Heat Diffusion Kernel and Distance on Surface Meshes and Point Sets

Giuseppe Patané[a], Michela Spagnuolo[a],

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

Abstract

The heat diffusion distance and kernel have gained a central role in geometry processing and shape analysis. This paper addresses a novel discretization and spectrum-free computation of the diffusion kernel and distance on a 3D shape $\mathcal{P}$ represented as a triangle mesh or a point set. After rewriting different discretizations of the Laplace-Beltrami operator in a unified way and using an intrinsic scalar product on the space of functions on $\mathcal{P}$, we derive a shape-intrinsic heat kernel matrix, together with the corresponding diffusion distances. Then, we propose an efficient computation of the heat distance and kernel through the solution of a set of sparse linear systems. In this way, we bypass the evaluation of the Laplacian spectrum, the selection of a specific subset of eigenpairs, and the use of multi-resolutive prolongation operators. The comparison with previous work highlights the main features of the proposed approach in terms of smoothness, stability to shape discretization, approximation accuracy, and computational cost.

Keywords: Spectral methods, heat diffusion equation, heat diffusion kernel and distance, Chebyshev approximation.

1. Introduction

The heat diffusion kernel and distance on manifolds play a central role in several applications such as spectral clustering, data classification, dimensionality reduction, kernel principal component analysis, and visualization. Among their main properties, we mention the intrinsic and multi-scale definition with respect to the input shape, the invariance to isometries, the shape-awareness, the robustness to noise and tessellation. The heat diffusion distance and kernel have been successfully applied to shape segmentation [10] and comparison [6, 14, 20, 23, 33]; to the computation of the gradient of discrete maps [38]; and to the multi-scale approximation of functions [24]. The diffusion distance and kernel also play a central role in several applications, such as dimensionality reduction with spectral embeddings [2, 39]; data visualization [2, 16, 27, 34], representation [7, 30, 40], and classification [22, 29, 32] through auto-diffusion maps [14] and diffusion distances [5, 9, 17].

Overview and contribution. This paper addresses the discretization and computation of the heat diffusion kernel and distance on 3D shapes represented as triangle meshes or point sets. For surfaces represented as triangle meshes or point sets, we firstly rewrite the discrete Laplace-Beltrami operator in a unified way as $\tilde{L} := B^{-1}L$, where L is a symmetric, semi-positive definite matrix and B is positive definite. For triangle meshes, L is the Laplacian matrix with cotangent weights and B is the diagonal matrix of the Voronoi areas [11], or the FEM mass matrix [26, 35], or the identity matrix [25]. For point sets, L is the Gram matrix associated to the exponential kernel and B is the identity matrix [3, 4] or the diagonal matrix of the Voronoi areas [19]. Under these assumptions, L represents the $\mathcal{L}^2$ scalar product on the space of 1-forms.

Successively, the diffusion kernel and distance are discretized with respect to the scalar product $\langle f, g\rangle_B := \mathbf{f}^\top B\mathbf{g}$ on the space $\mathcal{F}(\mathcal{P}) := \{f : \mathcal{P} \to \mathbb{R}, \mathbf{f} := f(\mathbf{p}_i)_{i=1}^n\}$ of functions defined on $\mathcal{P}$, which is either a triangle mesh or a point set. Here, the symmetric, positive-definite matrix B is chosen in such a way that the corresponding scalar product is intrinsic to the surface underlying $\mathcal{P}$ and is adapted to its local sampling. Using the B-scalar product, we derive the heat kernel $K_t := XD_tX^\top B$, $D_t := \mathrm{diag}(\exp(-\lambda_i t))_{i=1}^n$, where X is the eigenvectors' matrix associated to the generalized eigenproblem $LX = BX\Gamma$ and $\Gamma := \mathrm{diag}(\lambda_i)_{i=1}^n$, $0 \le \lambda_i \le \lambda_{i+1}$, is the eigenvalue matrix. Verifying that the heat kernel K_t is still the exponential $\exp(-tB^{-1}L)$ of the Laplacian matrix $\tilde{L} := B^{-1}L$, we motivate the inclusion of the matrix B in the definition of the heat kernel and the importance of the orthonormality of the Laplacian eigenvectors with respect to $\langle\cdot,\cdot\rangle_B$ for the validity of the exponential representation. Using the proposed discretization of the heat kernel and the shape-intrinsic scalar product, we also derive the corresponding heat diffusion distance, which are compared with respect to previous work.

If the matrix B is lumped to the positive diagonal matrix D, then the heat kernel K_t becomes equal to the discretization $K_t^\star := XD_tX^\top D$, which holds for Laplacians of type $L := D^{-1}W$ [6, 23, 28, 33, 37]. In this case, 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 points of $\mathcal{P}$ for both the Voronoi-cot and Voronoi-exp Laplacian weights. In previous work [14, 6, 23, 38, 33, 37], the Laplacian eigenvectors used for the computation of the diffusion distances on tri-

angle meshes are orthonormal with respect to the scalar product induced by the diagonal matrix D whose entries are the areas of the Voronoi regions associated to the mesh vertices. The linear FEM mass matrix B 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 proposed discretization of the heat kernel has also a higher robustness against topological and scale changes, irregular sampling, and noise.

Then, we compute any discretization of the heat diffusion distance through the Chebyshev approximation [8, 15, 21] of the weighted heat kernel matrix. In this case, the computation of the diffusion kernel and distance is achieved through the solution of sparse linear systems and a sequence of matrix-vector multiplications, without computing the Laplacian spectrum. To this end, we consider the (r, r)-degree rational function that is the best approximation of the exponential function with respect to the ℓ_∞ norm. According to [15, 36], the computation of the heat distance is reduced to the solution of r sparse linear systems in $O(rn)$ time with iterative solvers, such as the Jacobi, Gauss-Seidel, and minimum residual methods [15]. The overall cost of the computation of the value $K_t(\mathbf{p}_i, \mathbf{p}_j) = \mathbf{e}_i^\top K_t \mathbf{e}_j$ at k feature points varies from $O(kn)$ to $O(kn^2)$, according to the sparsity of the coefficient matrix. The approximation accuracy is lower than 10^{-r} (e.g., $r := 5, 7$) and it can be further reduced by slightly increasing the degree r of the rational Chebyshev polynomial used for the spectrum-free computation.

As main contribution with respect to previous work, our approach is independent of the evaluation of the Laplacian spectrum, the selection of a specific subset of eigenpairs, and the use of multiresolutive prolongation operators. To speed-up the computation of the values of the heat diffusion kernel and distances for several values of t, or between a large number of points, it is sufficient to apply iterative solvers of linear systems or a pre-factorization of the coefficient matrices. The proposed algorithm is also robust with respect to irregular sampling density, noise, mesh degeneracies; can be applied to different Laplacian weights; and is free of user-defined parameters. On the contrary, in previous work the resolution of the simplified approximation of the input surface, on which the Laplacian matrix is computed, and the number of Laplacian eigenpairs are tuned according to the target approximation accuracy. Since our approach works mainly with matrices, the evaluation of the spectral distances is independent of the discretization of the input surface as a triangle mesh or a point cloud.

