SIGGRAPH 2017 – Course “Laplacian Spectral Kernels and Distances: Theory, Computation, and Applications”

G. Patanè (CNR-IMATI, Italy)

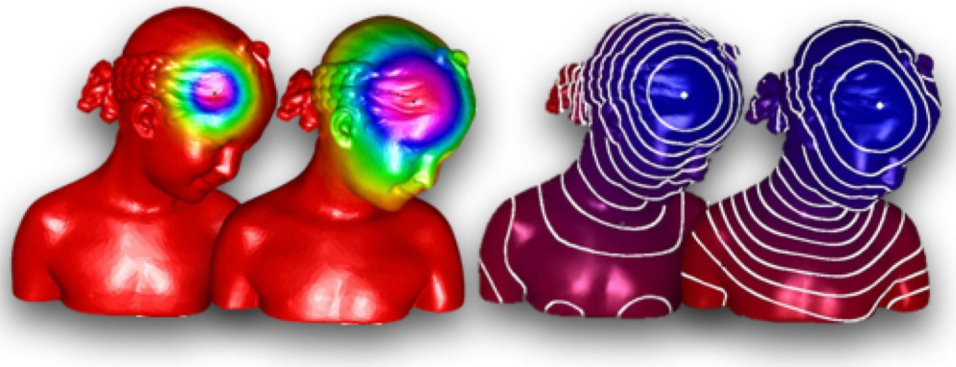

Table of Contents

SIGGRAPH 2017 – Course “Laplacian Spectral Kernels and Distances: Theory, Computation, and Applications”

G. Patanè (CNR-IMATI, Italy)

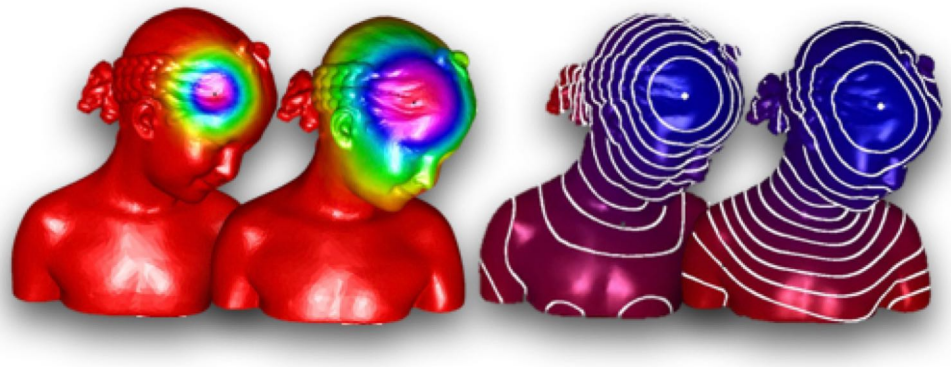

Course Overview

SIGGRAPH 2017 – Short Course
An Introduction to Laplacian Spectral Kernels and Distances: Theory, Computation, and Applications

Giuseppe Patané*

1 Course description

In geometry processing and shape analysis, several applications have been addressed through the properties of the spectral kernels and distances, such as commute-time, biharmonic, diffusion, and wave distances. Our course is intended to provide a background on the properties, discretization, computation, and main applications of the Laplace-Beltrami operator, the associated differential equations (e.g., harmonic equation, Laplacian eigenproblem, diffusion and wave equations), the Laplacian spectral kernels and distances (e.g., commute-time, biharmonic, wave, diffusion distances). While previous work has been focused mainly on specific applications of the aforementioned topics on surface meshes, we propose a general approach that allows us to review the Laplacian kernels and distances on surfaces and volumes, and for any choice of the Laplacian weights. All the reviewed numerical schemes for the computation of the Laplacian spectral kernels and distances are discussed in terms of robustness, approximation accuracy, and computational cost, thus supporting the reader in the selection of the most appropriate method with respect to shape representation, computational resources, and target applications.

Part I - Introduction

We present the outline and the main aims of this course on the theory, computation, and applications of the Laplacian spectral distances and kernels.

Part II - Laplace-Beltrami operator on surfaces and volumes

Firstly, we define a unified representation of the isotropic and anisotropic discrete Laplacian on surfaces and volumes; then, we introduce the associated differential equations. For the harmonic equation and the Laplacian eigenproblem, we focus on the stability and accuracy of numerical solvers, also presenting their main applications.

Part III - Heat equation and diffusion distances

Part II provides the background for a detailed analysis of the heat equation and allows us to identify the main limitations (e.g., computational cost, storage overhead, selection of user-defined parameters) of previous work on the approximation of the diffusion distances, which is based mainly on the evaluation of the Laplacian spectrum and on linear approximations of the exponential matrix. For the heat equation, we discuss the selection of the time scale and the main approaches for the computation of the solution to the heat equation, such as linear, polynomial, and rational approximations.

Part IV - Laplacian spectral kernels and distances

Filtering the Laplacian spectrum, we introduce the Laplacian spectral distances, which generalize the commute-time, biharmonic, diffusion and wave distances, and their discretization in terms of the Laplacian spectrum. The growing interest on these distances is motivated by their capability of encoding local geometric properties (e.g., Gaussian curvature, geodesic distance) of the input shape, their intrinsic and multi-scale definition with respect to the input shape, their invariance to isometries, shape-awareness, robustness to noise and tessellation. While previous work has been focused mainly on surfaces discretized as triangle meshes, we introduce a unified representation of the spectral distances and kernels, which is independent of the selected Laplacian weights, of the surface or volume representation as polygonal mesh, point set, tetrahedral or voxel grids. From this general representation, we show that the main properties of the spectral distances are guided mainly by the filter that is applied to the Laplacian eigenpairs.

The expensive cost for the computation of the Laplacian spectrum and the sensitiveness of multiple Laplacian eigenvalues to surface discretization generally preclude an accurate evaluation of the spectral kernels and distances on large data sets. To discuss these problems, we review and compare different methods for the numerical evaluation of the spectral distances and kernels. In particular, we detail their *spectrum-free computation*, which is defined through a polynomial or rational approximation of the filter function. The resulting computational scheme only requires the solution of sparse linear systems, is not affected by the Gibbs phenomenon, is independent of the representation of the input domain, the selected Laplacian weights, and the evaluation of the Laplacian spectrum.

Part V - Conclusions

As main applications, we detail the Laplacian smoothing and the definition of basis functions for geometry processing and shape analysis. Finally, we conclude our review with a discussion of open questions and challenges.

2 Course schedule

Part I – Introduction (10 *min.*)

1. Outline and motivations
2. Goals and contributions

Part II - Laplace-Beltrami operator on surfaces and volumes (15 *min.*)

1. Laplacian matrix and eigenproblem
2. Laplacian spectrum: computation and properties

Part III - Heat equation and diffusion distances (30 *min.*)

1. Heat diffusion equation
2. Diffusion kernel and distances: definition and properties
3. Computation of the diffusion kernel and distances

*Consiglio Nazionale delle Ricerche, Istituto di Matematica Applicata e Tecnologie Informatiche, Genova, Italy, patane@ge.imati.cnr.it

- Geometry-driven approaches: multi-resolution prolongation operator
- Spectrum-based approaches: truncated spectral approximation, Euler backward method, and power method
- Spectrum-free approaches: Padè-Chebyshev approximation, polynomial approximation, and Krylov subspace projection

4. Discussion
   - Approximation accuracy and stability
   - Robustness with respect to noise, discretisation, geometric and topologucal noise
   - Computational cost
5. Main applications
   - Approximation of geodesic and optimal transportation distances *via* heat kernel
   - Diffusion distances for shape analysis and manifold learning

Part IV - Laplacian spectral kernels and distances (25 *min.*)

1. Distance definition and properties: geometry-driven, functional, and mixed approaches
2. Laplacian spectral distances
   - Equivalent formulations: spectral operator, distance, and embedding
   - Filter selection and distance properties: smoothness, locality, and shape encoding
3. Main examples: commute-time, bi-harmonic, wave and diffusion distances
4. Discretisation and computation
   - Truncated spectral computation
   - Spectrum-free computation: polynomial and rational approaches
5. Main applications: spectral signatures for shape comparison

Part V - Conclusions (10 *min.*)

- Conclusions, Questions & Answers

3 Course Rationale

Target audience The target audience of this tutorial includes graduate students and researchers interested in numerical geometry processing and spectral shape analysis. Our course is intended to:

- present a *unified definition of the Laplacian spectral kernels and distances* with respect to the dimensionality and discretisation of the input domain, the discretisation of the Laplace-Beltrami operator, and the selected filter. In this way, we will provide a unified view on the definition and computation of well-known distances, such as random walks, heat diffusion, biharmonic, and wave kernel distances;
- provide a *common background* for those research areas and applications that apply the Laplacian spectral kernels and distances to geometry processing and shape analysis. In particular, shape segmentation and comparison with multi-scale and isometry-invariant signatures, diffusion geometry, dimensionality reduction with spectral embeddings, data visualisation, representation, and classification;
- introduce and discuss the *properties and applications of the Laplacian spectral kernels and distances* that are relevant for shape modelling and, more generally, computer graphics;
- discuss open problems and applications.

Prerequisites Knowledge about linear algebra, discrete geometry processing, and computer graphics.

Level of difficulty: Intermediate course.

Tutorial originality While previous courses have been focused mainly on the application of the Laplacian spectral kernels and distances, our focus is on their definition and computation, thus addressing their main common aspects and properties. Indeed, it is the first course that systematically presents the theory, algorithm, and applications of these topics.

Related tutorials organised by the lecturer This course proposal revises and extends our Eurographics 2016 STAR "*Laplacian Spectral Kernels and Distances for Geometry Processing and Shape Analysis*", which was attended by more than 120 participants. According to recent results of the authors and the feedback to the previous course, this SIGGRAPH course will include additional material on the definition of the Laplacian spectral kernels and distances in a more general setting, which will allow us to review a larger spectrum of research areas (e.g., graph theory, spectral geometry processing, manifold learning) and applications (e.g., shape analysis and comparison).

Previous tutorials on related topics Course C1 reviewed the main methods for the computation of the correspondences between geometric shapes. Tutorial C2 has addressed the definition and application of diffusion distances to shape analysis and comparison. Course C3 focused on the discrete exterior calculus and its relation with digital geometry processing and discrete differential geometry. Course C4 presented the main concepts behind spectral mesh processing on 3D shapes and its applications to filtering, shape matching, remeshing, segmentation, and parameterisation. Indeed, our tutorial is complementary to previous work and provides a common background for previous tutorial and research papers.

(C1) SIGGRAPH Asia 2016 Courses "*Computing and Processing Correspondences with Functional Maps*", M. Ovsjanikov, E. Corman, M. M. Bronstein, E. Rodolá, M. Ben-Chen, L. Guibas, F. Chazal, and A. M. Bronstein;

(C2) Eurographics Tutorial 2012, A. "*Diffusion geometry in shape analysis*, Bronstein, M., Castellani, U., Bronstein, A.;

(C3) SIGGRAPH'2013 Course "*Geometry Processing with Discrete Exterior Calculus* " (F. de Goes, K. Crane, M. Desbrun, P. Schroeder);

(C4) SIGGRAPH Asia'2010 "*Spectral Geometry Processing*" (B. Levy, R. H. Zhang).

4 Lecturer biography

Giuseppe Patanè

Affiliation	CNR-IMATI, Genova, Italy
e-mail	patanel@ge.imati.cnr.it
URL	http://www.ge.imati.cnr.it

Giuseppe Patanè is researcher at CNR-IMATI (2001-today). He received a Ph.D. in "Mathematics and Applications" from the University of Genova (2005) and a Post Lauream Degree Master from the "F. Severi National Institute for Advanced Mathematics" (2000). From 2001, his research activities have been focused on *numerical geometry processing*, *the modelling and analysis of 3D shapes and multi-dimensional data*, with applications to computer graphics and bio-medicine. Since 2013, he has organised the following SIGGRAPH and Eurographics courses

- Eurographics STAR 2016 "*Laplacian Spectral Kernels and Distances for Geometry Processing and Shape Analysis*" (G.Patanè);
- SIGGRAPH Asia 2014 Course "*An Introduction to Ricci Flow and Volumetric Approximation with Applications to Shape Modeling*" (G.Patanè, X.D. Gu, X.S. Li);
- SIGGRAPH Asia 2013 Course "*Surface-Based and Volume-Based Techniques for Shape Modeling and Analysis*" (G. Patanè, X.S. Li, X.D. Gu).

Since 2008, he was speaker of the following courses

- Shape Modeling International'2012 Tutorial "*Spectral, Curvature Flow Surface-Based and Volume-Based Techniques for Shape Modeling and Analysis*" (G. Patanè, X.D. Gu, X.S. Li, M. Spagnuolo);
- Eurographics 2007 Tutorial "*3D shape description and matching based on properties of real functions*" (S. Biasotti, B. Falcidieno, P. Frosini, D. Giorgi, C. Landi, S. Marini, G. Patanè, M. Spagnuolo);
- ICIAM2007 Mini-Symposium "*Geometric-Topological Methods for 3D Shape Classification and Matching*" (M. Spagnuolo, G. Patanè);
- SMI'2008 Mini-Symposium on "*Shape Understanding via Spectral Analysis Techniques*" (B. Levy, R. Zhang, M. Retuer, G. Patanè, M. Spagnuolo).

For more information, we refer to the personal web-page: http://pers.ge.imati.cnr.it/patane/Home.html.

Main author publications on the course topics

- Patanè G., *Accurate and Efficient Computation of Laplacian Spectral Distances and Kernels*. In: Computer Graphics Forum. In press, 2017.
- Patanè G., "*Laplacian Spectral Kernels and Distances for Geometry Processing and Shape Analysis*", STAR-State-of-the-Art Report. In: Computer Graphics Forum, 35(2): 599-624 (2016).
- Patanè G., *Volumetric Heat Kernel: Padè-Chebyshev Approximation, Convergence, and Computation*. In: Computer&Graphics, Volume 46, February 2015, pp. 64-71.
- Patanè G., *Diffusive Smoothing of 3D Segmented Medical Data*. In: Journal of Advanced Research, Elsevier, Volume 6, Issue 3, May 2015, pp. 425-431.
- Patanè G., *Laplacian spectral distances and kernels on 3D shapes*. In: Pattern Recognition Letters 47, pp. 102-110 (2014).
- Patanè G., *wFEM Heat Kernel: Discretization and Applications to Shape Analysis and Retrieval*. In: Computer Aided Geometric Design, Vol. 30, Issue 3, March 2013, pp. 276-295.
- Patanè G., Spagnuolo M., *An Interactive Analysis of Harmonic and Diffusion Equations on Discrete 3D Shapes*. In: Computer & Graphics, Vol. 37, Issue 5, August 2013, pp. 526-538.
- Patanè G., Spagnuolo M., *Heat Diffusion Kernel and Distance on Surface Meshes and Point Sets*. In: Computer & Graphics, Vol. 37, Issue 6, October 2013, pp. 676-686.

