i, i) := -\sum_{j \in N(i)} w_{ij}$ , and zero otherwise. Indeed, the solution  $\mathbf{F}(t) = \alpha_0 \mathbf{f} + \sum_{i=1}^r \mathbf{g}_i$  to the heat equation  $(\partial_t + \tilde{\mathbf{L}})\mathbf{F}(t) = \mathbf{0}$ ,  $\mathbf{F}(0) = \mathbf{f}$ , is the sum of the solutions of  $r$  sparse linear systems

$$(t\tilde{\mathbf{L}} + \theta_i \mathbf{B})\mathbf{g}_i = -\alpha_i \mathbf{B}\mathbf{f}, \quad i = 1, \dots, r. \quad (3)$$

Since the  $r$ -degree Padé-Chebyshev rational polynomial of the exponential map is a-priori known, for a given  $t$  the vectors in Eq. (3) are calculated as a minimum norm residual solution [15]. Our solver is free of user-defined parameters, and works with sparse, well-conditioned matrices [27]. It also regularizes noisy data while preserving local details (Fig. 2), which are over-smoothed by the truncated spectral approximation, and provides an efficient way to estimate the optimal scale (Fig. 3) by evaluating the map  $\varepsilon(\cdot)$  for several values of  $t$ .

### 4. Discussion and future work

We discuss the stability of the proposed computation to volume sampling and noise through a comparison with previous work on surfaces and adapted to the volumetric heat kernel. For the paper examples, the values of the initial condition  $f$  and the solution  $F(\cdot, t)$  have been normalized in  $[0, 1]$ . The level-sets on  $\partial\mathcal{M}$  are associated with iso-values uniformly sampled in the range

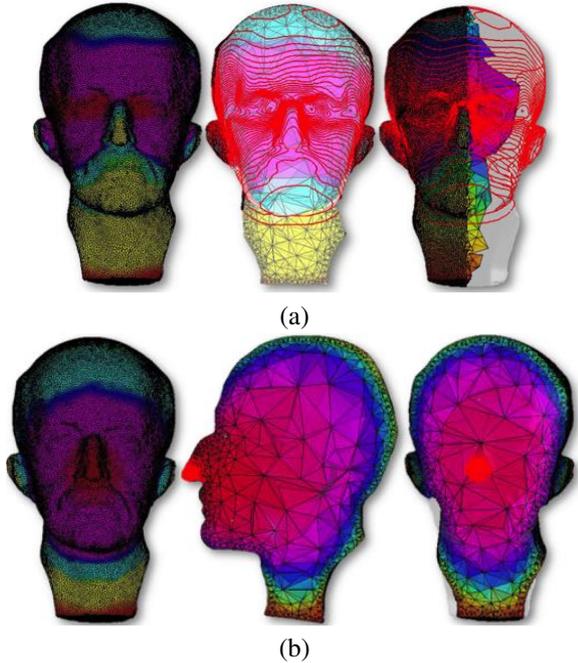


Figure 5: **Diffusion process with different initial conditions.** Diffusion process ( $r = 7$ ) generated by a seed point belonging to (a) the internal part and (b) the volume boundary.

of  $f$  or to the restriction of  $F(\cdot, t)$  on the boundary of  $\mathcal{M}$ ; in this case, the color map varies from blue (minimum) to red (maximum). For the visualization of  $F(\cdot, t)$  on a tetrahedral mesh, the color map varies the hue component of the hue-saturation-value color model; the colors begin with red, pass through yellow, green, cyan, blue, and magenta, and return to red. Finally, the initial condition  $F(\cdot, 0) = f$  and the solution  $F(\cdot, t)$  of the corresponding diffusion equation are always defined on the input tetrahedral mesh. In a similar way, the noise of the synthetic examples is added to the vertices of the tetrahedral mesh or to the function values at these points.

*Robustness to volume discretization.* We consider the solution  $\mathbf{F}(t) = \mathbf{K}_t \mathbf{e}_i$  to the volumetric heat equation, whose initial condition takes value 1 at the anchor point  $\mathbf{p}_i$  (black dot) and 0 otherwise. The spectrum-free computation of the volumetric heat kernel is stable with respect to the volume density; a lower volume sampling (Fig. 4) does not affect the shape and distribution of the level sets and a higher resolution of  $\mathcal{M}$  improves the sampling of the level-sets, which remain smooth and uniformly distributed around the anchor.