Related work. In previous work, the computation of the Laplacian spectrum is the main computational bottleneck for the evaluation of the heat diffusion kernel and distance; in fact, it takes from $O(n)$ to $O(n^3)$ time, according to the sparsity of the Laplacian matrix. Even though iterative solvers of sparse eigenproblems reduce the computational cost to super linear time [35], the computation of the whole spectrum is unfeasible. To overcome this drawback, the solution to the heat equation and the diffusion distance are approximated as

$$\begin{cases} \mathbf{F}_k(t) = \sum_{i=1}^{k} \exp(-\lambda_i t)\langle \mathbf{f}, \mathbf{x}_i \rangle_B \mathbf{x}_i, & (a) \\ d^2(\mathbf{p}_i, \mathbf{p}_j) = \sum_{l=1}^{k} \exp(-\lambda_l t)|x_l(\mathbf{p}_i) - x_l(\mathbf{p}_j)|^2, & (b) \end{cases} \quad (1)$$

where $x_j(\mathbf{p}_i)$ the ith component of the eigenvector $\mathbf{x}_j$ and k is the number of selected eigenpairs. Indeed, only a part of the Laplacian spectrum is used to approximate the heat kernel and distances through a truncated sum, which involves the contribution of the Laplacian eigenvectors related to the smaller eigenvalues. In this case, the heat diffusion distance and kernel can only be approximated and not exactly evaluated. To overcome these limitations of the partial spectral approximation of the heat diffusion distance and make the computation faster, previous work has also approximated the heat kernel by prolongating its values computed on a sub-sampling of the input shape [37], similarly to the computation of diffusion wavelets [9] through multiresolution decompositions. In all the previous approaches, the number of eigenpairs selected for the approximation of the heat diffusion kernel is heuristically adapted to the temporal parameter t or selected by the user. A small value of t generally requires the time-consuming computation of a large number of eigenvectors, whose capability of encoding local shape features depends on the mesh resolution used for the spectrum calculation. Furthermore, on point-sampled and non-manifold surfaces the definition of multi-resolutive and prolongation operators is generally ambiguous.

Paper organization. We introduce the proposed discretization of the heat diffusion kernel and distance (Sect. 2), their spectrum-free computation (Sect. 3), and an overview of their main properties (Sect. 4). Finally, we conclude the paper and present future work (Sect. 5).

2. Discrete heat diffusion equation and distances

The *scale-based representation* $F : \mathcal{M} \times \mathbb{R}^+ \to \mathbb{R}$ of the map $f : \mathcal{M} \to \mathbb{R}$, defined on a surface $\mathcal{M}$, is the solution to the *heat diffusion equation* $\partial_t F(\mathbf{p}, t) = -\Delta F(\mathbf{p}, t)$, $F(\cdot, t) = f$, $(\mathbf{p}, t) \in \mathcal{M} \times \mathbb{R}^+$, and it is written through the convolution operator $\star$ as

$$\begin{cases} F(\mathbf{p}, t) := k_t(\mathbf{p}, \cdot) \star h = \int_{\mathcal{M}} k_t(\mathbf{p}, \mathbf{q}) h(\mathbf{q}) d\mathbf{q} \\ k_t(\mathbf{p}, \mathbf{q}) := \sum_{i=1}^{+\infty} \exp(-\lambda_i t)\phi_i(\mathbf{p})\phi_i(\mathbf{q}), \end{cases}$$

where k_t is the *heat diffusion kernel* and (λ_i, ϕ_i) are the eigenpairs of the Laplace-Beltrami operator; i.e., $\Delta\phi_i = \lambda_i \phi_i$, $i \in \mathbb{N}$.

We discretize the heat diffusion equation on the space $\mathcal{F}(\mathcal{P}) := \{f : \mathcal{P} \to \mathbb{R}, \mathbf{f} := f(\mathbf{p}_i)_{i=1}^n\}$ of maps defined on $\mathcal{P}$. If $\mathcal{P}$ is the set of vertices of a triangle mesh, then the values of f are extended from the surface vertices along the edges and on the faces by using barycentric coordinates. If $\mathcal{P}$ is a point set, then f is defined only on $\mathcal{P}$. In $\mathcal{F}(\mathcal{P})$, we consider the scalar product $\langle f, g \rangle_B := \mathbf{f}^\top B \mathbf{g}$, induced by the symmetric, positive-definite matrix B, where $\mathbf{f} := (f(\mathbf{p}_i))_{i=1}^n$ and $\mathbf{g} := (g(\mathbf{p}_i))_{i=1}^n$ are the arrays of function values. As discussed in [12**?** , 35], the matrix B must be chosen in such a way that the corresponding scalar product is intrinsic to the surface underlying $\mathcal{P}$ and is adapted to its local sampling. For more details on the choice of B, we refer the reader to Sect. 4.

Let us now consider the Laplacian spectrum $\{(\lambda_i, \mathbf{x}_i)\}_{i=1}^n$ of the couple (L, B), which satisfies the generalized eigenproblem

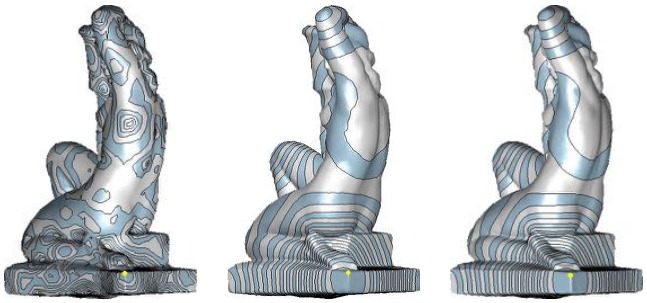

(a) Vor.-cot $k = 100$ (b) FEM $k = 100$ (c) FEM $k = 500$

Figure 1: Level sets of the map $\mathbf{F}_k(t)$ in Eq. (1a), which approximates $K_t\mathbf{e}_i$. The discrete map e_i takes value 1 at the anchor point (yellow dot) and 0 otherwise. In this example, we have selected the (a) Voronoi-cot and (b,c) the FEM weights, the same time parameter t, and k Laplacian eigenpairs.

$L\mathbf{x}_i = \lambda_i B\mathbf{x}_i$, $0 \leq \lambda_i \leq \lambda_{i+1}$, and the orthonormality conditions $\mathbf{x}_i^\top B\mathbf{x}_j = \delta_{ij}$, $i, j = 1, \dots, n$. These relations are rewritten in matrix form as

$$LX = BX\Gamma, \quad \Gamma := \mathrm{diag}(\lambda_i)_{i=1}^n, \quad X^\top BX = I, \tag{2}$$

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.