SIGGRAPH 2017 – Course “Laplacian Spectral Kernels and Distances: Theory, Computation, and Applications”

G. Patanè (CNR-IMATI, Italy)

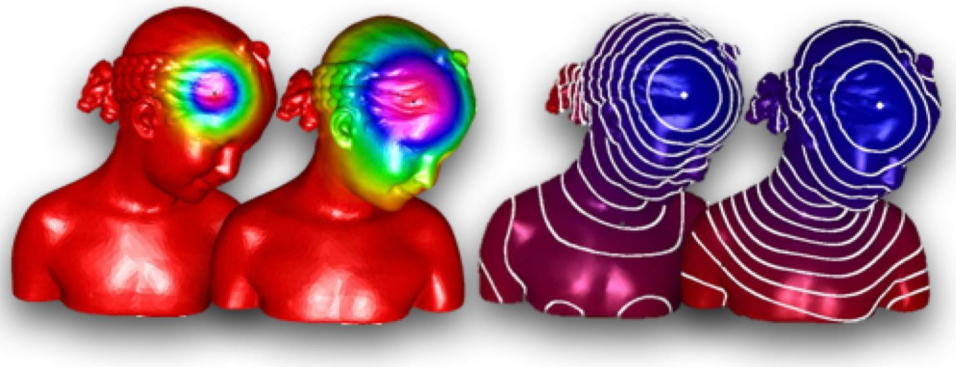

Course Notes

Contents

SIGGRAPH 2017 – Short Course
An Introduction to Laplacian Spectral Distances and Kernels: Theory, Computation, and Applications

Giuseppe Patané*

Abstract

In geometry processing and shape analysis, several applications have been addressed through the properties of the spectral kernels and distances, such as commute-time, biharmonic, diffusion, and wave distances. Our survey is intended to provide a background on the properties, discretization, computation, and main applications of the Laplace-Beltrami operator, the associated differential equations (e.g., harmonic equation, Laplacian eigenproblem, diffusion and wave equations), Laplacian spectral kernels and distances (e.g., commute-time, biharmonic, wave, diffusion distances). While previous work has been focused mainly on specific applications of the aforementioned topics on surface meshes, we propose a general approach that allows us to review Laplacian kernels and distances on surfaces and volumes, and for any choice of the Laplacian weights. All the reviewed numerical schemes for the computation of the Laplacian spectral kernels and distances are discussed in terms of robustness, approximation accuracy, and computational cost, thus supporting the reader in the selection of the most appropriate method with respect to shape representation, computational resources, and target application.

1 Introduction

In geometry processing and shape analysis, several applications have been addressed through the properties of the spectral kernels and distances, such as commute-time, biharmonic, diffusion, and wave distances. Spectral distances are easily defined through a filtering of the Laplacian eigenpairs and include random walks [Fouss et al. 2005; Ramani and Sinha 2013], heat diffusion [Bronstein et al. 2010a; Bronstein et al. 2011; Coifman and Lafon 2006; Gebal et al. 2009; Lafon et al. 2006; Luo et al. 2009], biharmonic [Lipman et al. 2010; Rustamov 2011b], and wave kernel [Bronstein and Bronstein 2011b; Aubry et al. 2011] distances. Laplacian spectral distances have been applied to shape segmentation [de Goes et al. 2008] and comparison [Bronstein et al. 2011; Gebal et al. 2009; Memoli 2009; Ovsjanikov et al. 2010; Sun et al. 2009] with multi-scale and isometry-invariant signatures [Dey et al. 2010b; Lafon et al. 2006; Mèmoli and Sapiro 2005; Memoli 2011; Raviv et al. 2010; Rustamov 2007; Mahmoudi and Sapiro 2009]. In fact, they are intrinsic to the input shape, invariant to isometries, multi-scale, and robust to noise and tessellation. Biharmonic [Lipman et al. 2010; Rustamov 2011b] and diffusion [Bronstein et al. 2010a; Bronstein et al. 2011; Coifman and Lafon 2006; Gebal et al. 2009; Lafon et al. 2006; Luo et al. 2009; Patanè and Spagnuolo 2013b] distances provide a trade-off between a nearly geodesic behavior for small distances and the encoding of global surface properties for large distances, thus guaranteeing an intrinsic and multi-scale characterization of the input shape. The heat kernel [Berard et al. 1994] is also central in diffusion geometry [Belkin and Niyogi 2003; Coifman and Lafon 2006; Gine and Koltchinskii 2006; Singer 2006], dimensionality reduction with spectral embeddings [Belkin and Niyogi 2003; Xiao et al. 2010], and data classification [Smola and Kondor 2003]. As main applications, we mention the multi-scale approximation of functions [Patanè and Falcidieno 2010] and gradients [Luo et al. 2009], shape segmentation and comparison through heat kernel shape descriptors, auto-diffusion functions, and diffusion distances. The diffusion kernel and distance also play a central role in several applications, such as dimensionality reduction with spectral embeddings [Belkin and Niyogi 2003; Xiao et al. 2010]; data visualization [Belkin and Niyogi 2003; Hein et al. 2005; Roweis and Saul 2000; Tenenbaum et al. 2000], representation [Chapelle et al. 2003; Smola and Kondor 2003; Zhu et al. 2003], and classification [Ng et al. 2001; Shi and Malik 2000; Spielman and Teng 2007].

Course topics and contributions Our survey is intended to provide a common background on the definition and computation of Laplacian spectral kernels and distances for geometry processing and shape analysis. All the reviewed numerical schemes are discussed and compared in terms of robustness, approximation accuracy, and computational cost, thus supporting the reader in the selection of the most appropriate with respect to shape representation, computational resources, and target application. Indeed, our review is complementary to previous work, which has been focused mainly on specific applications, such as mesh filtering [Taubin 1999], surface coding and spectral partitioning [Karni and Gotsman 2000], 3D shape deformation based on differential coordinates [Sorkine 2006], spectral methods [Zhang et al. 2007] and Laplacian eigenfunctions [Levy 2006] for geometry processing and diffusion shape analysis [Bronstein et al. 2012].

Firstly, we define a unified representation of the isotropic and anisotropic discrete Laplacian on surfaces and volumes (Sect. 2); then, we introduce the associated differential equations. For the harmonic equation (Sect. 3) and the Laplacian eigenproblem (Sect. 4), we focus on the stability and accuracy of numerical solvers, also presenting their main applications. This discussion provides the background for a detailed analysis of the heat equation (Sect. 5) and allows us to identify the main limitations (e.g., computational cost, storage overhead, selection of user-defined parameters) of previous work on the approximation of the diffusion distances, which is based mainly on the evaluation of the Laplacian spectrum and on linear approximations of the exponential matrix. For the heat equation, we discuss the selection of the time scale and the main approaches for the computation of the solution to the heat equation, such as linear, polynomial, and rational approximations.

Filtering the Laplacian spectrum, we introduce the Laplacian spectral distances (Sect. 6), which generalize the commute-time, biharmonic, diffusion and wave distances, and their discretization in terms of the Laplacian spectrum. The growing interest on these distances is motivated by their capability of encoding local geometric properties (e.g., Gaussian curvature, geodesic distance) of the input shape, their intrinsic and multi-scale definition with respect to the input shape, their invariance to isometries, shape-awareness, robustness to noise and tessellation. While previous work has been focused mainly on surfaces discretized as triangle meshes, we introduce a unified representation of the spectral distances and kernels, which is independent of the selected Laplacian weights, of the surface or volume representation as polygonal mesh, point set, tetrahedral or voxel grid. From this general representation, we show that the main properties of the spectral distances are guided mainly by

*Consiglio Nazionale delle Ricerche, Istituto di Matematica Applicata e Tecnologie Informatiche, Genova, Italy, patane@ge.imati.cnr.it

the filter that is applied to the Laplacian eigenpairs.

The expensive cost for the computation of the Laplacian spectrum and the sensitiveness of multiple Laplacian eigenvalues to surface discretization generally preclude an accurate evaluation of the spectral kernels and distances on large data sets. To discuss these problems, we review and compare different methods for the numerical evaluation of the spectral distances and kernels. In particular, we detail their *spectrum-free computation*, which is defined through a polynomial or rational approximation of the filter function. The resulting computational scheme only requires the solution of sparse linear systems, is not affected by the Gibbs phenomenon, is independent of the representation of the input domain, the selected Laplacian weights, and the evaluation of the Laplacian spectrum.

As main applications (Sect. 7), we detail the Laplacian smoothing and the definition of basis functions for geometry processing and shape analysis. Finally (Sect. 8), we conclude our review with a discussion of open questions and challenges.

2 Laplace-Beltrami operator and related equations

We review the isotropic and anisotropic Laplace-Beltrami operators and introduce a unified representation of the corresponding Laplacians for surfaces and volumes. Additional results have been presented in [Sorkine 2006; Taubin 1999; Karni and Gotsman 2000; Zhang et al. 2007].

Let $\mathcal{N}$ be a smooth surface, possibly with boundary, equipped with a Riemannian metric and let us consider the *scalar product* $\langle f,g\rangle_2 := \int_{\mathcal{N}} f(\mathbf{p})g(\mathbf{p})\mathrm{d}\mathbf{p}$ defined on the space $\mathscr{L}^2(\mathcal{N})$ of square integrable functions on $\mathcal{N}$ and the corresponding norm $\|\cdot\|_2$. Then, the *intrinsic smooth Laplace-Beltrami operator* $\Delta := -\mathrm{div}(\mathrm{grad})$ satisfies the following properties [Rosenberg 1997]:

- *self-adjointness*: $\langle \Delta f, g\rangle_2 = \langle f, \Delta g\rangle_2$, $\forall f,g$;
- *positive semi-definiteness*: $\langle \Delta f, f\rangle_2 \geq 0$, $\forall f$. In particular, the Laplacian eigenvalues are positive;
- *null eigenvalue*: the smallest Laplacian eigenvalue is null and the corresponding eigenfunction ϕ, $\Delta\phi = 0$, is constant;
- *locality*: the value $\Delta f(\mathbf{p})$ does not depend on $f(\mathbf{q})$, for any couple of distinct points $\mathbf{p}$, $\mathbf{q}$;
- *linear precision*: if $\mathcal{N}$ is planar and f is linear, then $\Delta f = 0$.

The *anisotropic Laplace-Beltrami operator* [Andreux et al. 2014] is defined as $\Delta_{\mathbf{D}} f = \mathrm{div}(\mathbf{D}\nabla f)$, where $\mathbf{D}$ is a 2×2 matrix applied to vectors belonging to the tangent plane and controls the direction and strength of the deviation from the isotropic case. The tensor $\mathbf{D} := \mathrm{diag}(\varphi_\alpha(\kappa_m), \varphi_\alpha(\kappa_M))$ takes into account the directions and the values κ_m, κ_M of low and high curvature, where the filter is $\varphi_\alpha(s) := (1+\alpha|s|)^{-1}$, $\alpha > 0$. As $\alpha \to 0$, we get the isotropic Laplace-Beltrami operator (i.e., $\mathbf{D} := \mathbf{I}$). The alternative definition [Kim et al. 2013] of the anisotropic Laplace-Beltrami operator applies a non-linear factor $\mathbf{D}(\mathbf{v})$, which modifies the magnitude of $\mathbf{D}(\mathbf{v})$ without changing its direction.

We now introduce a unified representation of the Laplacian matrix on surfaces and volumes, which is independent of the underlying discretization.

Discrete Laplacians and spectral properties Let us consider a (triangular, polygonal, volumetric) mesh $\mathscr{M} := (\mathscr{P}, T)$, which discretizes a domain $\mathcal{N}$, where $\mathscr{P} := \{\mathbf{p}_i\}_{i=1}^n$ is the set of n vertices and T is the connectivity graph (Fig. 1). On $\mathscr{M}$, a piecewise linear scalar function $f : \mathscr{M} \to \mathbb{R}$ is defined by linearly interpolating the values $\mathbf{f} := (f(\mathbf{p}_i))_{i=1}^n$ of f at the vertices using barycentric coordinates. For point sets, f is defined only at $\mathscr{P}$ and T is the k-nearest neighbor graph.

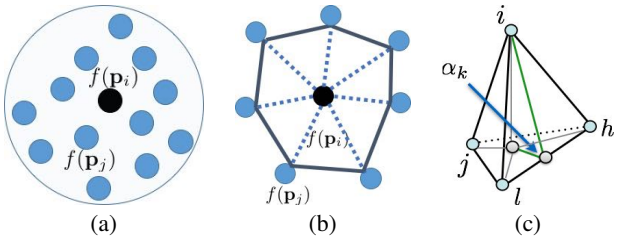

Figure 1: *Neighbor and Laplacian stencil for a (a) point set, (b) triangle and (c) tetrahedral mesh.*

We represent the Laplace-Beltrami operator on surface and volume meshes in a unified way as $\tilde{\mathbf{L}} := \mathbf{B}^{-1}\mathbf{L}$, where $\mathbf{B}$ is a sparse, symmetric, positive definite matrix (*mass matrix*) and $\mathbf{L}$ is sparse, symmetric, and positive semi-definite (*stiffness matrix*). We also assume that the entries of $\mathbf{B}$ are positive and that the sum of each row of $\mathbf{L}$ is null. In particular, we consider the $\mathbf{B}$-scalar product $\langle \mathbf{f},\mathbf{g}\rangle_{\mathbf{B}} := \mathbf{f}^\top\mathbf{B}\mathbf{g}$ and the induced norm $\|\mathbf{f}\|_{\mathbf{B}}^2 := \mathbf{f}^\top\mathbf{B}\mathbf{f}$. Analogously to the continuous case, the Laplacian matrix satisfies the following properties.

- *self-adjointness*: $\tilde{\mathbf{L}}$ is adjoint with respect to the $\mathbf{B}$-scalar product; i.e., $\langle \tilde{\mathbf{L}}\mathbf{f},\mathbf{g}\rangle_{\mathbf{B}} = \langle \mathbf{f},\tilde{\mathbf{L}}\mathbf{g}\rangle_{\mathbf{B}} = \mathbf{f}^\top\mathbf{L}\mathbf{g}$. If $\mathbf{B} := \mathbf{I}$, then this property reduces to the symmetry of $\mathbf{L}$;
- *positive semi-definiteness*: $\langle \tilde{\mathbf{L}}\mathbf{f},\mathbf{f}\rangle_{\mathbf{B}} = \mathbf{f}^\top\mathbf{L}\mathbf{f} \geq 0$. In particular, the Laplacian eigenvalues are positive;
- *null eigenvalue*: by construction, we have that $\tilde{\mathbf{L}}\mathbf{1} = \mathbf{0}$;
- *locality*: since the weight $w(i,j)$ is not null for each edge (i,j), the value $(\tilde{\mathbf{L}}\mathbf{f})_i$ depends only on the f-values at $\mathbf{p}_i$ and its 1-star neighbor $\mathcal{N}(i) := \{j : (i,j) \text{ edge}\}$.

For a detailed discussion of these properties with respect to the selected Laplacian weights, we refer the reader to [Wardetzky et al. 2007].

Laplacian matrix on graphs, triangle and polygonal meshes Associating a set $\{w(i,j)\}_{i,j}$ of positive weights with the edges (i,j) of T, the entries of the stiffness matrix are defined as $L(i,j) = -\sum_{k\neq i} w(i,k) + w(i,j)$. The entries of the mass matrix $\mathbf{B}$ are normalization coefficients that take into account the geometry of the input domain.