To discuss the capability of the diffusion smoothing to recover the function underlying a noise signal  $f$  and its local details on  $\mathcal{M}$ , let  $f$  be the initial condition of the

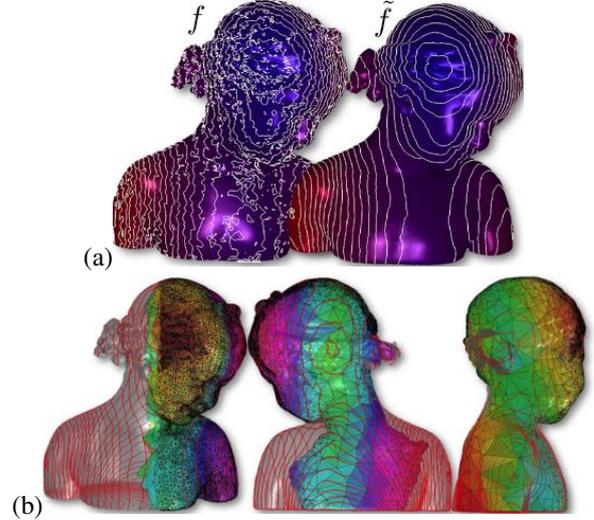


Figure 6: **Diffusion smoothing.** (a) Noisy volumetric map  $f$  on a tetrahedral mesh and (b) its diffusion smoothing ( $r = 7$ ) achieved by solving the heat equation with initial condition  $F(\cdot, 0) = f$ .

volumetric heat equation and  $f := \tilde{f} + e$  the perturbed initial condition, where  $e$  is a Gaussian noise. Indicating with  $\epsilon_\infty := \|F(\cdot, t) - \tilde{f}\|_\infty$  the  $\ell_\infty$  error between the ground-truth map  $f$  and the solution  $F(\cdot, t)$ , we get that  $\epsilon_\infty = 2.5 \times 10^{-6}$  ( $r = 5$ ) for the proposed approach;  $\epsilon_\infty = 8.9 \times 10^{-3}$ ,  $1.9 \times 10^{-4}$  for the spectral representation, with  $k = 130$  and  $k = 487$  eigenpairs, respectively;  $\epsilon_\infty = 2.7 \times 10^{-2}$  for the Euler backward method; and  $\epsilon_\infty = 4.1 \times 10^{-3}$  for the power method.

Our approximation scheme handles any map  $f$  defined on the vertices of  $\mathcal{M}$  or its boundary; for instance, Fig. 5(a) shows the diffusion process starting from a seed point  $\mathbf{p}_i$  belonging to the internal part of  $\mathcal{M}$ . In Figs. 5(b), 6, the seed point and the initial condition  $F(\cdot, 0) = f$  are defined on  $\partial\mathcal{M}$ .

*Comparison with ground-truth volumes.* For the evaluation of the approximation accuracy, we compare the computed (volumetric) diffusion distances, with correct results on the cylinder and sphere. We briefly recall that the diffusion distances are defined as [7]  $d^2(\mathbf{p}, \mathbf{q}) = \sum_{n=0}^{+\infty} \exp(-2\lambda_n t) |\phi_n(\mathbf{p}) - \phi_n(\mathbf{q})|^2$ ,  $\mathbf{p}, \mathbf{q} \in \mathcal{M}$ . According to [30], we consider analytic solutions for the Laplacian eigenvalues and eigenfunctions of the (volumetric) sphere and cylinder. Since we have an infinite number of eigenpairs, we select  $k$  such that the approximation  $d_k$  of the diffusion distance with  $k$  eigenpairs becomes stationary; i.e.,  $|d_{k+1}(\mathbf{p}, \mathbf{q}) - d_k(\mathbf{p}, \mathbf{q})| < \epsilon$ , where  $\epsilon$  is equal to the 1%. Once the exact Laplacian eigenpairs have been computed, we compare the ground-truth diffusion distance and the approximation