Introducing the vector $\mathbf{F}(t) := (F(\mathbf{p}_i, t))_{i=1}^n$, the heat diffusion equation is discretized in $\mathcal{F}(\mathcal{P})$ as $\partial_t\mathbf{F}(t) = -\tilde{L}\mathbf{F}(t)$, $\mathbf{F}(0) = \mathbf{f}$. Let us now express the function $\mathbf{F}(t) = \sum_{i=1}^n \alpha_i(t)\mathbf{x}_i$ in terms of the eigensystem of (L, B), where $\alpha(t) := (\alpha_i(t))_{i=1}^n$ is the unknown vector. Using the identity $\mathbf{f} = \sum_{i=1}^n \langle \mathbf{f}, \mathbf{x}_i\rangle_B \mathbf{x}_i$, the invertibility of the matrix B, and the linear independence of the Laplacian eigenvectors, each component $\alpha_i(t)$ satisfies the differential equation $\alpha_i'(t) + \lambda_i\alpha_i(t) = 0$, with boundary condition $\alpha_i(0) = \langle \mathbf{f}, \mathbf{x}_i\rangle_B$. Then, the *scale-based representation of* $f : \mathcal{P} \to \mathbb{R}$ is $\mathbf{F}(t) = \sum_{i=1}^n \exp(-\lambda_i t)\langle \mathbf{f}, \mathbf{x}_i\rangle_B \mathbf{x}_i$ and it is re-written in matrix form as $\mathbf{F}(t) = K_t\mathbf{f}$, where the *heat kernel matrix*

$$K_t := XD_tX^\top B, \quad D_t := \mathrm{diag}\left(\exp(-\lambda_i t)\right)_{i=1}^n, \tag{3}$$

is self-adjoint with respect to the B-scalar product. Recalling that the first eigenpair is $(0, \mathbf{1})$, K_t converges to the constant averaging operator, as $t \to +\infty$; i.e., $\lim_{t\to+\infty} K_t\mathbf{f} = (\mathbf{f}^\top B\mathbf{1})\mathbf{1}$. Finally, the representation

$$\begin{aligned} K_t &=_{(3)} X\mathrm{diag}\left(\sum_{k=0}^{+\infty} \frac{(-\lambda_i t)^k}{k!}\right)_{i=1}^n X^\top B \\ &= \sum_{k=0}^{+\infty} \frac{(-t)^k}{k!} X\Delta^k X^\top B \\ &= \sum_{k=0}^{+\infty} \frac{(-tB^{-1}L)^k}{k!}, \quad (B^{-1}L)^k =_{(2)} X\Delta^k X^\top B \\ &= \exp(-tB^{-1}L) \end{aligned} \tag{4}$$

of the heat kernel matrix as the exponential of the Laplacian matrix with respect to time motivates the inclusion of the matrix B in the heat kernel (3) and the importance of the orthogonality of the Laplacian eigenvectors for the validity of (4).

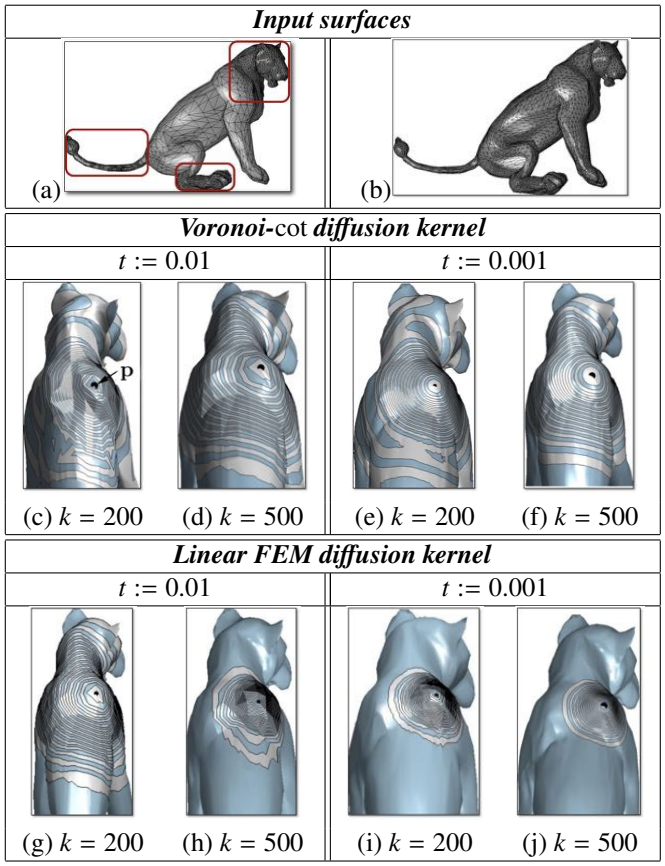

Figure 2: Approximation the map $\mathbf{F}(t) = K_t\mathbf{e}_i$ generated using k Laplacian eigenfunctions in Eq. (1a), with (c-f) Voronoi-cot and (g-j) FEM weights on a 3D shape with a (a) coarse and (b) fine sampling.

Heat diffusion distances. Using the identity

$$K_t^\top BK_t =_{(2)} BXD_{2t}X^\top B = BK_{2t}. \tag{5}$$

and according to [9], let us introduce the *diffusion distance*, with respect to K_t and the B-scalar product, as

$$\begin{aligned} d_B^2(\mathbf{p}_i, \mathbf{p}_j) :&= \left\|K_t(\mathbf{p}_i, \cdot) - K_t(\mathbf{p}_j, \cdot)\right\|_B^2 \\ &= \left\|K_t(\mathbf{e}_i - \mathbf{e}_j)\right\|_B^2 \\ &= (\mathbf{e}_i - \mathbf{e}_j)^\top K_t^\top BK_t(\mathbf{e}_i - \mathbf{e}_j) \\ &=_{(5)} (\mathbf{e}_i - \mathbf{e}_j)^\top BXD_{2t}X^\top B(\mathbf{e}_i - \mathbf{e}_j) \\ &= \sum_{l=1}^n \exp(-2\lambda_l t)\left|\langle \mathbf{x}_l, \mathbf{e}_i - \mathbf{e}_j\rangle_B\right|^2. \end{aligned} \tag{6}$$

Comparing this expression with the standard discretization of the heat distance (1b), we get that (6) is achieved from (1b) by replacing $x_k(\mathbf{p}_i) = \mathbf{e}_i^\top\mathbf{x}_k$ with $\mathbf{e}_i^\top B\mathbf{x}_k$ and rescaling the temporal variable. However, Eq. (1b) does not take into account the intrinsic B-scalar product, or its lumped approximation, thus disregarding the underlying generalized eigenproblem.