On *graphs* [Chung 1997], the weights of the stiffness matrix are equal to 1 for each edge and zero otherwise; each diagonal entry of the mass matrix is equal to the valence of the corresponding node. On *triangle meshes*, the *stiffness matrix* $\mathbf{L}$ and the *mass matrix* $\mathbf{B}$ of the linear FEM Laplacian weights [Reuter et al. 2006; Vallet and Lèvy 2008] are defined as

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

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

where $N(i)$ is the 1-star of the vertex i; α_{ij}, β_{ij} are the angles opposite to the edge (i,j) (Fig. 1b); t_r, t_s are the triangles that share the

edge (i,j); and $|t|$ is the area of the triangle t. Lumping the mass matrix $\mathbf{B}$ to the diagonal matrix $\mathbf{D}$, $D(i,i) = \frac{1}{3}\sum_{t \in N(i)} |t|$, whose entries are the areas of the Voronoi regions, $\tilde{\mathbf{L}}$ reduces to the Laplacian matrix $\mathbf{D}^{-1}\mathbf{L}$ with *Voronoi-cotangent weights* [Desbrun et al. 1999], which extend the cotangent weights introduced in [Pinkall and Polthier 1993] ($\mathbf{B} := \mathbf{I}$). The *mean-value weights* [Floater 2003] have been derived from the mean value theorem for harmonic functions and are always positive. In [Chuang et al. 2009], the weak formulation of the Laplacian eigenproblem is achieved by selecting a set of volumetric test functions, which are defined as $k \times k \times k$ B-splines (e.g., $k := 4$) and restricted to the input shape. For the *anisotropic Laplacian* [Andreux et al. 2014], the entries of $\mathbf{L}$ are a variant of the cotangent weights (i.e., with respect to different angles) and the entries of the diagonal mass matrix $\mathbf{B}$ are the areas of the Voronoi regions.

While the Laplace-Beltrami operator depends only on the Reimannian metric (*intrinsic property*), its discretization is generally affected by the quality of the input triangulation [Shewchuk 2002; Hildebrandt et al. 2006]. For instance, two (simplicial) isometric surfaces with two different triangulations are associated with two different Laplacian matrices. According to [Bobenko and Springborn 2007], the cotangent weights are non-negative if and only if the input triangulation is Delaunay and the corresponding Laplacian matrix is more accurate than the one evaluated on the original mesh. We briefly recall [Dyer et al. 2007; Liu et al. 2015a; Liu et al. 2015b] that a triangulation of a piecewise flat surface is a Delaunay triangulation if and only if all its interior edges are locally Delaunay (i.e., the sum of the angles opposite to an edge in the adjacent triangles does not exceed π). Furthermore, the minimum of the Dirichlet energy of a piecewise linear function, on all the possible triangulations of a piecewise flat surface $\mathscr{M}$, is attained at the Delaunay triangulation of $\mathscr{M}$ and the corresponding discrete Laplace-Beltrami operator is intrinsic to the input surface.

On *polygonal meshes*, the Laplacian discretization in [Alexa and Wardetzky 2011; Herholz et al. 2015] generalizes the Laplacian matrix with cotangent weights to surface meshes with non-planar, non-convex faces. Finally, an approximation of the Laplace-Beltrami operator with point-wise convergence has been proposed in [Belkin et al. 2008].

Laplacian matrix for point sets In [Belkin and Niyogi 2003; Belkin and Niyogi 2006; Belkin and Niyogi 2008; Belkin et al. 2009], the Laplace-Beltrami operator on a point set $\mathscr{P}$ has been discretized as the Laplacian matrix

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

To guarantee the sparsity of the Laplacian matrix, for each point $\mathbf{p}_i$ we consider only the entries $L(i,j)$ related to the points $\{\mathbf{p}_j\}_{j \in \mathscr{N}_{\mathbf{p}_i}}$ that are closest to $\mathbf{p}_i$ with respect to the Euclidean distance. In this case, we select either the k-nearest neighbor or the points that belong to a sphere centered at $\mathbf{p}_i$ and with radius σ. As described in [Dey and Sun 2005; Mitra and Nguyen 2003], the choice of σ can be adapted to the local sampling density $\varepsilon := k(\pi\sigma^2)^{-1}$ and the curvature of the surface underlying $\mathscr{P}$. The computation of the k- or σ-nearest neighbor graph takes $\mathscr{O}(n \log n)$-time [Arya et al. 1998; Bentley 1975], where n is the number of input points.

Starting from this approach, a new discretization [Liu et al. 2012] has been achieved through a finer approximation of the local geometry of the surface at each point through its Voronoi cell. More precisely, as $t \to 0$ the stiffness and mass matrix are defined as

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

and v_i is the area of the Voronoi cell associated with 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 $\mathscr{M}$ at $\mathbf{p}_i$. If $\mathbf{B} := \mathbf{I}$, then this approximation reduces to the previous one and both approaches converge to the Laplace-Beltrami operator, as $t \to 0^+$.

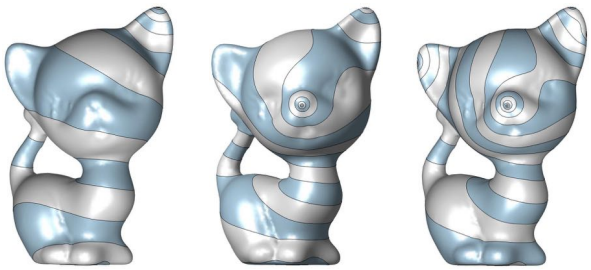

(a) $(1,1,2)$ (b) $(2,2,4)$ (c) $(3,3,6)$

Figure 2: *Level sets and critical points (m,M,s) of harmonic functions with (a) two, (b) four, and (c) six Dirichlet boundary conditions. The insertion of new initial constraints locally affects the resulting harmonic function.*

Laplacian matrix on volumes Representing the input domain as a tetrahedral mesh [Alliez et al. 2005; Liao et al. 2009; Tong et al. 2003], the entries of the stiffness matrix are (Fig. 1c) $L(i,j) := w(i,j) := \frac{1}{6}\sum_{k=1}^{n} l_k \cot \alpha_k$ for each edge (i,j), $L(i,i) := -\sum_{j \in N(i)} w(i,j)$, and zero otherwise; the diagonal mass matrix $\mathbf{B}$ encodes the tetrahedral volume at each vertex.

3 Harmonic equations

The *harmonic function* $h : \mathscr{N} \to \mathbb{R}$ is the solution of the Laplace equation $\Delta h = 0$ with Dirichlet boundary conditions $h|_{\mathscr{S}} = h_0$, $\mathscr{S} \subset \mathscr{N}$. We recall that a harmonic function

- minimizes the *Dirichlet energy* $\mathscr{E}(h) := \int_{\mathscr{N}} \|\nabla h(\mathbf{p})\|_2^2 \mathrm{d}\mathbf{p}$;
- satisfies the *locality property*; i.e., if $\mathbf{p}$ and $\mathbf{q}$ are two distinct points, then $\Delta h(\mathbf{p})$ is not affected by the value of h at $\mathbf{q}$;
- verifies $h(\mathbf{p}) = (2\pi R)^{-1} \int_{\Gamma} h(s) \mathrm{d}s = (\pi R^2)^{-1} \int_{\mathscr{B}} h(\mathbf{q}) \mathrm{d}\mathbf{q}$, where $\mathscr{B} \subseteq \mathscr{N}$ is a disc of center $\mathbf{p}$, radius R, and boundary Γ (*mean-value theorem*).

According to the *maximum principle* [Rosenberg 1997], a harmonic function has no local extrema other than at constrained vertices. In the case that all constrained minima are assigned the same global minimum value and all constrained maxima are assigned the same global maximum value, all the constraints will be extrema in the resulting field. Harmonic and poly-harmonic (i.e., $\Delta^i h = 0$) functions have been applied to volumetric parameterization [Li et al. 2007; Li et al. 2010], to the definition of shape descriptors with pairs of surface points [Zheng et al. 2013] and coupled biharmonic bases [Kovnatsky et al. 2013], to shape approximation [Feng and Warren 2012] and deformation [Joshi et al. 2007; Jacobson et al. 2014; Weber et al. 2012].

Discrete harmonic functions The harmonic equation is approximated at the vertices of $\mathscr{M}$ as the homogeneous linear system

$\mathbf{Lf} = \mathbf{0}$, with initial conditions $f(\mathbf{p}_i) = a_i$, $i \in \mathscr{I} \subseteq \{1,\ldots,n\}$. According to the *Euler formula* $\chi(\mathscr{M}) = m - s + M$, the number of minima m, maxima M, and saddles s of a harmonic function depends on the Dirichlet boundary conditions, which determine the maxima and minima of the resulting harmonic function. In particular, a harmonic function with one maximum and one minimum has a minimal number of $2g$ saddles, where g is the genus of $\mathscr{M}$ (Fig. 2). Harmonic functions are efficiently computed in $\mathscr{O}(n)$ time with iterative solvers of sparse linear systems; their computation is stable for the mean-value weights while negative Voronoi cotangent weights generally induce local undulations in the resulting harmonic function. Main applications include surface quadrangulation [Dong et al. 2005; Ni et al. 2004], the definition of volumetric mappings [Li et al. 2009; Li et al. 2010; Martin et al. 2008; Martin and Cohen 2010], and biharmonic distances [Ovsjanikov et al. 2012; Lipman et al. 2010; Rustamov 2011b] (Sect. 6.2.2).

In the paper examples, the level sets of a given function, or kernel, or distance are associated with iso-values uniformly sampled in its range. For spectral distances, the minimum and the maximum values are depicted in blue and red, respectively. For all the other functions, colors begin with red, pass through yellow, green, cyan, blue, and magenta, and return to red. Finally, the color coding represents the same scale of values for multiple shapes.

4 Laplacian eigenproblem

We introduce the Laplacian eigenpairs (Sect. 4.1), their discretization (Sect. 4.2), and the stability of their computation (Sect. 4.3).

4.1 Laplacian eigenfunctions

Since the Laplace-Beltrami operator is self-adjoint and positive semi-definite, it has an *orthonormal eigensystem* $\mathscr{B} := \{(\lambda_n, \phi_n)\}_{n=0}^{+\infty}$, $\Delta\phi_n = \lambda_n\phi_n$, in $\mathscr{L}^2(\mathscr{N})$. In the following, we assume that the Laplacian eigenvalues are increasingly ordered; in particular $\lambda_0 = 0$. Using the orthonormality and completeness of the Laplacian eigenfunctions in $\mathscr{L}^2(\mathscr{N})$, any function can be represented as a linear combination of the eigenfunctions as $f = \sum_{n=0}^{+\infty}\langle f,\phi_n\rangle_2\phi_n$, where $\langle f,\phi_n\rangle_2\phi_n$ is the projection of f on ϕ_n. Furthermore, the function Δf is expressed in terms of the Laplacian spectrum as $(\Delta f)(\mathbf{p}) = \sum_{n=0}^{+\infty}\lambda_n\langle f,\phi_n\rangle_2\phi_n(\mathbf{p})$ (*spectral decomposition theorem*). A deeper discussion of the analogies between the heat kernel, the Fourier analysis, and wavelets has been presented in [Hammond et al. 2011; Boscaini et al. 2015a].

The Laplacian eigenfunctions are intrinsic to the input shape and those ones related to larger eigenvalues correspond to smooth and slowly-varying functions. Increasing the eigenvalues, the corresponding eigenfunctions generally show rapid oscillations (Fig. 3). From the Laplacian spectrum, we can estimate geometric and topological properties of the input shape. For instance, we can compute the surface area, as the sum of the Laplacian eigenvalues; estimate the Euler characteristic of a surface with genus $g \geq 2$ through the relation [Nadirashvili 1988] $m_j \leq 2j - 2\chi(\mathscr{M}) + 3$, where m_j is the multiplicity of λ_j; and evaluate the total Gaussian curvature [Reuter et al. 2006]. If two shapes are isometric, then they have the same Laplacian spectrum (*iso-spectral property*); however, the viceversa does not hold [Gordon and Szabo 2002; Zeng et al. 2012] and we cannot recover the metric of a given surface.

4.2 Discrete Laplacian eigenpairs

To introduce the discrete Laplcian eigenpairs, the Laplacian eigenproblem is converted to its weak formulation $\langle\Delta\phi,\psi\rangle_2 = \lambda\langle\phi,\psi\rangle_2$ [Allaire 2007], where ψ is a test function. The weak formulation is then discretized as $\mathbf{Lx} = \lambda\mathbf{Bx}$. Here, $\mathbf{L}$, $L(i,j) := \langle\Delta\psi_i,\psi_j\rangle_2$, is the stiffness matrix and $\mathbf{B}$, $B(i,j) := \langle\psi_i,\psi_j\rangle_2$, is the mass matrix. The *generalized Laplacian eigensystem* $\{(\lambda_i,\mathbf{x}_i)\}_{i=1}^n$ $(\lambda_1 = 0)$ satisfies the identity $\mathbf{Lx}_i = \lambda_i\mathbf{Bx}_i$ and the eigenvectors are orthonormal with respect to the $\mathbf{B}$-scalar product; i.e., $\langle\mathbf{x}_i,\mathbf{x}_j\rangle_\mathbf{B} = \mathbf{x}_i^\top\mathbf{Bx}_j = \delta_{ij}$. In particular, the spectral decomposition theorem becomes $\tilde{\mathbf{L}}\mathbf{f} = \sum_{i=1}^n\lambda_i\langle\mathbf{f},\mathbf{x}_i\rangle_\mathbf{B}\mathbf{x}_i = \mathbf{X}\Gamma\mathbf{X}^\top\mathbf{Bf}$, where $\mathbf{X}$ is the eigenvectors' matrix and Γ is the diagonal matrix of the eigenvalues. The discrete Laplacian eigenfunctions generally have a global support (i.e., they are null only at some isolated points) and eigenfunctions with a compact support can be calculated by minimizing the corresponding ℓ_1 norm [Neumann et al. 2014].

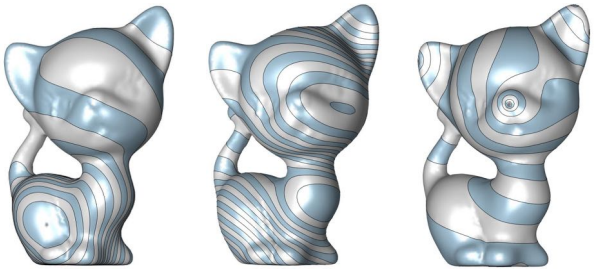

ϕ_1: (2,2,4) ϕ_2: (4,4,8) ϕ_3: (5,3,8)

Figure 3: *Level sets and number of critical points of different Laplacian eigenfunctions (linear FEM weights).*

For the computation of the Laplacian eigenvectors, numerical methods generally exploit the sparsity of the Laplacian matrix and reduce the high-dimensional eigenproblem to one of lower dimension, by applying a coarsening step. The solution is efficiently calculated in the low-dimensional space and then mapped back to the initial dimension through a refinement step. Main examples include the algebraic multi-grid method [Falgout 2006], Arnoldi iterations [Lehoucq and Sorensen 1996; Sorensen 1992], and the Nystrom method [Fowlkes et al. 2004]. Even though the eigenvalues and eigenvectors are computed in super-linear time [Vallet and Lèvy 2008], this computational cost and the required $\mathscr{O}(n^2)$ storage are expensive for densely sampled domains. Indeed, modifications of the Laplacian eigenproblem are applied to locally compute specific sub-parts of the Laplacian spectrum. For instance, the shift method evaluates the spectrum $(\lambda_i - \lambda, \mathbf{x}_i)_{i=1}^n$ of $(\tilde{\mathbf{L}} - \lambda\mathbf{I})$ to calculate the eigenpairs associated with a spectral band centered around a value λ. To compute the larger eigenvalue and the corresponding eigenvector, the inverse method considers the spectrum $(\lambda_i^{-1},\mathbf{x}_i)_{i=2}^n$ of the pseudo-inverse $\tilde{\mathbf{L}}^\dagger$. The power method computes the eigenpairs $(\lambda_i^k,\mathbf{x}_i)_{i=2}^n$ of the sequence of matrices $(\tilde{\mathbf{L}}^k)_{k\geq 1}$ and controls the convergence speed through the selection of k. Finally, pre-conditioners of the Laplacian matrix tailored to computer graphics' applications have been proposed in [Krishnan et al. 2013].