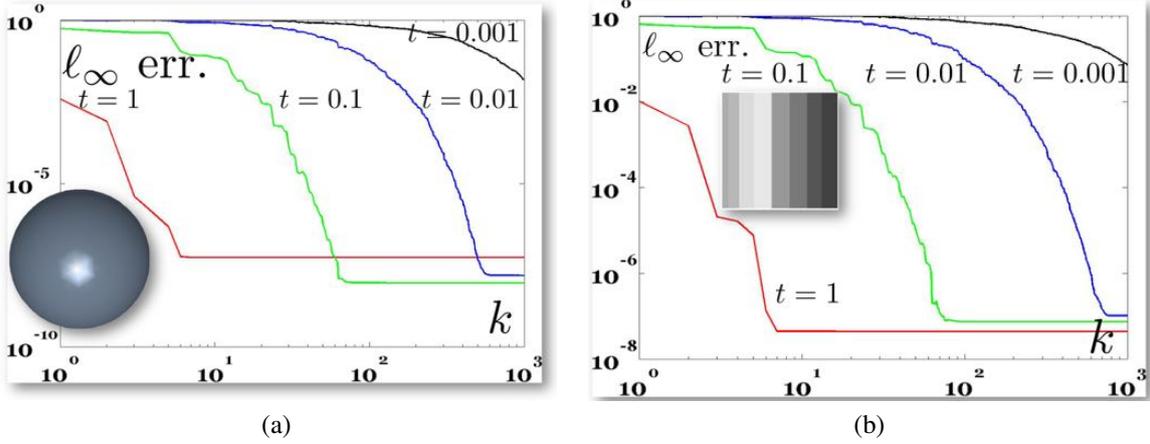


Figure 7: **Padé-Chebyshev and truncated spectral approximations of the diffusion distances.**  $\ell_\infty$  error (y-axis) between the diffusion distance computed with the truncated spectral approximation with  $k$  (x-axis) Laplacian eigenpairs on a tetrahedralization of (a) the sphere and (b) the cylinder. For the Padé-Chebyshev method ( $r = 5$ ) and all the time values, the  $\ell_\infty$  error with respect to the ground-truth is lower than  $8.9 \times 10^{-6}$ .

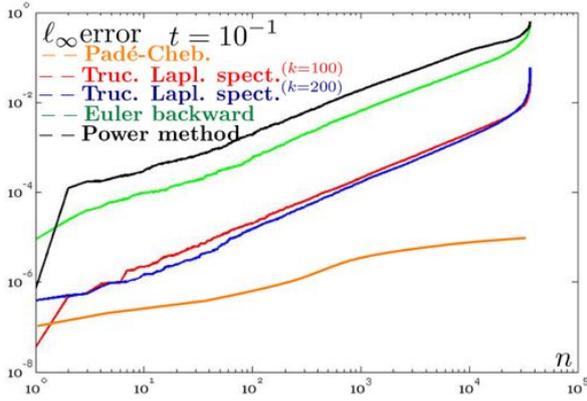


Figure 8: **Comparison of the accuracy of different approximations of the diffusion distances.**  $\ell_\infty$  error (y-axis) between the ground-truth volumetric diffusion distances on the cylinder, with a different sampling (x-axis) and scales, and their computation with the Padé-Chebyshev ( $r = 5$ ), Euler backward (green line), power (black line) methods, and the truncated spectral approximation with 100 and 200 eigenpairs (red, blue line).

provided by previous work.

For small scales (Fig. 7,  $t = 10^{-2}, 10^{-3}$ ), the  $\ell_\infty$  discrepancy (y-axis) between the approximation of the diffusion distance computed with the Padé-Chebyshev method and  $k$  eigenpairs remains higher than  $10^{-2}$ , with  $k \leq 280$ . For large scales (e.g.,  $t = 1, 10^{-1}$ ), increasing  $k$  reduces the approximation error until it becomes almost constant and close to zero. In fact, local shape features encoded by the heat kernel are recovered for a small  $t$  using the eigenvectors associated with high frequencies, thus requiring the computation of a large part of the Laplacian spectrum and without achieving the approx-