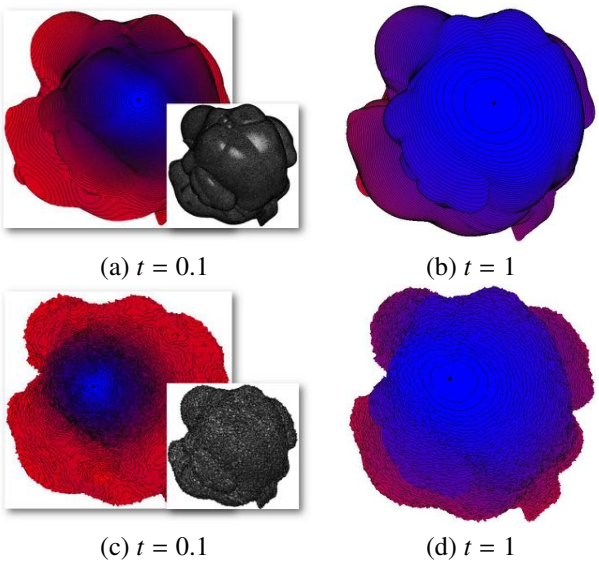

(a) $t = 0.1$ (b) $t = 1$

(c) $t = 0.1$ (d) $t = 1$

Figure 3: Level sets of the linear FEM diffusion distance (6), computed using the Chebyshev approximation ($r := 7$) in Eq. (9), from a source point (black dot), with different values of t, on a (a,b) smooth and (c,d) noisy surface.

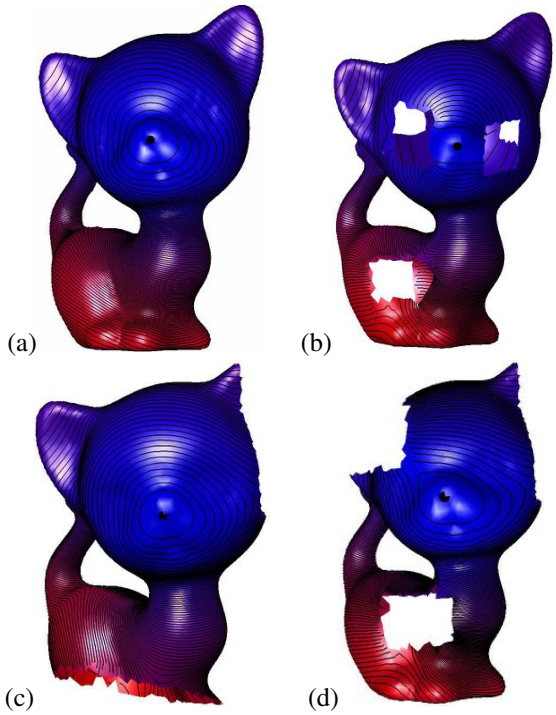

(a) (b) (c) (d)

Figure 4: (b-d) Level sets of the linear FEM diffusion distances (6), computed using the Chebyshev method ($r := 7$), from a source point (black dot) on partially-sampled surfaces. The behavior of the level sets remains almost unchanged and coherent with respect to the original shape in (a).

3. Computation of the discrete diffusion kernel and distance

For the computation of the heat distance, previous work overcomes the high computational cost for the evaluation of Laplacian spectrum by considering only the contribution of a fixed number of eigenpairs in Eq. (1b). This choice is motivated by the exponential decay of the filter factor $\exp(-\lambda_i t)$, which increases with λ_i and reduces the contribution of the corresponding eigenvector to the diffusion distance and kernel. We now describe how the heat diffusion distance and kernel are approximated through the Chebyshev approximation and without computing the Laplacian spectrum; to this end, we reduce both problems to evaluate the vector $K_t\mathbf{f}$, for specific choices of $\mathbf{f}$.

Computation of the heat diffusion kernel and distances. For the evaluation of $K_t(\mathbf{p}_i, \mathbf{p}_j) = \mathbf{e}_i^\top K_t \mathbf{e}_j$, we firstly compute the vector $K_t\mathbf{e}_i$ through the Chebyshev method; then, we consider its ith component. Rewriting the heat distance (6) in terms of the kernel $H_t := BK_{2t}$ as

$$\begin{aligned} d_B^2(\mathbf{p}_i, \mathbf{p}_j) &= \|K_t\mathbf{e}_i\|_B^2 - 2\langle K_t\mathbf{e}_i, K_t\mathbf{e}_j\rangle_B + \|K_t\mathbf{e}_j\|_B^2 \\ &=_{(5)} \mathbf{e}_i^\top BK_{2t}\mathbf{e}_i - 2\mathbf{e}_i^\top BK_{2t}\mathbf{e}_j + \mathbf{e}_j^\top BK_{2t}\mathbf{e}_j, \end{aligned}$$

we get that $d_B^2(\mathbf{p}_i, \mathbf{p}_j)$ is evaluated by computing $K_{2t}\mathbf{e}_l$, $K_{2t}\mathbf{e}_j$, which are then multiplied by $\mathbf{e}_i^\top B$ or $\mathbf{e}_j^\top B$. Indeed, the entries of the heat diffusion kernel and the corresponding distances are computed through the Chebyshev approximation and without extracting the Laplacian spectrum. Finally, we notice that the proposed computation of both the heat diffusion kernel and distance is independent of the discretization of the input surface as a triangle mesh or a point cloud.

Evaluation of $K_t\mathbf{f}$ through the Chebyshev approximation. To evaluate $\mathbf{F}(t) := K_t\mathbf{f}$, for any $\mathbf{f} \in \mathbb{R}^n$, we apply the Chebyshev method [8, 15, 21] to the weighted heat kernel. Using the relation (4), which expresses K_t as the exponential of the weighted Laplacian matrix with respect to time, we compute the (r, r)-degree rational function $c_{rr}(x) := a_{rr}(x)/b_{rr}(x)$ that is the best approximation of the exponential function with respect to the ℓ_∞ norm over the semi-axis $[0, +\infty)$. According to [15, 36], this function is $c_{rr}(x) = \alpha_0 + \sum_{i=1}^r \alpha_i/(x - \theta_i)$ and the exponential matrix is approximated by

$$\exp(C) \approx \alpha_0 I + \sum_{i=1}^{r} \alpha_i (C - \theta_i I)^{-1}. \tag{7}$$