Laplacian eigenfunctions on surfaces In *spectral graph theory*, the Laplacian eigenvectors have been applied to graph partitioning [Fiedler 1973; Mohar and Poljak 1993; Koren 2003] into sub-graphs, which are handled in parallel [Alpert et al. 1999], to graph/mesh layout [Dìaz et al. 2002; Koren 2003], to the reduction of the bandwidth of sparse matrices [Barnard et al. 1993]. In *machine learning*, the Laplacian spectrum have been used for clustering [Schoelkopf and Smola 2002] (§ 14) and dimensionality reduction [Belkin and Niyogi 2003; Xiao et al. 2010] with spectral embeddings. For instance, a common way to measure the dissimilarity between two graphs is to compute the corresponding spectral de-

composition in their own [Lee and Duin 2008] or joint [Umeyama 1988; Caelli and K. 2004] eigenspaces.

In *geometry processing*, the spectral properties of the uniform discrete Laplacian have been used to design low-pass filters [Taubin 1995]. Successively, this formulation has been refined to include the local geometry of the input surface [Desbrun et al. 1999; Kim and Rossignac 2005; Pinkall and Polthier 1993] and it has been applied to implicit mesh fairing [Desbrun et al. 1999; Kim and Rossignac 2005; Zhang and Fiume 2003] and to fairing functionals [Kobbelt et al. 1998; Mallet 1989], which optimize the triangles' shape and/or the surface smoothness [Nealen et al. 2006]. Further applications include mesh watermarking [Ohbuchi et al. 2001; Ohbuchi et al. 2002], geometry compression [Karni and Gotsman 2000; Sorkine et al. 2003], the computation of the gradient [Luo et al. 2009] and the multi-scale approximation of functions [Patanè and Falcidieno 2009; Patanè 2013; Patanè and Spagnuolo 2013a; Patanè and Falcidieno 2009]. The Laplacian eigenvectors have been also used for embedding a surface of arbitrary genus into the plane [Zhou et al. 2004; Zigelman et al. 2002] and mapping a closed genus zero surface into a spherical domain [Gotsman et al. 2003].

In *shape analysis*, the Laplacian spectrum has been applied to shape [Liu and Zhang 2007; Zhang and Liu 2005] segmentation and analysis through nodal domains [Reuter et al. 2009a], correspondence [Jain and Zhang 2007; Jain et al. 2007], and comparison [Marini et al. 2011; Reuter et al. 2006; Jain and Zhang 2007]. Mesh Laplacian operators are also associated with a set of differential coordinates for surface deformation [Sorkine et al. 2004] and quadrangulation with Laplacian eigenfunctions [Dong et al. 2005]. As detailed in Sect. 6, the Laplacian spectrum is also fundamental to define random walks [Ramani and Sinha 2013], commute-time [Bronstein and Bronstein 2011b], biharmonic [Ovsjanikov et al. 2012; Rustamov 2011b], wave kernel [Bronstein and Bronstein 2011b; Aubry et al. 2011], and diffusion distances [Bronstein et al. 2010a; Bronstein et al. 2011; Coifman and Lafon 2006; Gebal et al. 2009; Lafon et al. 2006; Luo et al. 2009; Patanè and Spagnuolo 2013b].

Laplacian eigenfunctions on volumes Laplacian eigenfunctions on a discrete volumetric domain $\mathscr{M}$ are computed either by diagonalizing the corresponding Laplacian matrix or by extending the values of the eigenfunctions computed on the boundary of $\mathscr{M}$ to its interior with barycentric coordinates or non-linear methods (e.g., moving least-squares, radial basis functions) [Patanè et al. 2009; Patané and Spagnuolo 2012]. The computational cost, which is generally high in case of volumetric meshes, is effectively reduced but associated with a lower approximation accuracy. Volumetric Laplacian eigenfunctions have been applied to shape retrieval [Jain and Zhang 2007] and to the definition of volumetric [Rustamov 2011a] shape descriptors.

4.3 Stability of the Laplacian spectrum

Theoretical results on the sensitivity of the Laplacian spectrum against geometry changes, irregular sampling density and connectivity have been presented in [Hildebrandt et al. 2006; Xu 2007]. Here, we briefly recall that the instability of the computation of the Laplacian eigenpairs is generally due to repeated or close eigenvalues, with respect to the numerical accuracy of the solver of the eigen-equation. While repeated eigenvalues are quite rare and typically associated with symmetric shapes, numerically close or switched eigenvalues can be present in the spectrum and in spite of the regularity of the input discrete surface. The following discussion will be useful also for the definition of the conditions on the filter function that induces the spectral distances (Sect. 6.4).

To show that the computation of single eigenvalues is numerically stable, we perturb the Laplacian matrix $\tilde{\mathbf{L}}$ by $\varepsilon\mathbf{E}$, $\varepsilon \to 0$, and compute the eigenpair $(\lambda(\varepsilon),\mathbf{x}(\varepsilon))$ of the new problem $(\mathbf{B}^{-1}\mathbf{L}+\varepsilon\mathbf{E})\mathbf{x}(\varepsilon)=\lambda(\varepsilon)\mathbf{x}(\varepsilon)$, with initial conditions $\mathbf{x}(0)=\mathbf{x}$, $\lambda(0)=\lambda$. The size of the derivative of $\lambda(\varepsilon)$ indicates the variation that it undergoes when the matrix $\tilde{\mathbf{L}}$ is perturbed in the direction $(\mathbf{E},\varepsilon)$. By differentiating the previous equation and evaluating the result at $\varepsilon=0$, we obtain that $\mathbf{BEx}+\mathbf{Lx}'(0)=\lambda'(0)\mathbf{Bx}+\lambda\mathbf{Bx}'(0)$. Multiplying both sides of this last relation with $\mathbf{x}^\top$, the perturbed eigenvalue $|\lambda'(0)|=|\mathbf{x}^\top\mathbf{BEx}|\leq\|\mathbf{Ex}\|_\mathbf{B}$ is bounded by the $\mathbf{B}$-norm of $\mathbf{Ex}$. Indeed, the computation of the Laplacian eigenvalue with multiplicity one is stable.

Assuming that λ_k is an eigenvalue with multiplicity m_k and rewriting the characteristic polynomial as $p_{\tilde{\mathbf{L}}}(\lambda)=(\lambda-\lambda_k)^{m_k}q(\lambda)$, where $q(\cdot)$ is a polynomial of degree $n-m_k$ and $q(\lambda_k)\neq 0$, we get that $(\lambda-\lambda_k)^{m_k}=O(\varepsilon)/q(\lambda)$; i.e., $\lambda=\lambda_k+O(\varepsilon^{\frac{1}{m_k}})$. It follows that a perturbation $\varepsilon:=10^{-m_k}$ produces a change of order 0.1 in λ_k and this amplification becomes more and more evident while increasing the multiplicity of the eigenvalue. According to [Golub and VanLoan 1989] (§ 7), repeated eigenvalues are generally associated with a numerical instability in the computation of the corresponding eigenvectors; in fact, the ℓ_2-norm of the difference between the generalized eigenvectors $\mathbf{x}_i$, $\mathbf{x}_j$ related to the eigenvalues λ_i, λ_j is bounded as

$$\|\mathbf{x}_i-\mathbf{x}_j\|_2\leq\varepsilon\sum_{j\neq i}\left|\frac{\mathbf{x}_i^\top\mathbf{E}\mathbf{x}_j}{\lambda_i-\lambda_j}\right|+O(\varepsilon^2).$$

Indeed, the computation of the eigenvectors related to multiple or close Laplacian eigenvalues might be unstable. Finally, the Laplacian eigenvalues might be locally switched (i.e., we are not able to numerically distinguish two consecutive eigenvalues) and this situation happens independently of the quality of the discretized surface in terms of point density, angles, and connectivity.

5 Heat and wave equations

We introduce the heat (Sect. 5.1), wave and mean curvature flow (Sect. 5.2) equations; then, we discuss their discretization (Sect. 5.3), the selection of the time scale (Sect. 5.4), and the computation of their solution (Sect. 5.5).

5.1 Heat equation

The *scale-based representation* $H:\mathscr{N}\times\mathbb{R}^+\to\mathbb{R}$ of the function $h:\mathscr{N}\to\mathbb{R}$ is the solution to the *heat diffusion equation* $(\partial_t+\Delta)H(\mathbf{p},t)=0$, $H(\cdot,0)=h$. The function $H(\mathbf{p},t)$ represents the heat distribution at the point $\mathbf{p}$ and at time t, where h is the initial distribution. The solution to the heat equation is written as

$$H(\mathbf{p},t)=\langle K_t(\mathbf{p},\cdot),h\rangle_2=\sum_{n=0}^{+\infty}\exp(-\lambda_n t)\langle h,\phi_n\rangle_2\phi_n(\mathbf{p}),\quad(1)$$

where $K_t(\mathbf{p},\mathbf{q})=\sum_{n=0}^{+\infty}\exp(-\lambda_n t)\phi_n(\mathbf{p})\phi_n(\mathbf{q})$ is the spectral representation of the *heat diffusion kernel*. The heat diffusion and the Laplace-Beltrami operators have the same eigenfunctions $\{\phi_n\}_{n=0}^{+\infty}$ and $(\exp(-\lambda_n t))_{n=0}^{+\infty}$ are the eigenvalues of the heat operator. The heat kernel is invariant to isometries and verifies the *semi-group* $\langle K_{t_1},K_{t_2}\rangle_2=K_{t_1+t_2}$ and *inversion* $K_t^{-1}=K_{-t}$ *properties*. The spectral representation (1) shows the smoothing effect on the initial condition h; as the scale increases, the component of h along the eigenfunctions associated with the larger Laplacian eigenvalue becomes null. We also notice that the normalized function $\mathscr{A}_\mathscr{N}^{-1}H(\cdot,t)$

with respect to the surface area $\mathscr{A}_\mathscr{N}$ minimizes the weighted least-squares error $\int_\mathscr{N} K_t(\mathbf{p},\mathbf{q})|h(\mathbf{q})-g(\mathbf{p})|^2 \mathrm{d}\mathbf{q}$ on $\mathscr{L}^2(\mathscr{N})$, for a given h.

Heat equation on surfaces On surfaces, the heat kernel satisfies the following properties [Sun et al. 2009; Grigoryan 2006]:

- for an isometry $\Phi : \mathscr{N} \to \mathscr{Q}$ between two manifolds $\mathscr{N}$, $\mathscr{Q}$,

$$K_t^\mathscr{N}(\mathbf{p},\mathbf{q}) = \mathscr{K}_t^\mathscr{Q}(\Phi(\mathbf{p}),\Phi(\mathbf{q})), \qquad \forall \mathbf{p},\mathbf{q} \in \mathscr{N}, \forall t \in \mathbb{R}^+; \tag{2}$$

- if Φ is surjective and Eq. (2) holds, then Φ is an isometry;
- if D is a compact set of $\mathscr{N}$, then $\lim_{t\to 0} K_t^D(\mathbf{p},\mathbf{q}) = K_t^\mathscr{N}(\mathbf{p},\mathbf{q})$;
- if $D_1 \subseteq D_2 \subseteq \mathscr{N}$, then $K_t^{D_1}(\mathbf{p},\mathbf{q}) \leq K_t^{D_2}(\mathbf{p},\mathbf{q})$;
- on smooth and polygonal surfaces, the heat kernel fully determines the Riemannian metric [Zeng et al. 2012].

For small values of t [Sun et al. 2009; Varadhan 1967], the auto-diffusivity function

$$K_t(\mathbf{p},\mathbf{p}) \approx \begin{cases} (4\pi t)^{-1}(1+1/3t\kappa(\mathbf{p})) + \mathscr{O}(t^2), \\ (4\pi t)^{3/2}(1+1/6s(\mathbf{p})), \end{cases} \quad t \to 0,$$

encodes the Gaussian $\kappa(\mathbf{p})$ and total $s(\mathbf{p})$ curvature at $\mathbf{p}$. For large t, $K_t(\mathbf{p},\mathbf{q})$ is dominated by the Fiedler vector ϕ_1 [Fiedler 1973], which encodes the global structure of the input shape. According to [Sun et al. 2009; de Goes et al. 2008], the surface $\mathscr{N}$ at $\mathbf{p}$ can be characterized in terms of the average squared diffusion distance at $\mathbf{p}$ (*eccentricity*), which is defined as

$$\mathrm{ecc}_t(\mathbf{p}) = \mathscr{A}_\mathscr{N}^{-1} \int_\mathscr{N} d_t(\mathbf{p},\mathbf{q})\mathrm{d}\mathbf{q} = K_t(\mathbf{p},\mathbf{p}) + E_\mathscr{N}(t) - 2\mathscr{A}_\mathscr{N}^{-1}, \quad t \to 0,$$

where $E_\mathscr{N}(t) := \sum_{n=0}^{+\infty} \exp(-\lambda_n t)$ is the sum of the eigenvalues of the heat kernel. Since the area and trace are independent of the evaluation point, the functions ecc_t and $K_t(\mathbf{p},\cdot)$ have the same level sets and extrema on $\mathscr{N}$. In particular, for small scales the extrema of the eccentricity are localized at the curvature extrema.

Heat equation on volumes The analytical representation of the volumetric heat kernel $K_t(\mathbf{p},\mathbf{q}) := (4\pi t)^{-3/2}\exp(-\|\mathbf{p}-\mathbf{q}\|_2^2/4t)$ allows us to solve the heat equation as $F(\cdot,t) = \langle K_t, f\rangle_2$ and without computing the Laplacian spectrum (Sect. 5.5.4).