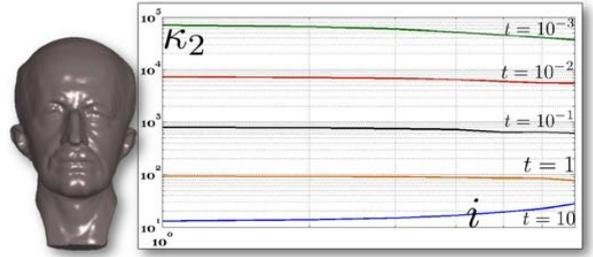


Figure 9: **Numerical stability of the Padé-Chebyshev approximation.** With reference to Fig. 5, conditioning number  $\kappa_2$  (y-axis) of the matrices  $\{(t\mathbf{L} + \theta_i\mathbf{B})\}_{i=1}^7$ , for different time values  $t$ ; the indices of the coefficients  $\{\theta_i\}_{i=1}^7$  are reported on the x-axis.

imation accuracy of the spectrum-free approach, which remains lower than  $8.9 \times 10^{-6}$  for all the scales.

*Comparison with previous work.* We briefly recall that the “power” method applies the identity  $(\mathbf{K}_{t/m})^m = \mathbf{K}_t$ , where  $m$  is chosen in such a way that  $t/m$  is sufficiently small to guarantee that the approximation  $\mathbf{K}_{t/m} \approx (\mathbf{I} - \frac{t}{m}\tilde{\mathbf{L}})$  is accurate. However, the selection of  $m$  and its effect on the approximation accuracy cannot be estimated a-priori. In [9, 13], the solution to the heat equation is computed through the Euler backward method  $(t\tilde{\mathbf{L}} + \mathbf{I})\mathbf{F}_{k+1}(t) = \mathbf{F}_k(t)$ ,  $\mathbf{F}_0 = \mathbf{f}$ . The resulting functions are over-smoothed and converge to a constant map, as  $k \rightarrow +\infty$ . Fig. 8 compares the accuracy of the diffusion distances computed with (i) the proposed approach; (ii) the spectral representation of the heat kernel, with  $k$  eigenpairs; (iii) the Euler backward method; and (iv) the power method. The approximation accuracy of the Padé-Chebyshev method is higher than

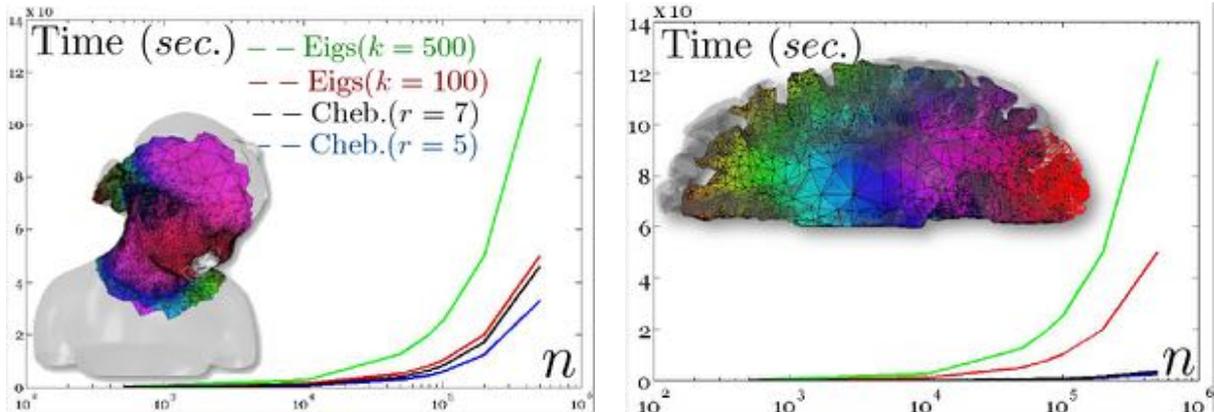


Figure 10: **Computational cost.** Timings (in seconds) for the evaluation of the heat kernel on tetrahedral meshes with  $n$  points, approximated with  $k$  eigenpairs (*Eigs*) and the Padé-Chebyshev method. Tests have been performed on a 2.7 GHz Intel Core i7 Processor, with 8 GB memory.