In this representation, the poles $\{\theta_i\}_{i=1}^r$ and the coefficients $\{\alpha_i\}_{i=1}^r$ have been computed for $r := 5, 7$ [13]. For a general degree r and a fixed value of t, the coefficients of the rational approximation of the exponential function are computed using the Padé method [15], which is implemented in standard numerical software packages. Through Eq. (7), $\exp(C)\mathbf{f}$ is approximated as $\exp(C)\mathbf{f} \approx \alpha_0\mathbf{f} + \sum_{i=1}^r \alpha_i (C - \theta_i I)^{-1}\mathbf{f}$; i.e., $\exp(C)\mathbf{f}$ is the sum of the solutions of r sparse linear systems

$$(C - \theta_i I)\mathbf{g}_i = \alpha_i \mathbf{f}, \qquad i = 1, \ldots, r. \tag{8}$$

Since we cannot explicitly invert the matrix B and apply the scheme to $C := -tB^{-1}L$, we notice that each vector in Eq. (8) solves the system $(tB^{-1}L + \theta_i I)\mathbf{g}_i = -\alpha_i\mathbf{f}$ if and only if $(tL + \theta_i B)\mathbf{g}_i = -\alpha_i B\mathbf{f}$. For any $i = 1, \ldots, r$, $\mathbf{g}_i$ is now calculated as the solution of a sparse linear system and $K_t\mathbf{f}$ is recovered as

$$K_t\mathbf{f} \approx \alpha_0\mathbf{f} + \sum_{i=1}^{r} \mathbf{g}_i = \alpha_0\mathbf{f} - \sum_{i=1}^{r} \alpha_i (tL + \theta_i B)^{-1} B\mathbf{f}. \tag{9}$$

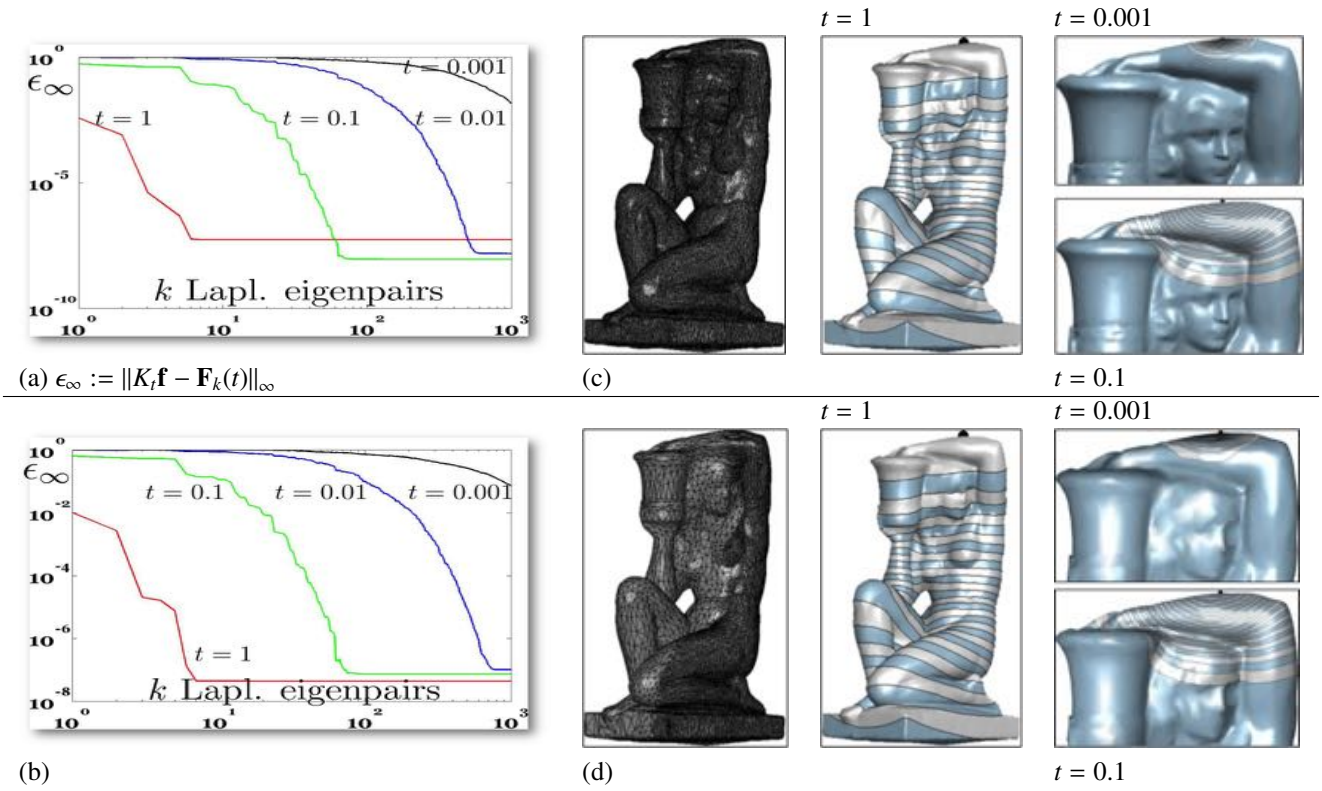

Figure 6: (a,b) ℓ_∞ error (y-axis) $\epsilon_\infty := \|K_t\mathbf{f} - \mathbf{F}_k(t)\|_\infty$ between the Chebyshev approximation ($r := 7$) of $\mathbf{F}(t) = K_t\mathbf{e}_i$ in Eq. (9) and the partial spectral representation $\mathbf{F}_k(t)$ in Eq. (1a), computed on a (c) regularly and (d) irregularly sampled shape, with respect to a different number k (x-axis, $k \leq 10^3$) of eigenpairs and values of t.

The solution $\mathbf{g}_i$ is computed by an iterative solver, which exploits the sparsity of the coefficient matrix $(tL + \theta_i B)$, without pre-factorizing the matrices L and B. Among the main solvers, we mention the Jacobi, Gauss-Seidel, and minimum residual methods (*minres*) [15]. Our implementation uses the *minres* procedure, which computes a minimum norm residual solution to the input linear system, whose coefficient matrix is symmetric, large, and sparse but not necessarily positive definite. Then, the overall cost of the computation of the value $K_t(i,i) = \mathbf{e}_i^\top K_t \mathbf{e}_i$ at k feature points varies from $O(kn)$ to $O(kn^2)$, according to the sparsity of the coefficient matrix. To speed-up the computation of the values of the heat diffusion kernel and distances for several values of t, or between a large number of points, it is sufficient to solve the linear systems in Eq. (8) with iterative methods or pre-factorizing the matrices (L, B) [15].