5.2 Wave equation and mean curvature flow

The heat equation is strictly related to the *Schroedinger (wave) equation* $(i\Delta + \partial_t)H(\cdot,t) = 0$, with initial condition $H(\cdot,0) = h$, which represents the physical model of a quantum particle with initial energy h. The spectral representation of the solution is $H(\cdot,t) = \sum_{n=0}^{+\infty} \exp(i\lambda_n t)\langle h,\phi_n\rangle_2 \phi_n$; i.e., a complex wave function with oscillatory behavior. This periodic effect is due to the real and complex parts of the filter $\exp(i\lambda_n t) = \cos(\lambda_n t) + i\sin(\lambda_n t)$. The norm of the solution is the probability $P_t(\mathbf{p})$ to find a point $\mathbf{p}$ after a time t; in fact, the following identity holds

$$\begin{aligned} P_t(\mathbf{p}) &= \lim_{T\to+\infty} \int_0^T |H(\mathbf{p},t)|^2 \mathrm{d}t \\ &= \sum_{n=0}^{+\infty} |\langle h,\phi_n\rangle_2|^2 |\phi_n(\mathbf{p})|^2 = \|H(\mathbf{p},t)\|_2^2. \end{aligned}$$

Figure 4: *Anisotropic heat kernel centered at a (black) seed point on a coarse triangle mesh.*

Table 1: *Main properties of the discrete heat kernel: sparsity, positive definiteness, and symmetry. The full • and empty ∘ circle means that the corresponding property is or is not satisfied, respectively.*

Heat Ker.	Matrix $\mathbf{K}_t$	Sp.	Pos. Def.	Sym.	Cov.	Inv.
Std	$\mathbf{X}\mathbf{D}_t\mathbf{X}^\top$	∘	•	•	∘	∘
Vor.-cot	$\mathbf{X}\mathbf{D}_t\mathbf{X}^\top\mathbf{D}$	∘	•	∘	•	∘
wFEM	$\mathbf{X}\mathbf{D}_t\mathbf{X}^\top\mathbf{B}$	∘	•	∘	•	∘

Finally, the heat equation is related to the *mean curvature flow* [Crane et al. 2013a; Kazhdan et al. 2012] $(\partial_t + \Delta_t)\Phi_t = 0$, where $\Phi_t : \mathscr{M} \to \mathbb{R}^3$ is a family of immersions and Δ_t is the Laplace-Beltrami operator associated with the metric induced by the immersion at time t.

5.3 Discrete heat equation and kernel

We briefly introduce the weak formulation [Allaire 2007] of the heat equation; similar results apply to the equations previously introduced. Chosen a set $\mathscr{B} := \{\psi_i\}_{i=1}^n$ of linearly independent functions on $\mathscr{N}$, we approximate the solution $\tilde{F}(\mathbf{p},t) := \sum_{i=1}^n a_i(t)\psi_i(\mathbf{p})$ to the weak heat equation as $\langle \partial_t \tilde{F}(\cdot,t), \psi_i\rangle_2 + \langle \Delta\tilde{F}(\cdot,t), \psi_i\rangle_2 = 0$, $i = 1,\ldots,n$. Introducing the matrices $\mathbf{L} := (\langle \Delta\psi_i, \psi_j\rangle_2)_{i,j=1}^n$ and $\mathbf{B} := (\langle \psi_i, \psi_j\rangle_2)_{i,j=1}^n$, the discrete heat equation becomes $(\mathbf{B}\partial_t + \mathbf{L})\mathbf{a}(t) = \mathbf{0}$, $\mathbf{a}(t) := (a_i(t))_{i=1}^n$. An analogous relation can be derived for the boundary condition $F(\mathbf{p},0) = h(\mathbf{p})$. Since $\mathbf{B}$ is the Gram matrix associated with $\mathscr{B}$, it is invertible and the previous system of equations is $(\partial_t + \mathbf{B}^{-1}\mathbf{L})\mathbf{a}(t) = \mathbf{0}$, with initial condition $\mathbf{F}(0) = \mathbf{f}$. Then, the solution to the discrete heat equation is expressed as a linear combination of the Laplacian eigensystem as $\mathbf{F}(t) = \sum_{i=1}^n \exp(-\lambda_i t)\langle \mathbf{f}, \mathbf{x}_i\rangle_\mathbf{B}\mathbf{x}_i$.

Properties The solution to the discrete heat equation is $\mathbf{F}(t) = \mathbf{K}_t\mathbf{f}$ (Fig. 4), where $\mathbf{K}_t := \mathbf{X}\mathbf{D}_t\mathbf{X}^\top\mathbf{B}$, $\mathbf{D}_t := \mathrm{diag}\,(\exp(-\lambda_i t))_{i=1}^n$, is the *heat kernel matrix* (Table 1). Lumping the linear FEM mass matrix $\mathbf{B}$, the heat kernel becomes equal to the *Voronoi-cotangent heat kernel* $\mathbf{K}_t^\star := \mathbf{X}\mathbf{D}_t\mathbf{X}^\top\mathbf{D}$, $\mathbf{L}\mathbf{X} = \mathbf{X}\Gamma$. Choosing $\mathbf{B} := \mathbf{I}$, we get the *heat kernel* $\tilde{\mathbf{K}}_t := \mathbf{X}\mathbf{D}_t\mathbf{X}^\top$ with cotangent weights. Analogously to the results in Sect. 5.1, the heat kernel matrix satisfies the following relations: $\mathbf{K}_{t_1} \times \mathbf{K}_{t_2} = \mathbf{K}_{t_2} \times \mathbf{K}_{t_1} = \mathbf{K}_{t_1+t_2}$ (*commutative* and *semi-group properties*), $\mathbf{K}_t^{-1} = \mathbf{K}_{-t}$ (*inversion property*). If $\mathbf{B}$ is the linear FEM mass matrix or the diagonal matrix of the Voronoi areas, then the heat kernel matrix $\mathbf{K}_t$ is intrinsically *scale-covariant*; i.e., rescaling the points of $\mathscr{M}$ by a factor α, $\alpha > 0$, and indicating the new surface as $\alpha\mathscr{M}$, we get that only the time component of the kernel is rescaled. In fact, the rescaling changes the matrix $\mathbf{B}$ and the eigensystem $\{(\lambda_i, \mathbf{x}_i)\}_{i=1}^n$

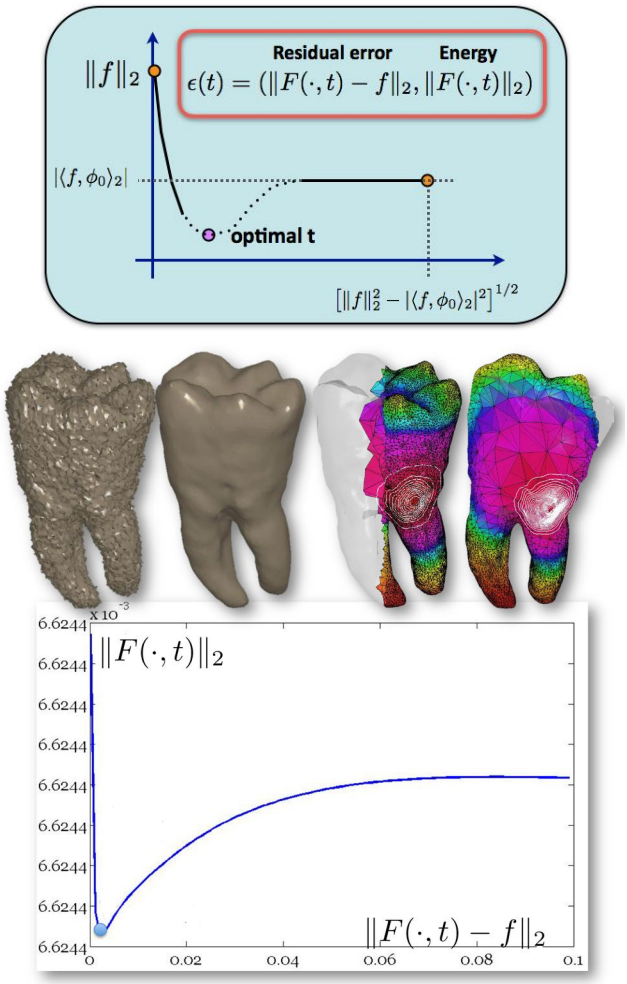

Figure 5: *First row: behavior of the L-curve. Second row: selection of the optimal scale ($t_{opt} = 0.0032$) and corresponding volumetric diffusion smoothing, Padé-Chebyshev approximation of degree $r = 7$) on the noisy volumetric model of a teeth.*

of $\mathscr{M}$ into $\alpha^2\mathbf{B}$ and $\{(\alpha^{-2}\lambda_i, \alpha^{-1}\mathbf{x}_i)\}_{i=1}^n$, respectively. Indeed, $\mathbf{K}_t(\alpha\mathscr{M}) = \mathbf{K}_{\alpha^{-2}t}(\mathscr{M})$ without an *a-posteriori* normalization. The scale-covariance of $\mathbf{K}_t$ is guaranteed by the mass matrix, which changes according to the surface rescaling and compensates the variation of the corresponding Laplacian spectrum. The kernel becomes *scale-invariant* (i.e., $\mathbf{K}_t(\alpha\mathscr{M}) = \mathbf{K}_t(\mathscr{M})$) by normalizing each eigenvalue by λ_n, which is efficiently computed using the inverse method [Golub and VanLoan 1989; Vallet and Lèvy 2008]. Alternatively, the scale-invariance and covariance of the heat kernel is achieved in the Fourier domain [Bronstein and Kokkinos 2010]. In [Bronstein et al. 2010b; Bronstein et al. 2010c], the matching performances of heat kernel descriptors have been tested against shape transformation, sampling, and noise.

5.4 Selection of the time scale

For shape analysis, the real line is uniformly sampled in order to consider both small and large scales. For geometry processing, the *optimal time value* is defined as the value of t that provides the best compromise between a small residual error $\|F(\cdot,t) - f\|_2^2 = \sum_{n=0}^{+\infty} |1 - \exp(-2\lambda_n t)|^2 |\langle f, \phi_n \rangle_2|^2$ and a low energy $\|F(\cdot,t)\|_2^2 = \sum_{n=0}^{+\infty} \exp(-2\lambda_n t) |\langle f, \phi_n \rangle_2|^2$. If t tends to zero, then the residual becomes null and the energy converges to $\|f\|_2$. If t becomes large, then the residual increases until it converges to $|\langle f, \phi_0 \rangle_2|$ and the solution norm decreases until it converges to $(\|f\|_2^2 - |\langle f, \phi_0 \rangle_2|^2)^{1/2}$. According to these properties, the plot (*L-curve*) of the energy (y-axis) versus the residual (x-axis) is L-shaped [Hansen and O'Leary 1993] and its minimum provides the optimal time value (Fig. 5). For the computation of the optimal time value, we mention the corner detection based on cubic B-splines approximation [Hansen and O'Leary 1993], the evaluation of the curvature of the graph of the L-curve or its adaptive pruning [Hansen and O'Leary 1993].

5.5 Computation of the discrete heat kernel

For the computation of the solution to the discrete heat equation and kernel, we consider linear (Sect. 5.5.1), polynomial (Sect. 5.5.2), and rational (Sect. 5.5.3) approximations of the exponential filter. On volumes (Sect. 5.5.4), we discuss the solution to the heat equation based on the analytic representation of the heat kernel. With the exception of the truncated spectral method, all the previous approximations are independent of the evaluation of the Laplacian spectrum and reduce to a set of sparse linear systems (Table 2). The polynomial and rational approximations generally provide the best compromise between approximation accuracy and computational cost.

5.5.1 Linear approximation

For the solution to the heat equation, we review the truncated spectral approximation, the Euler backward method, the first order Taylor approximation, the Krylov and Schur methods.

Truncated spectral approximation and power method The computational bottleneck for the evaluation of the whole Laplacian spectrum imposes on us to consider only a small subset of the Laplacian spectrum. Since the decay of the *filter factor* $\exp(-\lambda_i t)$ increases with λ_i, in the spectral representation of the solution to the heat equation we consider only the contribution related to the first k eigenpairs; i.e., $\mathbf{F}_k(t) = \sum_{i=1}^{k} \exp(-\lambda_i t) \langle \mathbf{f}, \mathbf{x}_i \rangle_{\mathbf{B}} \mathbf{x}_i$. The truncated approximation is accurate only if the exponential filter decays fast (e.g., large values of time) and the effect of the selected eigenpairs on the approximation accuracy cannot be estimated without computing the whole spectrum. The *multi-resolution prolongation operators* [Vaxman et al. 2010] prolongate the values of the truncated spectral approximation, computed on a low-resolution representation of the input shape, to the initial resolution through a hierarchy of simplified meshes. In this case, the number of eigenpairs are heuristically adapted to the surface resolution and its global/local features.

Euler backward method In [Clarenz et al. 2000; Desbrun et al. 1999], 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 function, as $k \to +\infty$.

First order Taylor approximation and power method Since the derivative of $\mathbf{K}_t$ at $t = 0$ equals the Laplacian matrix (i.e., $(\mathbf{I} - \mathbf{K}_t)/t \to \mathbf{B}^{-1}\mathbf{L}$, $t \to 0$), the heat kernel $\mathbf{K}_t$ is approximated by $(\mathbf{I} - t\mathbf{B}^{-1}\mathbf{L})$ and $\mathbf{F}(t) = \mathbf{K}_t\mathbf{f}$ solves the sparse linear system $\mathbf{B}(\mathbf{K}_t\mathbf{f}) = (\mathbf{B} - t\mathbf{L})\mathbf{f}$. This last relation gives an approximation of $\mathbf{F}(t)$ that is independent of the Laplacian spectrum and is valid only for small values of t. For an arbitrary value of t, the "power"

Table 2: *Numerical computation of the solution to the heat equation; $\tau(n)$ is the cost for the solution of a sparse linear system.*

Method	Numerical scheme	Scales	Comput. cost	References
Linear approximation				
Trunc. spec. approx.	$\mathbf{F}_k(t) = \sum_{i=1}^{k} \exp(-\lambda_i t)\langle \mathbf{f}, \mathbf{x}_i \rangle_{\mathbf{B}} \mathbf{x}_i$	Any	$\mathcal{O}(kn)$	[Golub and VanLoan 1989; Vaxman et al. 2010]
Euler backw. approx.	$(t\tilde{\mathbf{L}}+\mathbf{I})\mathbf{F}_{k+1}(t) = \mathbf{F}_k(t)$	Small	$\mathcal{O}(\tau(n))$	[Clarenz et al. 2000; Desbrun et al. 1999; Zhang and Hancock 2008]
I order Taylor approx.	$\mathbf{BF}(t) = (\mathbf{B} - t\mathbf{L})\mathbf{f}$	Small	$\mathcal{O}(\tau(n))$	[Clarenz et al. 2000; Desbrun et al. 1999]
Krylov/Schur approx.	Projection on $\{\mathbf{g}_i := (\mathbf{B}^{-1}\mathbf{L})^i \mathbf{f}\}_{i=1}^{m}$	Any	$\mathcal{O}(m\tau(n))$, $\mathbf{B} \neq \mathbf{I}$ $\mathcal{O}(n)$, $\mathbf{B} = \mathbf{I}$	[Golub and VanLoan 1989; Saad 1992; Zhang and Hancock 2008]
Polynomial approximation				
Power approx.	$\mathbf{F}(t) = \sum_{i=0}^{m} \mathbf{g}_i / i!$ $\mathbf{g}_i := \tilde{\mathbf{L}}^i \mathbf{f}$	Any	$\mathcal{O}(m\tau(n))$, $\mathbf{B} \neq \mathbf{I}$ $\mathcal{O}(n)$, $\mathbf{B} = \mathbf{I}$	[Golub and VanLoan 1989]
Rational approximation				
Padé-Cheb. approx.	$\mathbf{F}(t) = \alpha_0 \mathbf{f} + \sum_{i=1}^{r} \mathbf{g}_i$ $(t\mathbf{L} + \theta_i \mathbf{B})\mathbf{g}_i = -\alpha_i \mathbf{Bf}$	Any	$\mathcal{O}(r\tau(n))$	[Carpenter et al. 1984; Sidje 1998; Saad 1992] [Patanè 2013; Patanè 2014; Patané 2016]
Contour integral approx.	$\mathbf{F}(t) = \sum_{i=1}^{r} \alpha_i \mathbf{g}_i$ $(\alpha_i)_{i=1}^{r}$ quadr. coeff.	Any	$\mathcal{O}(r\tau(n))$	[Pusa 2011]

Algorithm 1: *Padé-Chebyshev approximation of the solution to the heat equation.*

Require: A function $f : \mathscr{P} \to \mathbb{R}$, $\mathbf{f} := (f(\mathbf{p}_i))_{i=1}^{n}$.
Ensure: The approximate solution $\mathbf{F}(t) = \mathbf{K}_t \mathbf{f}$ of $\mathbf{f}$ to the heat equation.
1: Select the value of t (e.g., optimal value, Sect. 5.4).
2: **for** $i = 1, \ldots, r-1$ **do**
3: Compute $\mathbf{g}_i$: $(t\mathbf{L} + \theta_i \mathbf{B})\mathbf{g}_i = -\alpha_i \mathbf{Bf}$.
4: **end for**
5: Approximate $\mathbf{K}_t \mathbf{f}$ as $\alpha_0 \mathbf{f} + \sum_{i=1}^{r} \mathbf{g}_i$. =0

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 $\mathbf{K}_{t/m} \approx (\mathbf{I} - t/m\tilde{\mathbf{L}})$ is accurate. However, the selection of m and its effect on the approximation accuracy cannot be estimated a-priori.