Teeth (Figs. 4)							
$n$	Cheb.	$k=50$	$\times$	$k=100$	$\times$	$k=500$	$\times$
10K	12.45	12.24	<b>0.98</b>	26.31	<b>2.11</b>	118.34	<b>9.50</b>
50K	37.02	33.25	<b>0.89</b>	68.23	<b>1.84</b>	289.73	<b>7.82</b>
80K	56.57	51.23	<b>0.90</b>	99.44	<b>1.75</b>	512.67	<b>9.06</b>
100K	0.54	0.48	<b>0.88</b>	1.17	<b>2.16</b>	5.10	<b>9.44</b>
Bimba (Figs. 6)							
5K	0.11	0.07	<b>0.63</b>	0.10	<b>0.90</b>	0.53	<b>4.81</b>
25K	3.45	0.71	<b>0.20</b>	1.56	<b>0.45</b>	8.23	<b>2.38</b>
35K	7.01	1.23	<b>0.17</b>	2.50	<b>0.35</b>	0.91	<b>0.12</b>
50K	14.01	2.45	<b>0.17</b>	5.32	<b>0.37</b>	0.98	<b>0.06</b>
Max Planck (Figs. 8)							
70K	28.56	27.32	<b>0.95</b>	57.78	<b>2.02</b>	131.23	<b>4.59</b>
100K	68.23	52.34	<b>0.76</b>	110.32	<b>1.61</b>	557.13	<b>8.16</b>
120K	99.21	75.23	<b>0.75</b>	156.78	<b>1.58</b>	698.23	<b>7.03</b>
150K	155.23	114.01	<b>0.73</b>	231.23	<b>1.48</b>	893.78	<b>5.75</b>

Table 1: Timings (in seconds) for the evaluation of the heat diffusion kernels on 3D volumes with  $n$  points, approximated with  $k=50, 100, 500$  eigenpairs (*Eigs*) and the Padé-Chebyshev approximation (*Cheb.*,  $r=7$ ). Column ' $\times$ ' indicates the number of times the computational cost is reduced. Tests have been performed on a 2.7 GHz Intel Core i7 Processor, with 8 GB memory.

the truncated Laplacian spectrum with  $k$  eigenpairs,  $k=1, \dots, 10^3$ , the Euler backward method, and the power method. Reducing the scale, the accuracy of the Padé-Chebyshev remains almost unchanged while the other methods are affected by a larger discrepancy and tend to have an analogous behavior ( $t=10^{-4}$ ). Finally, the coefficient matrices in Eq. (3) are well-conditioned, as also confirmed by our experiments (Fig. 9).

**Computational cost.** The spectrum-free approach reduces the heat equation to solve  $r$  sparse, symmetric, linear systems (c.f., Eq. (3)), whose coefficient matrices have the same structure and sparsity of the connectivity matrix of the grid underlying the discrete volume. Applying an iterative and sparse linear solver (e.g., Gauss-Seidel method, conjugate gradient) [15] (Ch. 10), the computational cost for the evaluation of the heat kernel and the diffusion distance between two points is  $O(r\tau(n))$ , where  $O(\tau(n))$  is the computational cost of

the selected solver (Fig. 10, Table 1). Here, the function  $\tau(n)$ , which depends on the number  $n$  of samples and the sparsity of the coefficient matrix, typically varies from  $\tau(n) = n$  to  $\tau(n) = n \log n$ . In fact,  $O(n \log n)$  is the average computational cost of the aforementioned iterative solvers of sparse linear systems.

Finally, the spectrum-free computation of the one-to-all (volumetric) distances  $d(\mathbf{p}_i, \mathbf{p}_j)$ ,  $j=1, \dots, n$ , takes  $O(rn\tau(n))$  time; in fact, we solve the linear system (3) with  $n$  different right-hand vectors  $\mathbf{f} := \mathbf{e}_i - \mathbf{e}_j$ ,  $j=1, \dots, n$ . Computing a fixed number  $k$  of eigenpairs in  $O(kn^2)$  time, the one-to-all distance is evaluated in super linear time [37]; indeed, the computational cost of the spectrum-free computation is lower than the truncated spectral approximation with  $k(n) \geq rn\tau(n)/n$  eigenpairs. The truncated spectral approximation is accurate only if the exponential filter decays fast (e.g., large values of time). Otherwise, a larger number of eigenpairs is needed and the resulting computational cost varies from  $O(kn^2)$  to  $O(n^3)$  time, according to the sparsity of the Laplacian matrix. Furthermore, the effect of the selected eigenpairs on the approximation accuracy cannot be estimated without computing the whole spectrum. A similar discussion applies to convolutions with heat kernels for several time values.