Approximation accuracy and numerical stability. Assuming exact arithmetic, the ℓ_2 approximation error between $\exp(-tC)$ and its rational approximation $c_{rr}(tC)$ is lower than the uniform rational Chebyshev constant σ_{rr} [36]. Since this constant is known, independent of t, and related to the degree of the rational Chebyshev polynomial by the relation $\sigma_{rr} \approx 10^{-r}$, $r := 7$ provides an error lower than 10^{-7}, which is satisfactory for the approximation of $K_t\mathbf{f}$ on 3D shapes. If necessary, a higher approximation accuracy is achieved by slightly increasing the degree r of the Chebyshev rational polynomial. According to [21], the Chebyshev approximation of the matrix $\exp(-tC)$ might be numerically unstable if $\|tC\|_2$ becomes large. From the upper bound $\|tB^{-1}L\|_2 \leq t\lambda_{\max}(L)\lambda^{-1}_{\min}(B)$, we get that a well-conditioned mass matrix B guarantees that $\|tB^{-1}L\|_2$ is bounded. These considerations and our experiments confirm that the Chebyshev method provides a good approximation accuracy and numerical stability for the computation of the discrete heat diffusion kernel.

4. Results and discussion

In the following, we represent the Laplace-Beltrami operator on triangle meshes and point sets in a unified way as $\tilde{L} := B^{-1}L$, where B is a positive definite matrix and L is symmetric, semi-positive definite. Then, we discuss the main features of the proposed approach in terms of smoothness, stability to shape discretization, approximation accuracy, and computational cost.

Laplacian matrix for triangle and polygonal meshes. Assuming that the input shape is discretized as a triangle mesh, whose set of vertices is $\mathcal{P} := \{\mathbf{p}_i\}_{i=1}^n$, the Laplacian matrix is defined as $\tilde{L} := B^{-1}L$, where L is the Laplacian matrix with cotangent weights and B is the diagonal matrix whose entries are the areas of the Voronoi regions of the mesh vertices (*Voronoi-cot weights*) [11]. Alternatively, B is the FEM mass matrix (*linear FEM weights*) [26, 35], which codes the variation of the triangle areas. On polygonal meshes, we consider the Laplacian discretization proposed in [1], which provides a generalization of the Laplacian matrix with cot-weights to surface meshes with non-planar, non-convex faces.

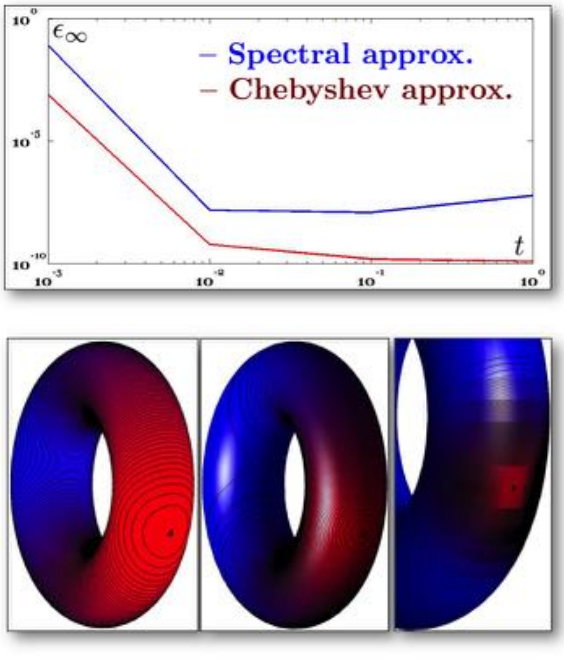

Figure 5: (a) Variation of the ℓ_∞ approximation error ϵ_∞ (y-axis) between the ground-truth diffusion distances and its approximation with (blue line) $k = 500$ Laplacian eigenpairs (red line) and the Chebyshev approximation ($r := 7$), with different values of t (x-axis). (b) Level sets of the diffusion distance from a source point (black dot) with different values of t.

Laplacian matrix for point sets. In [2, 3, 4], the Laplace-Beltrami operator for a function $f : \mathcal{P} \to \mathbb{R}$ defined on a point set $\mathcal{P} := \{\mathbf{p}_i\}_{i=1}^n$, is discretized as the linear operator $(\mathcal{L}f(\mathbf{p}_i))_{i=1}^n = L\mathbf{f}$, where the Laplacian matrix is

$$L(i,j) := \frac{1}{nt(4\pi t)^{3/2}} \begin{cases} \exp\left(-\frac{\|\mathbf{p}_i - \mathbf{p}_j\|_2}{4t}\right) & i \neq j, \\ -\sum_{k \neq i} \exp\left(-\frac{\|\mathbf{p}_i - \mathbf{p}_k\|_2}{4t}\right) & i = j \end{cases}$$

and $\mathbf{f} := (f(\mathbf{p}_i))_{i=1}^n$ is the array of the f-values at $\mathcal{P}$. We briefly recall that $L\mathbf{f}$ converges to Δf, on $\mathcal{P}$, as $t \to 0^+$.

Recently [19], a new discretization of the Laplace-Beltrami operator has been proposed using a finer approximation of the local geometry of the surface at each point through its Voronoi cell. More precisely, the Laplacian matrix is defined as $\tilde{L} := B^{-1}L$, where the entries of the symmetric matrix L and the diagonal matrix B are defined as follows

$$L(i,j) := \begin{cases} \frac{1}{4\pi t^2} \exp\left(-\frac{\|\mathbf{p}_i - \mathbf{p}_j\|_2^2}{4t}\right) & i \neq j, \\ -\sum_{j \neq i} L(i,j) & i = j, \end{cases} \qquad B(i,i) = v_i,$$

and v_i is the area of the Voronoi cell associated to the point $\mathbf{p}_i$. The Voronoi cell of $\mathbf{p}_i$ is approximated by projecting the points of a neighbor of $\mathbf{p}_i$ on the estimated tangent plane to $\mathcal{M}$ at $\mathbf{p}_i$. This discretization still converges to the Laplace-Beltrami operator of the underlying manifold, as the sampling density increases and t tends to zero.

Examples and discussion. For our tests, we consider the solution $K_t\mathbf{e}_i$ to the heat diffusion process, whose initial condition takes value 1 at the anchor point $\mathbf{p}_i$ and 0 otherwise. On irregularly sampled data, the linear FEM heat kernel (Figs. 1, 2) provides smooth level sets that are well-distributed around the anchor point $\mathbf{p}_i$; on the contrary, the Voronoi-cot heat kernel is more sensitive to the surface sampling. On noisy (Fig. 3) and partially-sampled data (Figs 4), the analogous behavior of the level sets and color maps also confirms the robustness of the linear FEM heat distances.