Krylov and Schur approximations The Krylov subspace projection [Golub and VanLoan 1989; Saad 1992] computes an approximation of $\exp(-t\mathbf{A})\mathbf{f}$ in the space generated by the vectors $\mathbf{f}, \mathbf{Af}, \ldots, \mathbf{A}^{m-1}\mathbf{f}$, thus processing a $m \times m$ matrix instead of a $n \times n$ matrix, where m is much lower than n (e.g., $m \approx 20$). This approximation [Zhang and Hancock 2008] becomes computationally expensive when the dimension of the Krylov space increases, still remaining much lower than n (e.g., $n \approx 5K$). In both cases, the vector $\tilde{\mathbf{L}}^i \mathbf{f} = (\mathbf{B}^{-1}\mathbf{L})^i \mathbf{f}$ must be computed without inverting the mass matrix; to this end, we notice that the vector $\mathbf{g}_i := (\mathbf{B}^{-1}\mathbf{L})^i \mathbf{f}$ satisfies the linear system $\mathbf{Bg}_i = \mathbf{Lg}_{i-1}$, $\mathbf{Bg}_1 = \mathbf{Lf}$. Since the coefficient matrix $\mathbf{B}$ is sparse, symmetric, and positive definite, the vectors $(\mathbf{g}_i)_{i=1}^{m}$ are evaluated in linear time by applying iterative solvers (e.g., conjugate gradient) or pre-factorizing $\mathbf{B}$.

5.5.2 Polynomial approximations

The exponential of a matrix $\mathbf{A}$ is defined as the exponential *power series* $\exp(\mathbf{A}) = \sum_{n=0}^{+\infty} \mathbf{A}^n / n!$, which converges for any square matrix $\mathbf{A}$. Even though the input matrix $\mathbf{A}$ is sparse, its exponential $\exp(-t\mathbf{A})$ is always full ($t \neq 0$) and can be computed or stored only if $\mathbf{A}$ has a few hundred rows and columns only. In particular, for computer graphics applications we can consider 3D shapes only with a small number of samples (i.e., few hundreds) or evaluate the heat kernel on a set of seed points that are representative of the geometry and features of the input shape.

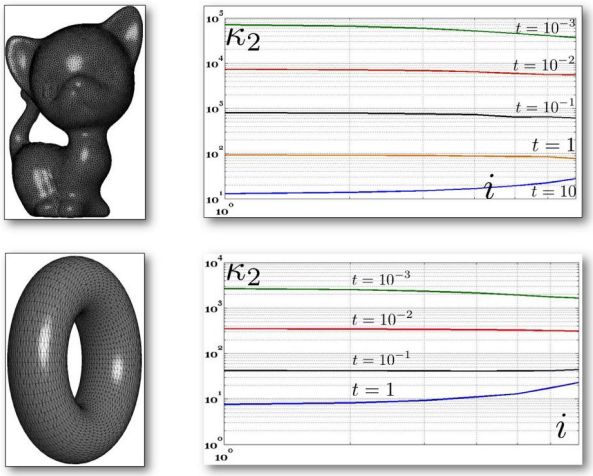

Figure 6: *Conditioning number κ_2 (y-axis) of the matrices $\{(t\mathbf{L} + \theta_i \mathbf{B})\}_{i=1}^{7}$, for several values the time parameter t; the indices of the coefficients $\{\theta_i\}_{i=1}^{7}$ are reported on the x-axis.*

5.5.3 Rational approximation

The exponential of an arbitrary matrix $\mathbf{A}$ is equal to the complex contour integral $\exp(t\mathbf{A}) = (2\pi i)^{-1} \int_{\Gamma} \exp(z)(z\mathbf{I} - t\mathbf{A})^{-1} \mathrm{d}z$, where Γ is a closed contour winding once around the spectrum of $t\mathbf{A}$ [Golub and VanLoan 1989] (§ 11), [Rudin 1987] (§ 10). From this identity, we introduce two accurate and computationally efficient approximations of the exponential of the Laplacian matrix.

Padé-Chebyshev approximation The rational approximation of the exponential function of order (k,k) and with simple poles is $r_{kk}(z) := p_k(z)/q_k(z) = \alpha_0 + \sum_{i=1}^{k} \alpha_i (z - \theta_i)^{-1}$, where p_k, q_k are polynomials of order k, $\alpha_0 = \lim_{z \to +\infty} r_{kk}(z)$, θ_i is a pole, and α_i is the residual at θ_i. Applying this last relation to $t\mathbf{A}$, we get $\exp(t\mathbf{A}) = \alpha_0 \mathbf{I} + \sum_{i=1}^{k} \alpha_i (t\mathbf{A} - \theta_i \mathbf{I})^{-1}$. Among the rational approximations of the exponential function, we focus on its best approximation $r_{kk}(\cdot)$ of order k with respect to the ℓ_∞ norm; i.e., the unique $r_{kk}(z) = p_k(z)/q_k(z)$ that minimizes the error $\|\pi(z) - \exp(-z)\|_\infty$ in the space $\Pi_{kk} \ni \pi$ of rational polynomials of order k. Here, the main difficulty is the evaluation of the coefficients and poles of the rational approximation of the exponential function for a given k, which is generally affected by the ill-conditioned computation of

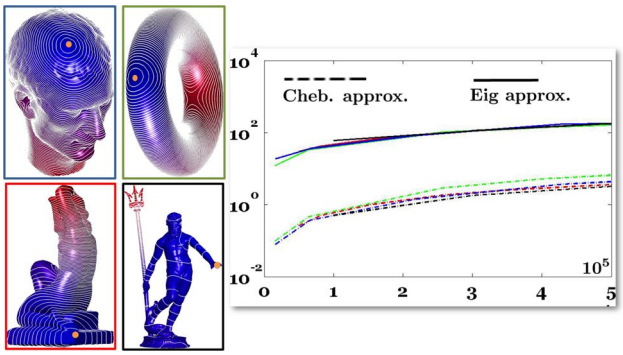

Figure 7: *Cost (in seconds, y-axis,* log*-scale) for the evaluation of the diffusion distances on 3D shapes with n samples (x-axis), approximated with $k = 500$ eigenpairs and the Padé-Chebyshev approximation. Colors from the source (orange) point vary from blue (null distance) to red (maximum distance).*

the polynomial roots. These coefficients and poles have been computed with a different accuracy and for different orders of the rational polynomial [Carpenter et al. 1984; Cody et al. 1969; Moler and Van Loan 2003; Sidje 1998; Saad 1992]. These approximations are also included in standard numerical libraries for signal processing. Finally, we recall that in spectral graph theory [Orecchia et al. 2012], the Padé-Chebyshev and the Lanczos methods have been applied to the approximation of $\exp(-\mathbf{A})\mathbf{f}$, where $\mathbf{A}$ is a symmetric and positive semi-definite matrix.

The idea behind the *spectrum-free computation* [Patanè 2013; Patanè 2014] is to apply the (r,r)-degree Padé-Chebyshev rational approximation to the exponential representation $\mathbf{F}(t) = \exp(-t\tilde{\mathbf{L}})\mathbf{f}$ of the solution to the heat equation $(\partial_t + \tilde{\mathbf{L}})\mathbf{F}(t) = \mathbf{0}$, $\mathbf{F}(0) = \mathbf{f}$ (Algorithm 1). In this case, the solution $\mathbf{F}(t) = \alpha_0\mathbf{f} + \sum_{i=1}^{r}\mathbf{g}_i$ is the sum of the solutions of r sparse linear systems $(t\mathbf{L} + \theta_i\mathbf{B})\mathbf{g}_i = -\alpha_i\mathbf{B}\mathbf{f}$, $i = 1,\ldots,r$. The resulting approximation belongs to the linear space generated by $\mathbf{f}$ and $\{\mathbf{g}_i\}_{i=1}^{r}$, which are calculated as a minimum norm residual solution [Golub and VanLoan 1989], depend on the input domain, the initial condition, and the selected time value. 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 $\mathbf{F}(t)$ with respect to t and $\mathbf{f}$.

This approximation is independent of the computation of the Laplacian spectrum, user-defined parameters, and multi-resolutive prolongation operators [Vaxman et al. 2010], which heuristically adapt the number of eigenpairs to the surface resolution. The sparse and well-conditioned matrices of the previous linear systems have the same structure and sparsity of the connectivity matrix of the input domain, properly encode the local and global features in the heat kernel, and can be computed for any representation of the input domain and of the Laplacian weights. Finally, the accuracy of the Padé-Chebychev approximation is lower than 10^{-r} (e.g., $r = 5,7$).

The value of t influences the conditioning number of the matrices $(t\mathbf{L} + \theta_i\mathbf{B})$, $i = 1,\ldots,r$. Experiments (Fig. 6,[Patanè 2014]) have shown that the linear systems associated with the Padé-Chebyshev approximation are generally well-conditioned; in any case, preconditioners and regularization techniques [Golub and VanLoan 1989] can be applied to attenuate numerical instabilities. Finally, timings on surfaces and volumes (Fig. 7) are reduced from 20 up to 1200 times with respect to the approximation based on a fixed number of Laplacian eigenpairs. Laplacian eigenvectors have been computed with the Arnoldi iteration method [Lehoucq and Sorensen 1996; Sorensen 1992].

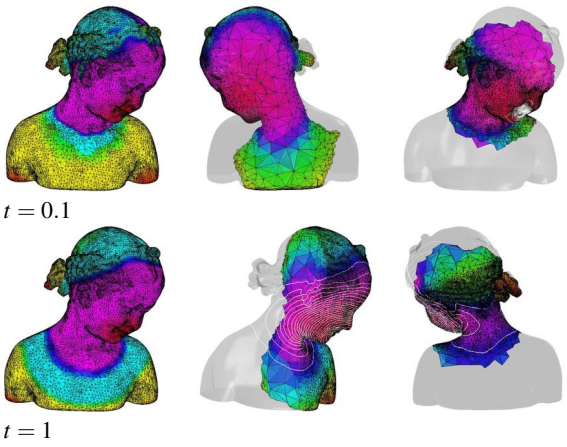

Figure 8: *Volumetric heat kernel ($r = 7$). Level-sets correspond to iso-values uniformly sampled in the range of the solution restricted to the volume boundary.*

Rational approximation from contour integrals Since the exponential factor rapidly decays to zero as $Re(z) \to +\infty$, in [Pusa 2011] the complex contour integral has been efficiently computed with quadrature rules. In this case, $\alpha_0 = 0$, the poles $\theta_i := \phi(x_i)$ are evaluated at the quadrature points $\{x_i\}_i$, $\alpha_i := -(2\pi i)^{-1}h\exp(\phi(x_i))\phi'(x_i)$ are the weights of the quadrature rules, and h is the interval length in the quadrature scheme. The resulting approximation accuracy is guided by the degree of the quadrature rule; low degrees (e.g., $k = 2$, $k = 4$) generally provide a satisfactory approximation accuracy.

5.5.4 Special case: heat equation on volumes

On a volume, the function $\mathbf{F}(t) = \sum_{i=1}^{n}\alpha_i V_i K_t(\mathbf{p}_i,\cdot)f(\mathbf{p}_i)$ is approximated as a linear combination of the basis functions $\{K_t(\mathbf{p}_i,\cdot)\}_{i=1}^{n}$. Here, $\mathbf{V} = \mathrm{diag}(V_i)_{i=1}^{n}$ is the diagonal matrix of the volumes V_i at $\mathbf{p}_i$, $\mathbf{K}_t$ is the Gram matrix for the Gaussian kernel, and the unknowns $\alpha = (\alpha_i)_{i=1}^{n}$ are determined by imposing the condition $F(\mathbf{p}_i,0) = f(\mathbf{p}_i)$, $i = 1,\ldots,n$. To overcome the time-consuming solution of the $n \times n$ linear system $\mathbf{V}\mathbf{K}_t\alpha = \mathbf{f}$, the number of conditions is reduced or the coefficient matrix is sparsified according to the exponential decay of its entries. Alternatively, the volumetric heat equation is solved by discretizing the Laplace-Beltrami operator with finite elements [Allaire 2007; Reuter et al. 2009b], or with finite differences on a 6-neighborhood stencil [Litman et al. 2011; Litman et al. 2012; Raviv et al. 2010], or with a geometry-driven approximation of the gradient field [Liao et al. 2009; Tong et al. 2003].