As future work, we plan to apply the proposed approach to volumetric mesh fairing and to the evaluation of volume-based descriptors in shape analysis.

**Acknowledgements** Special thanks are given to the anonymous Reviewers, Bianca Falcidieno, Michela Spagnuolo, and the Shape Modeling Group at CNR-IMATI for their valuable comments. This work has been partially supported by the EU FP7 Integrated Project

IQmulus (FP7-ICT-2011-318787). Shapes are courtesy of the AIM@SHAPE Repository; tetrahedral meshes have been generated with the TETGEN Software.

## References

- [1] G.G. Allaire. *Numerical Analysis and Optimization: An Introduction to Mathematical Modelling and Numerical Simulation*. Numerical Mathematics and Scientific Computation. Oxford University Press, Incorporated, 2007.
- [2] P. Alliez, D. Cohen-Steiner, M. Yvinec, and M. Desbrun. Variational tetrahedral meshing. *ACM Trans. on Graphics*, 24(3):617–625, 2005.
- [3] C. L. Bajaj and G. Xu. Anisotropic diffusion of subdivision surfaces and functions on surfaces. *ACM Trans. on Graphics*, 22:4–32, 2002.
- [4] M. Belkin and P. Niyogi. Towards a theoretical foundation for Laplacian-based manifold methods. *Journal of Computer System Sciences*, 74(8):1289–1308, 2008.
- [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. *Intern. Journal of Computer Vision*, 2-3:266–286, 2010.
- [6] A. M. Bronstein, M. M. Bronstein, M. Ovsjanikov, and L. J. Guibas. Shape Google: geometric words and expressions for invariant shape retrieval. *ACM Trans. on Graphics*, 30(1), 2011.
- [7] M. M. Bronstein and A. M. Bronstein. Shape recognition with spectral distances. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 33(5):1065–1071, 2011.
- [8] A.J. Carpenter, A. Ruttan, and R.S. Varga. Extended numerical computations on the 1/9 conjecture in rational approximation theory. In *Rational Approximation and Interpolation*, volume 1105 of *Lecture Notes in Mathematics*, pages 383–411. Springer Berlin Heidelberg, 1984.
- [9] U. Clarenz, U. Diewald, and M. Rumpf. Anisotropic geometric diffusion in surface processing. In *IEEE Visualization*, pages 397–405, 2000.
- [10] R. R. Coifman and S. Lafon. Diffusion maps. *Applied and Computational Harmonic Analysis*, 21(1):5–30, 2006.
- [11] K. Crane, C. Weischedel, and M. Wardetzky. Geodesics in heat: a new approach to computing distance based on heat flow. *ACM Siggraph*, 32(14), 2013.
- [12] F. de Goes, S. Goldenstein, and L. Velho. A hierarchical segmentation of articulated bodies. *Computer Graphics Forum*, 27(5):1349–1356, 2008.
- [13] 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.
- [14] K. Gebal, J. A. Berentzen, H. Aanæs, and R. Larsen. Shape analysis using the auto diffusion function. *Computer Graphics Forum*, 28(5):1405–1413, 2009.
- [15] G. Golub and G.F. VanLoan. *Matrix Computations*. John Hopkins University Press, 2nd Edition, 1989.
- [16] I. Guskov, W. Sweldens, and P. Schröder. Multiresolution signal processing for meshes. *ACM Siggraph*, pages 325–334, 1999.
- [17] D. K. Hammond, P. Vandergheynst, and R. Gribonval. Wavelets on graphs via spectral graph theory. *Applied and Computational Harmonic Analysis*, 30(2):129 – 150, 2011.
- [18] P. C. Hansen and D. P. O’Leary. The use of the l-curve in the regularization of discrete ill-posed problems. *SIAM Journal of Scientific Computing*, 14(6):1487–1503, 1993.