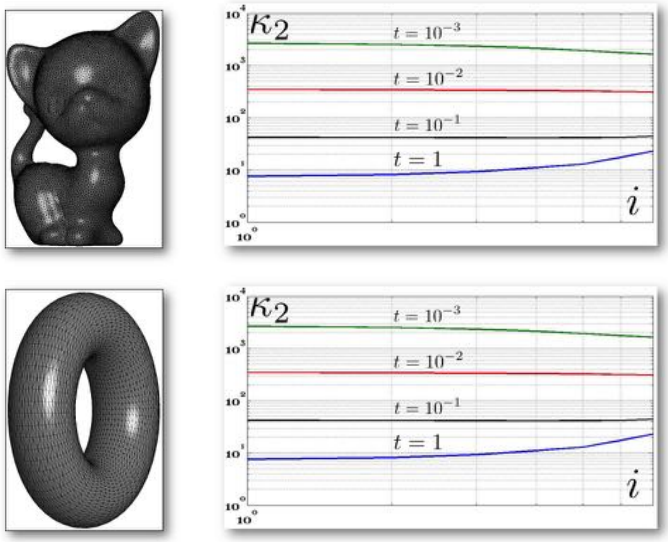

Figure 7: Conditioning number κ_2 (y-axis) of the matrices $\{(tL + \theta_i B)\}_{i=1}^7$ in Eq. (8), for several values the time parameter t; the indices of the coefficients $\{\theta_i\}_{i=1}^7$ are reported on the x-axis.

Fig. 5 shows the ℓ_∞ approximation error between the ground-truth heat diffusion distances from a source point and its approximation with $k = 500$ Laplacian eigenpairs and the Chebyshev approximation. While the approximation error of the diffusion distance with the same number of Laplacian eigenpairs decreases and becomes more sensible to local noise as t diminishes, the Chebyshev approximation provides a lower approximation error for any value of t. We have further analyzed the different accuracy (Fig. 6) of the spectral and Chebyshev approximation of the heat kernel by measuring the ℓ_∞ approximation error (y-axis) between the spectral representation of the heat kernel $K_t(\cdot,\cdot)$, computed using a different number k (x-axis) of eigenfunctions, and the corresponding Chebyshev approximation (c.f., Eq. (9)). For small values of t, the partial spectral representation requires a large number k of Laplacian eigenvectors to recover local details. For instance (Fig. 6(a,b)), selecting $1K$ eigenpairs the approximation error remains higher than 10^{-2}; in fact, local shape features encoded by $K_t(\cdot,\cdot)$ for a small t are recovered using the eigenvectors associated with high frequencies, thus requiring the computation of a large part of the Laplacian spectrum. For large values of t, increasing k strongly reduces the approximation error until it becomes almost constant and close to zero. In this case, the behavior of the heat kernel is mainly influenced by the Laplacian eigenvectors related to the eigenvalues of smaller magnitude. We conclude that the spectral representation generally requires a high num-

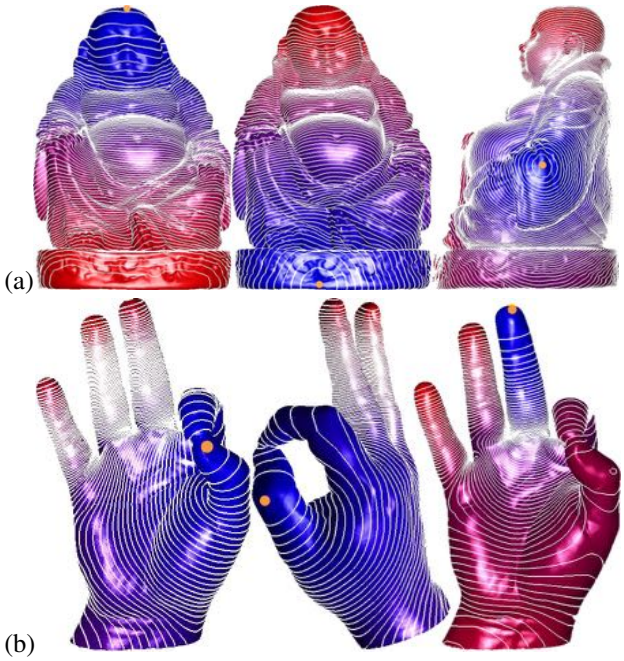

Figure 8: Level sets of the diffusion distance from a source point (orange) on (a) point set (150K points) and (b) a bordered (bottom). In (a), the level sets are rendered on the underlying triangle mesh. For the computation, we have used the Chebyshev approximation of order $r := 7$.

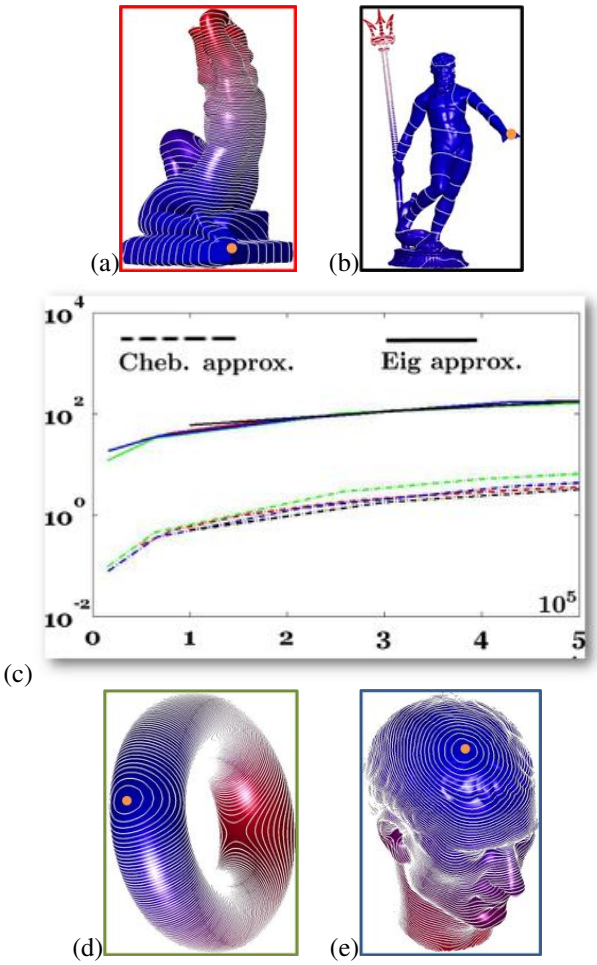

Figure 9: (c) Computational cost (in seconds, y-axis, log-scale) for the evaluation of the heat kernel on (a,b;d,e) 3D shapes with n samples (x-axis), approximated with (straight line) with $k = 500$ eigenpairs and the Chebyshev approximation (dotted line). Colors of diffusion distances from the source (orange) point vary from blue (null distance) to red (maximum distance). Timings are reported in Table 1.

ber of eigenpairs without achieving an accuracy of the same order of the spectrum-free approximation through the Chebyshev method. The proposed approach guarantees the smoothness of the heat diffusion distance at small and large scales and it is not affected by the irregular surface sampling.