While a discretization of the heat kernel on a voxel grid is accurate enough for the evaluation of diffusion descriptors [Litman et al. 2011; Raviv et al. 2010], which are quantized and clustered in bags-of-features, the computation of the solution to the volumetric heat equation generally requires a more accurate discretization of the input domain. The prolongation of the Laplacian [Rustamov 2011a; Rustamov 2011b], harmonic [Li et al. 2010; Martin et al. 2008], and diffusion functions from the volume boundary to its interior, through barycentric coordinates or non-linear approximation, achieves a low accuracy of the solution in a neighbor of the boundary. The multi-resolution simplification of the input volume is also time-consuming, and the selection of the volume resolution with respect to the expected approximation accuracy are generally guided by heuristics. Indeed, these methods do not intend to approximate the heat kernel quantitatively, but provide alternative approaches

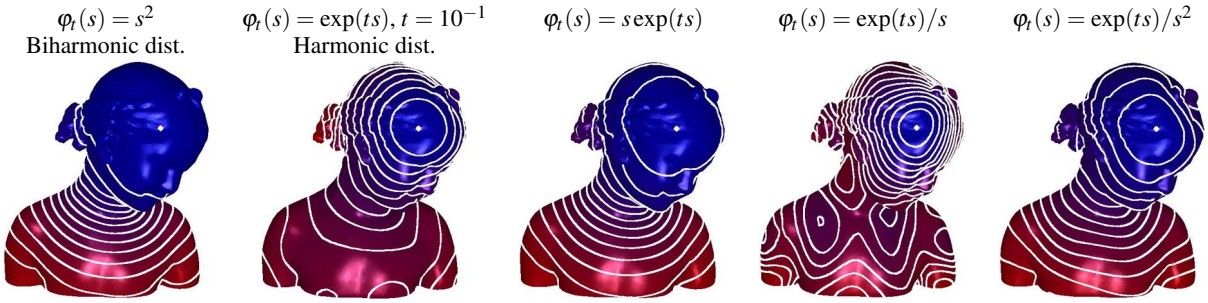

Figure 9: *Level-sets of the spectral distances from a source point (white dot) induced by the filter φ and evaluated with the Padé-Chebyshev approximation ($r = 5$).*

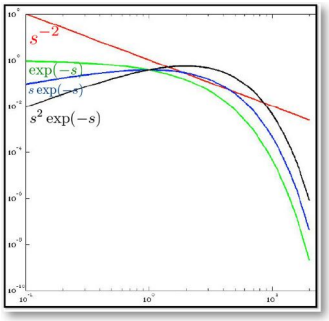

φ^2	**Spect. dist.**	$d^2(\mathbf{p},\mathbf{q}) = \sum_{n=0}^{+\infty} \varphi^2(\lambda_n)\lvert\phi_n(\mathbf{p}) - \phi_n(\mathbf{q})\rvert^2$	
		Associated equation	**Kernel**
s^{-1} s^{-2} s^{-k}	Comm.-time Biharm. Poly-harm.	$(\partial_t + \Delta^r)\, F(\cdot,t) = 0$ *Poly-harm. eq.*	$\sum_{n=0}^{+\infty} t^k e^{-\lambda_n t}\lvert\phi_n(\mathbf{p}) - \phi_n(\mathbf{q})\rvert^2$
e^{-st}	Heat diff.	$(\partial_t + \Delta)\, F(\cdot,t) = 0$ *Heat diffusion eq.*	$\sum_{n=0}^{+\infty} e^{-\lambda_n t}\lvert\phi_n(\mathbf{p}) - \phi_n(\mathbf{q})\rvert^2$
e^{-ist}	Wave ker.	$(\partial_t + i\Delta)\, F(\cdot,t) = 0$ *Schroedinger eq.*	$\sum_{n=0}^{+\infty} e^{-i\lambda_n t}\lvert\phi_n(\mathbf{p}) - \phi_n(\mathbf{q})\rvert^2$

Figure 10: *Spectral distances and kernels induced by the filter function φ (*log*-scale on the t- and y-axis) applied to the Laplacian eigenvalues.*

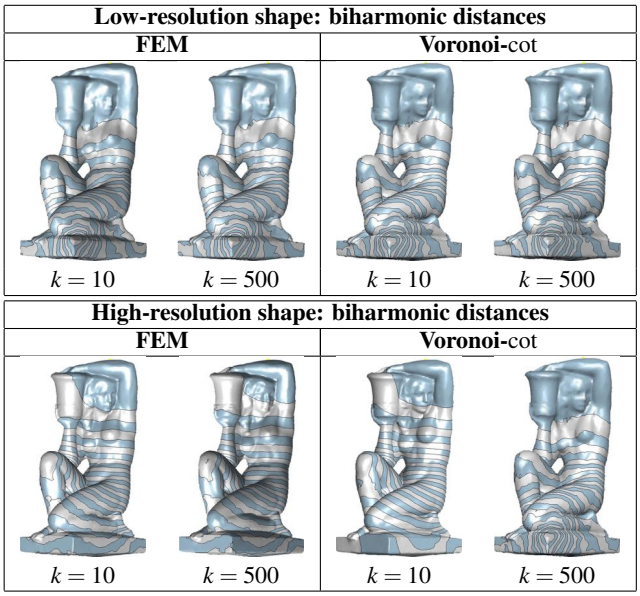

Figure 11: *Biharmonic distance on a surface at different resolutions, with different Laplacian weights and k eigenpairs.*

that qualitatively behave like the heat kernel on volumes. To improve the accuracy, we consider the volumetric Laplacian matrix of the input domain and compute the Padé-Chebyshev approximation of the induced heat kernel (Fig. 8).

6 Laplacian spectral distances

Distances can be defined directly on the input domain $\mathscr{M}$ (e.g., geodesic distances) or in the space of functions on $\mathscr{M}$ (e.g., random walks, biharmonic distances, diffusion and wave distances). For geometry processing and shape analysis, the distance

- satisfies the following properties: *positivity* ($d(\mathbf{p},\mathbf{q}) \leq 0$); *nullity* ($d(\mathbf{p},\mathbf{q}) = 0$ if and only if $\mathbf{p} \equiv \mathbf{q}$); *symmetry* ($d(\mathbf{p},\mathbf{q}) = d(\mathbf{q},\mathbf{p})$); *triangular inequality* ($d(\mathbf{p},\mathbf{q}) \leq d(\mathbf{p},\mathbf{r}) + d(\mathbf{r},\mathbf{q})$);
- should be *multi-scale* and *geometry-aware*, through the encoding of local/global features and geometric properties;
- should be *isometry-invariant* through a proper filtering of the Laplacian spectrum, *robust* to noise and domain discretization.

We introduce the spectral distances (Sect. 6.1), by filtering the Laplacian spectrum and as a generalization of the commute-time, biharmonic, diffusion and wave distances (Sect. 6.2). Then, we discuss their spectral discretization (Sect. 6.3), computation (Sect. 6.4), and comparison (Sect. 6.5).

6.1 Laplacian spectral kernels and distances

Starting from recent work on the geodesic and heat diffusion distances [Crane et al. 2013b; Patanè 2013], we address the definition of spectral distances on a manifold $\mathscr{N}$ by filtering its Laplacian spectrum [Bronstein and Bronstein 2011a; Patanè 2014]. Given a strictly positive *filter function* $\varphi : \mathbb{R}^+ \to \mathbb{R}$, let us consider the power series $\varphi(s) = \sum_{n=0}^{+\infty} \alpha_n s^n$. Noting that

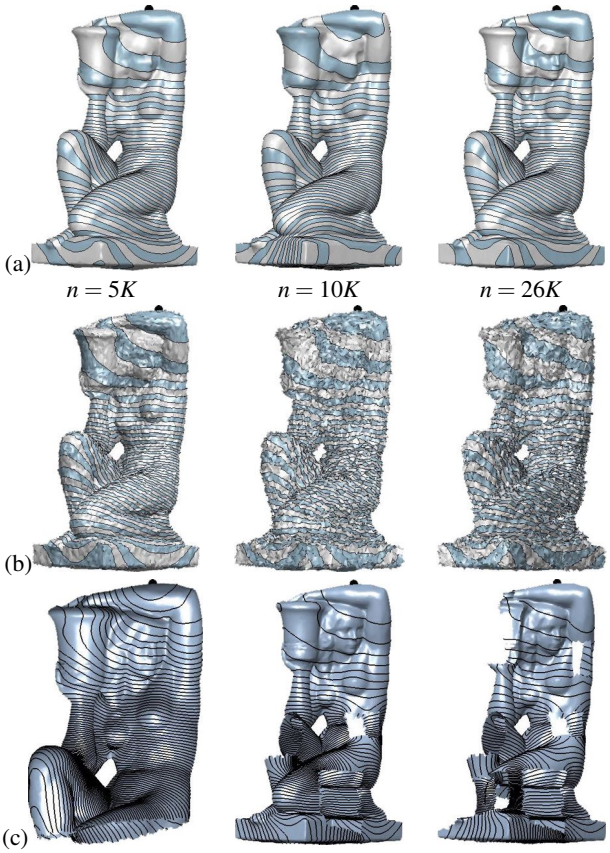

Figure 12: *Stability of the biharmonic distance from a source (black) point with respect to (a) sampling, (b) noise, (c) holes.*

$\Delta^i f = \sum_{n=0}^{+\infty} \lambda_n^i \langle f, \phi_n \rangle_2 \phi_n$, we define the *spectral operator* as

$$\Phi(f) = \sum_{n=0}^{+\infty} \alpha_n \Delta^n f = \sum_{n=0}^{+\infty} \varphi(\lambda_n) \langle f, \phi_n \rangle_2 \phi_n. \tag{3}$$

According to [Patané 2016], if the function $\tilde{\varphi}(s) := s^{1/2}\varphi(s)$ is integrable on $\mathbb{R}^+$ then the spectral operator is well-defined, linear, continuous, and $\Phi(f) = \langle K, f \rangle_2$, where $K(\mathbf{p},\mathbf{q}) = \sum_{n=0}^{+\infty} \varphi(\lambda_n)\phi_n(\mathbf{p})\phi_n(\mathbf{q})$ is the *spectral kernel*. Through the spectral operator, in $\mathscr{L}_2(\mathscr{N})$ we introduce the *spectral scalar product* and *distance* as

$$\begin{cases} \langle f, g \rangle := \langle \Phi(f), \Phi(g) \rangle_2 = \sum_{n=0}^{+\infty} \varphi^2(\lambda_n) \langle f, \phi_n \rangle_2 \langle g, \phi_n \rangle_2 \\ d^2(f,g) = \|f - g\|^2 = \sum_{n=0}^{+\infty} \varphi^2(\lambda_n) |\langle f - g, \phi_n \rangle|^2. \end{cases} \tag{4}$$

Indicating with $\delta_{\mathbf{p}}$ the function that takes value 1 at $\mathbf{p}$ and 0 otherwise, the *spectral distance* between $\mathbf{p}$, $\mathbf{q}$ is (Fig. 9)

$$\begin{aligned} d^2(\mathbf{p},\mathbf{q}) :&= \|\delta_{\mathbf{p}} - \delta_{\mathbf{q}}\|^2 \\ &= \sum_{n=0}^{+\infty} \varphi^2(\lambda_n) |\phi_n(\mathbf{p}) - \phi_n(\mathbf{q})|^2 \\ &= \|K(\mathbf{p},\cdot) - K(\mathbf{q},\cdot)\|_2^2, \end{aligned}$$

where the first row provides the spectral representation and the second row expresses the spectral distances in terms of the corresponding kernel.

Properties Analogously to the diffusion kernel, the spectral kernel satisfies the following properties: *non-negativity* ($K(\mathbf{p},\mathbf{p}) \geq 0$), *symmetry* ($K(\mathbf{p},\mathbf{q}) = K(\mathbf{q},\mathbf{p})$), and *positive semi-definiteness*:

$$\begin{aligned} 0 \leq \langle \Phi(f), f \rangle_2 &= \int_{\mathscr{N}\times\mathscr{N}} K(\mathbf{p},\mathbf{q}) f(\mathbf{p}) f(\mathbf{q}) \mathrm{d}\mathbf{p}\mathrm{d}\mathbf{q} \\ &= \sum_{n=0}^{+\infty} \varphi(\lambda_n) |\langle f, \phi_n \rangle_2|^2. \end{aligned}$$

We also mention the *square integrability* $\|K\|_2^2 = \sum_{n=0}^{+\infty} |\varphi(\lambda_n)|^2$, which is equivalent to the Parseval's equality and the *conservation*: $\int_{\mathscr{N}} K(\mathbf{p},\mathbf{q}) \mathrm{d}\mathbf{p} = 1$, which is a consequence of the Perron-Frobenious theorem.

Through the selected filter function and the Laplacian spectrum, we define the *spectral embedding* $\mathscr{E} : \mathscr{M} \to \ell_2$, which maps each point $\mathbf{p}$ to the sequence $\mathscr{E}(\mathbf{p}) := (\varphi(\lambda_n)\phi_n(\mathbf{p}))_{n=0}^{+\infty}$. The equality $d(\mathbf{p},\mathbf{q}) = \|\mathscr{E}(\mathbf{p}) - \mathscr{E}(\mathbf{q})\|_2$ shows that the spectral distances can be interpreted as Euclidean distances in the embedding space. Finally, the *spectral shape descriptor* $SD(\mathbf{p}) := \sum_{n=0}^{+\infty} |\varphi(\lambda_n)|^2 \phi_n^2(\mathbf{p})$ and *signature* $SE(\mathbf{p}) := (\varphi^{-1/2}(\lambda_n)\phi_n(\mathbf{p}))_{n=0}^{+\infty}$ generalize the diffusion descriptor and signature [Dey et al. 2010a; Sun et al. 2009] (Sect. 6.2.1).

Selection of the filter function The filter function is learned from a training data set [Aflalo et al. 2011] or chosen in such a way that the corresponding spectral distances satisfy the properties introduced at the beginning of Sect. 6. For instance (Fig. 10), selecting $\varphi_t(s) := \exp(-st), \exp(-ist)$ or $\varphi(s) := s^{-k/2}, s^{-1/2}$, we get the *heat diffusion*, *wave*, or *poly-harmonic*, *commute-time distances*, respectively. Mexican hat wavelets [Hou and Qin 2012] are generated by the filter $\varphi(s) := s^{1/2}\exp(-s^2)$ and in [Bronstein and Bronstein 2011b; Aubry et al. 2011] the filter function $\varphi(s) := \exp(is)$, $s \in [0, 2\pi]$, defines the wave kernel signature. The spectral distances associated with this periodic filter identify local shape features by separating the contribution of different frequencies and of the corresponding eigenfunctions.

Similarly to random walks [Ramani and Sinha 2013], we introduce multi-scale kernels by integrating the moment of order k of the differential operator $\Delta^\alpha \exp(-t\Delta^\alpha)$. In this case, the filter function is $\varphi(s) = t^k s^\alpha \exp(-ts^\alpha)$, where k scales the rate of diffusion and α controls the decay of the Laplacian eigenvalues to zero. The selection of the parameters α, k makes the multi-scale kernels more robust to geometric and topological noise; the integral over time also avoids the selection of the heat diffusion rate. The filter functions $\varphi_t(s) := [\cos^{-1/2}(\sqrt{s}t), s^{-1/4}\sin^{1/2}(\sqrt{s}t)]$ and $\varphi(s,t) = \exp(s^r t)$ are associated with the diffusion equations $(\partial_t^2 + \Delta)F(\cdot,t) = 0$ and $(\partial_t + \Delta^r)F(\cdot,t) = 0$, respectively. Finally, the filter function can be learned from a set of retrieval examples [Aflalo et al. 2011; Boscaini et al. 2015b].

6.2 Main examples of spectral distances

As special cases, we consider the diffusion (Sect. 6.2.1), commute-time and biharmonic distances (Sect. 6.2.2), and the approximation of the geodesic and transportation distances (Sect. 6.2.3).