- [19] A. Kuprat. Modeling microstructure evolution using gradient-weighted moving finite elements. *SIAM Journal on Scientific Computing*, 22(2):535–560, 2000.
- [20] A. Kuprat, A. Khamayseh, D. George, and L. Larkey. Volume conserving smoothing for piecewise linear curves, surfaces, and triple lines. *Journal of Computational Physics*, 172(1):99 – 118, 2001.
- [21] X. Li, H. Xu, S. Wan, Z. Yin, and W. Yu. Feature-aligned harmonic volumetric mapping using MFS. *Computers & Graphics*, 34(3):242–251, 2010.
- [22] S. Liao, R. Tong, J. Dong, and F. Zhu. Gradient field based inhomogeneous volumetric mesh deformation for maxillofacial surgery simulation. *Computers & Graphics*, 33(3):424 – 432, 2009.
- [23] R. Litman, A.M. Bronstein, and M.M. Bronstein. Diffusion-geometric maximally stable component detection in deformable shapes. *Computers & Graphics*, 35(3):549–560, 2011.
- [24] T. Martin, E. Cohen, and M. Kirby. Volumetric parameterization and trivariate B-spline fitting using harmonic functions. In *Proc. of the ACM Symposium on Solid and Physical Modeling*, pages 269–280, 2008.
- [25] F. Memoli. Spectral Gromov-Wasserstein distances for shape matching. In *Workshop on Non-Rigid Shape Analysis and Deformable Image Alignment*, pages 256–263, 2009.
- [26] M. Ovsjanikov, Q. Mérigot, F. Memoli, and L. Guibas. One point isometric matching with the heat kernel. *ACM Symposium on Discrete Algorithms*, pages 650–663, 2010.
- [27] G. Patanè. wFEM heat kernel: Discretization and applications to shape analysis and retrieval. *Computer Aided Geometric Design*, 30(3):276–295, 2013.
- [28] G. Patanè and M. Spagnuolo. Heat diffusion kernel and distance on surface meshes and point sets. *Computers & Graphics*, 37(6):676–686, 2013.
- [29] D. Raviv, M. M. Bronstein, A. M. Bronstein, and R. Kimmel. Volumetric heat kernel signatures. In *Proceedings of the ACM Workshop on 3D Object Retrieval*, pages 39–44, 2010.
- [30] M. Reuter, F.-E. Wolter, M. E. Shenton, and M. Niethammer. Laplace-beltrami eigenvalues and topological features of eigenfunctions for statistical shape analysis. *Computer-Aided Design*, 41(10):739–755, 2009.
- [31] R. M. Rustamov. Interpolated eigenfunctions for volumetric shape processing. *The Visual Computer*, 27(11):951–961, 2011.
- [32] R. M. Rustamov. Multiscale biharmonic kernels. *Computer Graphics Forum*, 30(5):1521–1531, 2011.
- [33] J. R. Shewchuk. Delaunay refinement algorithms for triangular mesh generation. *Computational Geometry Theory and Applications*, 22(1-3):21–74, 2002.
- [34] A. Spira, R. Kimmel, and N. Sochen. A short-time beltrami kernel for smoothing images and manifolds. *Trans. on Image Processing*, 16(6):1628–1636, 2007.
- [35] 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.
- [36] Y. Tong, S. Lombeyda, A. N. Hirani, and M. Desbrun. Discrete multiscale vector field decomposition. *ACM Trans. on Graphics*, 22(3):445–452, 2003.
- [37] B. Vallet and B. Lévy. Spectral geometry processing with manifold harmonics. *Computer Graphics Forum*, 27(2):251–260, 2008.
- [38] R.S. Varga. *Scientific computation on mathematical problems and conjectures*. SIAM, CBMS-NSF regional conference series in applied mathematics, 1990.
- [39] A. Vaxman, M. Ben-Chen, and C. Gotsman. A multi-resolution approach to heat kernels on discrete surfaces. *ACM Trans. on Graphics*, 29(4):1–10, 2010.
- [40] A. P. Witkin. Scale-space filtering. In *Proc. of the Intern. Joint Conference on Artificial Intelligence*, pages 1019–1022, 1983.