According to Eq. (9), the value of t influences the conditioning number of the coefficient matrices $(tL + \theta_i B)$, $i = 1, \ldots, r$. Our experiments (Fig. 7) have shown that the linear systems in Eq. (8) are generally well-conditioned; in any case, preconditioners and regularization techniques [15] can be applied to attenuate numerical instabilities. Since our approach works mainly on matrices, the computation of the heat diffusion distance and kernel is independent of the discretization of the input surface as a manifold/non-manifold polygonal [1] mesh or a point cloud. The spectrum-free computation on point-sampled surfaces or non-manifold meshes (Fig. 8) is one of the novelties of the proposed approach with respect to previous work, which uses multi-resolutive and prolongation operators on manifold triangle meshes. Timings (Table 1, Fig. 9) are reduced from 20 up to 1200 times with respect to the approximation based on a fixed number of Laplacian eigenpairs. For the computation of the Laplacian eigenvectors, we have used the Arnoldi iteration method [18, 31]. Selecting $B := I$ or $B := D$, the Chebyshev method also provides a new computation of the discrete heat diffusion kernel associated to the Laplacian matrix with cot and Voronoi-cot weights.

5. Conclusions and future work

This paper has presented an efficient and spectrum-free computation of the heat diffusion distance and kernel through the rational Chebyshev approximation, which involves the solution of a set of sparse, symmetric, well-conditioned linear systems and a sequence of matrix-vector multiplications. With respect to previous work, we provide an efficient computation on both triangle meshes and point-sampled surfaces, thus avoiding the computation of the Laplacian spectrum, the selection of a specific subset of eigenpairs, and the use of multiresolutive prolongation operators. Furthermore, the Chebyshev computation is efficient, accurate, and robust to noise, missing and irregularly sampled areas. As future work, we foresee the generalization of the proposed approach to a larger class of spectral distances.

Acknowledgements. Special thanks are given to the anonymous Reviewers for their valuable comments. 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

Table 1: Timings (in seconds, Fig. 9(c)) for the evaluation of the heat diffusion kernels on 3D shapes with n points, approximated with $k = 500$ eigenpairs *(Eigs)* and the Chebyshev approximation *(Cheb.)*. Column '×' 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.

Acquarius Fig. 9(a)				**Neptune Fig. 9(b)**			
n	Eigs	Cheb.	×	n	Eigs	Cheb.	×
5K	30.06	0.26	115	10K	59.65	0.50	119
25K	97.25	1.83	53	30K	111.28	1.78	62
35K	130.39	2.61	49	50K	176.47	3.21	54
50K	173.78	3.60	48	100K	372.16	7.44	50
Torus Fig. 9(c)				**Julius Fig. 9(d)**			
n	Eigs	Cheb.	×	n	Eigs	Cheb.	×
2K	12.00	0.01	1200	2K	18.47	0.08	230
6K	33.28	0.45	73	7K	35.89	0.37	97
26K	100.88	2.89	34	22K	82.47	1.42	58
49K	140.00	5.14	33	43K	173.52	3.71	46
58K	186.06	7.92	23	50K	174.89	4.34	40

Modalities for Physiological Human Articulation". 3D shapes are courtesy of the AIM@SHAPE Repository.

References

[1] M. Alexa and M. Wardetzky. Discrete laplacians on general polygonal meshes. *ACM Transactions on Graphics*, 30(4), 2011.

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

[3] M. Belkin and P. Niyogi. Convergence of laplacian eigenmaps. In *NIPS*, pages 129–136, 2006.

[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. *International 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 Transactions on Graphics*, 30(1), 2011.

[7] O. Chapelle, J. Weston, and B. Schölkopf. Cluster Kernels for Semi-Supervised Learning. In *Neural Information Processing Systems*, volume 15, pages 585–592, 2003.

[8] W. J. Cody, G. Meinardus, and R. S. Varga. Chebyshev rational approximations to $\exp(-z)$ in $(0, +\infty)$ and applications to heat-conduction problems. *Journal of Approximation Theory*, 2:50–65, 1969.

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

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

[11] 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.

[12] M. Desbrun, M. Meyer, P. Schroeder, and A. H. Barr. Discrete differential-geometry operators for triangulated 2-manifolds. *VisMath*, '02:35–57, 2002.

[13] E. Gallopoulos and Y. Saad. Efficient solution of parabolic equations by krylov approximation methods. *SIAM Journal of Scientific Statistical Computation*, 13:1236–1264, 1992.

[14] 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.

[15] G. Golub and G.F. VanLoan. *Matrix Computations*. John Hopkins University Press, 2nd Edition, 1989.

[16] 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.

[17] 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.

[18] R.B. Lehoucq and D. C. Sorensen. Deflation techniques for an implicitly re-started arnoldi iteration. *SIAM J. Matrix Anal. Appl*, 17:789–821, 1996.

[19] Y. Liu, B. Prabhakaran, and X. Guo. Point-based manifold harmonics. *IEEE Transactions on Visualization and Computer Graphics*, 18(10):1693 –1703, 2012.

[20] F. Memoli. Spectral Gromov-Wasserstein distances for shape matching. In *Workshop on Non-Rigid Shape Analysis and Deformable Image Alignment*, pages 256–263, 2009.

[21] C. Moler and C. Van Loan. Nineteen dubious ways to compute the exponential of a matrix, twenty-five years later. *SIAM Review*, 45(1):3–49, 2003.

[22] 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.

[23] 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.

[24] G. Patanè and B. Falcidieno. Multi-scale feature spaces for shape processing and analysis. In *Proc. of Shape Modeling International*, pages 113–123, 2010.

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

[26] 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.

[27] S. T. Roweis and L. K. Saul. Nonlinear dimensionality reduction by locally linear embedding. *Science*, 290:2323–2326, 2000.

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

[29] J. Shi and J. Malik. Normalized cuts and image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 22(8):888–905, 2000.

[30] A. J. Smola and R. I. Kondor. Kernels and regularization on graphs. In *Conference on Learning Theory*, pages 144–158, 2003.

[31] D. C. Sorensen. Implicit application of polynomial filters in a k-step arnoldi method. *SIAM J. Matrix Anal. Appl.*, 13(1):357–385, January 1992.

[32] 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.

[33] 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.

[34] J. B. Tenenbaum, V. Silva, and J. C. Langford. A Global Geometric Framework for Nonlinear Dimensionality Reduction. *Science*, 290(5500):2319–2323, 2000.

[35] B. Vallet and B. Lévy. Spectral geometry processing with manifold harmonics. *Comput. Graph. Forum*, 27(2):251–260, 2008.

[36] R.S. Varga. *Scientific computation on mathematical problems and conjectures*. SIAM, 1990.

[37] 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.

[38] Y. Wang. Approximating gradients for meshes and point clouds via diffusion metric. *Computer Graphics Forum*, 28:1497–1508(12), 2009.

[39] 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.

[40] 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.