6.2.1 Diffusion distances

The heat kernel induces the *diffusion distances*, whose spectral representation is $d_t^2(\mathbf{p},\mathbf{q}) = \sum_{n=0}^{+\infty} \exp(-\lambda_n t)|\phi_n(\mathbf{p}) - \phi_n(\mathbf{q})|^2$. Recalling that the heat kernel is self-adjoint with respect to the scalar product induced by the mass matrix $\mathbf{B}$, we define the *diffusive scalar product* $\langle \mathbf{f}, \mathbf{g} \rangle_t := \langle \mathbf{K}_t \mathbf{f}, \mathbf{g} \rangle_{\mathbf{B}}$ and express the *discrete*

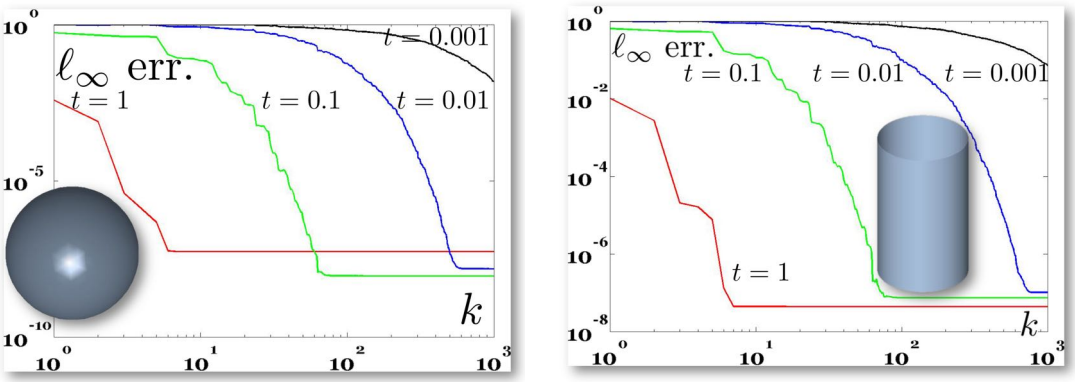

Figure 13: *ℓ_∞ error (y-axis) for the diffusion distance approximated with k (x-axis) Laplacian eigenpairs. For the Padé-Chebyshev method ($r = 5$) and all the scales, the ℓ_∞ error with respect to the ground-truth is lower than 8.9×10^{-6}.*

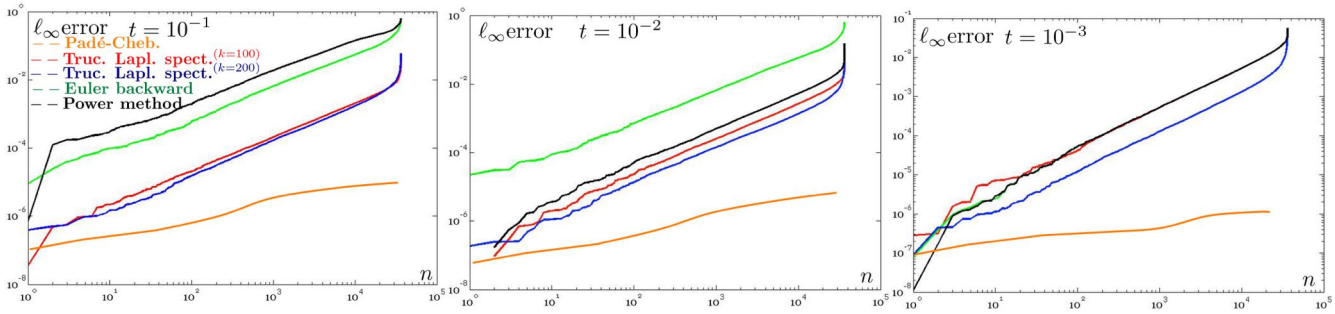

Figure 14: *ℓ_∞ error (y-axis) between the ground-truth diffusion distances on the cylinder, with a different sampling (x-axis). For different scales, the accuracy of the Padé-Chebyshev method ($r = 5$, orange) remains almost unchanged and higher than the truncated approximation with 100 and 200 eigenpairs (red, blue), the Euler backward (green) and power (black) methods.*

diffusion distances as $d_t(\mathbf{p}_i, \mathbf{p}_j) = \|\delta_{\mathbf{p}_i} - \delta_{\mathbf{p}_j}\|_t = \|\mathbf{K}_t(\mathbf{e}_i - \mathbf{e}_j)\|_{\mathbf{B}}$, where $\|\mathbf{f}\|_t^2 = \mathbf{f}^\top \mathbf{X}\mathbf{D}_t\mathbf{X}^\top\mathbf{B}\mathbf{f} = \sum_{i=1}^n \exp(-\lambda_i t)|\langle \mathbf{f}, \mathbf{x}_i\rangle_{\mathbf{B}}|^2$ is the *diffusion norm*.

Through the heat kernel, a shape is associated with a diffusion metric that measures the rate of connectivity among its points with paths of length t and characterizes the local/global geometric behavior with small/large values of t. This property has been used to define a multi-scale and isometry-invariant signatures [Bronstein et al. 2010a; Bronstein and Kokkinos 2010; Bronstein et al. 2011; Coifman and Lafon 2006; Dey et al. 2010b; Gebal et al. 2009; Lafon et al. 2006; Mèmoli and Sapiro 2005; Memoli 2009; Memoli 2011; Ovsjanikov et al. 2010; Raviv et al. 2010; Rustamov 2007; Mahmoudi and Sapiro 2009; Sun et al. 2009] and to rewrite the shape similarity problem as the comparison of two metric spaces. Main examples include the *heat kernel signature* $HKS(\mathbf{p}) := \sum_{n=0}^{+\infty} \exp(-\lambda_n t)|\phi_n(\mathbf{p})|^2$ and *descriptor* $HKD(\mathbf{p}) := (\lambda_n^{-1/2}\phi_n(\mathbf{p}))_{n=0}^{+\infty}$, and the *wave kernel signature* $WKS(\mathbf{p}) := \sum_{n=0}^{+\infty} \exp(-i\lambda_n t)|\phi_n(\mathbf{p})|^2$. Furthermore, the heat diffusion distance and kernel have been successfully applied to shape segmentation [de Goes et al. 2008]; the computation of the gradient of discrete functions [Luo et al. 2009]; and the multi-scale approximation of functions [Patanè and Falcidieno 2010]. The diffusion distance and kernel also play a central role in several applications, such as dimensionality reduction with spectral embeddings [Belkin and Niyogi 2003; Xiao et al. 2010]; data visualization [Belkin and Niyogi 2003; Hein et al. 2005; Roweis and Saul 2000; Tenenbaum et al. 2000], representation [Chapelle et al. 2003; Smola and Kondor 2003; Zhu et al. 2003], and classification [Ng et al. 2001; Shi and Malik 2000; Spielman and Teng 2007].

6.2.2 Commute-time and biharmonic distances

Integrating the diffusion distances with respect to t, we get the *commute-time distance*

$$d^2(\mathbf{p}, \mathbf{q}) = \frac{1}{2}\int_0^{+\infty} d_t^2(\mathbf{p}, \mathbf{q})\mathrm{d}t = \sum_{n=0}^{+\infty} \lambda_n^{-1}|\phi_n(\mathbf{p}) - \phi_n(\mathbf{q})|^2,$$

which is induced by the filter $\varphi(s) := s^{1/2}$ and is scale-invariant. We notice that this series can also be unbounded. While the diffusion distance estimates the connection of two points with respect to any random walk of length t, the commute-time distance measures this connection with respect to arbitrary random walks. The *biharmonic distances* [Ovsjanikov et al. 2012; Lipman et al. 2010; Rustamov 2011b] are induced by $\varphi(s) := s$ and provide a trade-off between a nearly-geodesic behavior for small distances and global shape-awareness for large distances, thus guaranteeing an intrinsic multi-scale characterization of the input shape. In Fig. 11, the approximation of the biharmonic kernel and distance with a subset of the Laplacian spectrum presents local artifacts, which are represented by isolated level sets and are reduced by increasing the number of eigenpairs without disappearing. In Fig. 12, the smooth and uniform distribution of the level sets of the biharmonic distance around the anchor point (black dot) confirms the stability of the spectrum-free approximation with respect to surface sampling, noise, and missing parts.

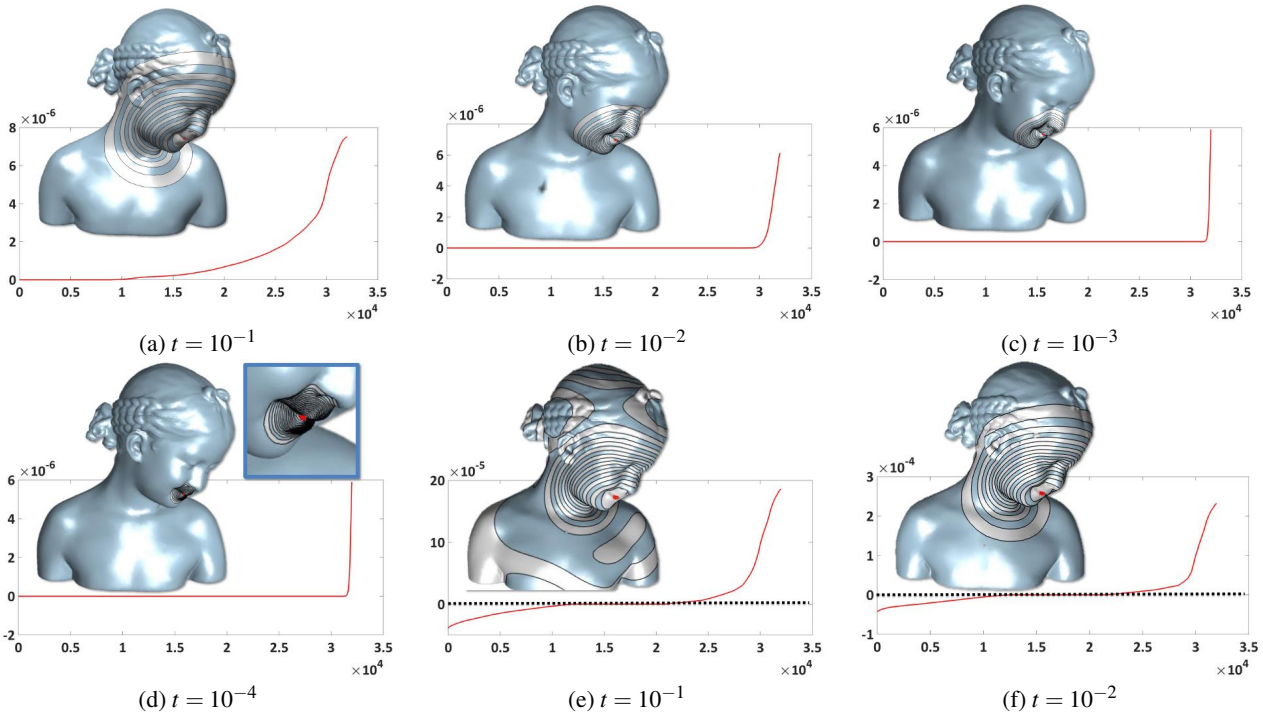

(a) $t = 10^{-1}$ (b) $t = 10^{-2}$ (c) $t = 10^{-3}$

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

Figure 15: *(a-d) Robustness of the Padé-Chebyshev approximation of the diffusion distances and (e,f) sensitiveness of truncated spectral approximation to the Gibbs phenomenon. At all scales (a-d), the distance values (red curve) computed with the Padé-Chebyshev approximation are positive; at large scales (e,f), the truncated spectral approximation is affected by the Gibbs phenomenon, as represented by the part of the plot below the zero line (black curve).*

6.2.3 Approximating geodesics and transportation distances with the heat kernel

According to [Varadhan 1967], the geodesic distance can be approximated as $d_G(\mathbf{p},\mathbf{q}) = -\lim_{t\to 0}(4t\log K_t(\mathbf{p},\mathbf{q}))$, where $K_t(\mathbf{p},\mathbf{q})$ is computed with the Padé-Chebyshev method. In fact, the truncated spectral approximation generally does not provide a value of $K_t(\mathbf{p},\mathbf{q})$, $t\to 0$, enough accurate to apply the Varadhan formula. Alternatively [Crane et al. 2013b], the geodesic distance $d_G(\mathbf{p},\mathbf{q})$ from the heat kernel values $K_t(\mathbf{p},\mathbf{q})$ as the scale tends to zero. More precisely, the geodesic distance d_G on $\mathscr{N}$ is approximated by computing the solution $F(\cdot,t)$ to the heat equation $(\partial_t+\Delta)F(\cdot,t)=0$ on $\mathscr{N}$, as $t\to 0^+$, normalizing the corresponding gradient $X = \nabla F(\cdot,t)/\|\nabla F(\cdot,t)\|_2$, and solving the equation $\Delta d_G = \mathrm{div}(X)$. This approximation is computationally efficient, capable of identifying different types of features by selecting different diffusion models, and robust to noise. The main difficulty is the tuning of the time scale with respect to the shape features; in fact, the selection of a large scale is generally associated with an over-smoothing of the geodesic values and local details.

In [Solomon et al. 2015], the *optimal transportation distances* have been approximated using the iterative Sinkhorn's method [Sinkhorn 1964] and the entropic regularization, thus reducing their computation to the solution of two sparse matrix equations that involve the heat kernel matrix. Instead of approximating the heat kernel with an implicit Euler integration [Desbrun et al. 1999], we can apply the Padé-Chebyshev approximation (Sect. 5.5.3) in order to improve the approximation accuracy at small scales and without modifying the overall approach.

6.3 Discrete spectral distances

Inserting the generalized eigensystem $\mathbf{LX} = \mathbf{BX\Lambda}$, with orthonormal eigenvectors $\mathbf{X}^\top\mathbf{BX} = \mathbf{I}$, in Eq. (3), its discretization is $\mathbf{K} = \mathbf{X}\varphi(\Lambda)\mathbf{X}^\top\mathbf{B}$, $\varphi(\Lambda) := \mathrm{diag}(\varphi(\lambda_i))_{i=1}^n$, and the corresponding *discrete spectral distances* are

$$d^2(\mathbf{p}_i,\mathbf{p}_j) = \|\mathbf{K}(\mathbf{e}_i-\mathbf{e}_j)\|_{\mathbf{B}}^2 = \sum_{l=1}^{n}\varphi^2(\lambda_i)|\langle\mathbf{x}_l,\mathbf{e}_i-\mathbf{e}_j\rangle_{\mathbf{B}}|^2. \quad (5)$$

The spectral representation of the kernel provides its link between the Laplacian matrix; i.e., $\tilde{\mathbf{L}}$ and $\mathbf{K}$ have the same eigenvectors and $(\varphi(\lambda_i))_{i=1}^n$ are the (filtered) Laplacian eigenvalues of $\mathbf{K}$.

In previous work, the spectral distances have been discretized as $d(\mathbf{p}_i,\mathbf{p}_j) = \|\mathbf{K}^\star(\mathbf{e}_i-\mathbf{e}_j)\|_2$ and with respect to the Euclidean scalar product, where $\mathbf{K}^\star := \mathbf{X}\varphi(\Lambda)\mathbf{X}^\top$ is the corresponding kernel. This last discretization does not take into account the intrinsic $\mathbf{B}$-scalar product, thus disregarding the geometry of the input data and the underlying generalized eigenproblem. Considering the linear FEM mass matrix $\mathbf{B}$ and noting that $B(i,j) = \langle\delta_{\mathbf{p}_i},\delta_{\mathbf{p}_j}\rangle_2$, where $\delta_{\mathbf{p}}$ is the function that takes value 1 at $\mathbf{p}$ and 0 otherwise, the $\mathbf{B}$-scalar product is the counterpart of the $\mathscr{L}_2(\mathscr{N})$ scalar product on the space of discrete functions on $\mathscr{M}$. The orthogonality of the Laplacian eigenvectors with respect to the $\mathbf{B}$-scalar product is crucial to encode the geometry of the surface underlying $\mathscr{M}$ in the spectral distances and makes its evaluation robust to surface sampling.

6.4 Computation of the spectral distances

Recalling that the computation of the Laplacian eigenpairs is numerically unstable in case of repeated eigenvalues (Sect. 4.3), the

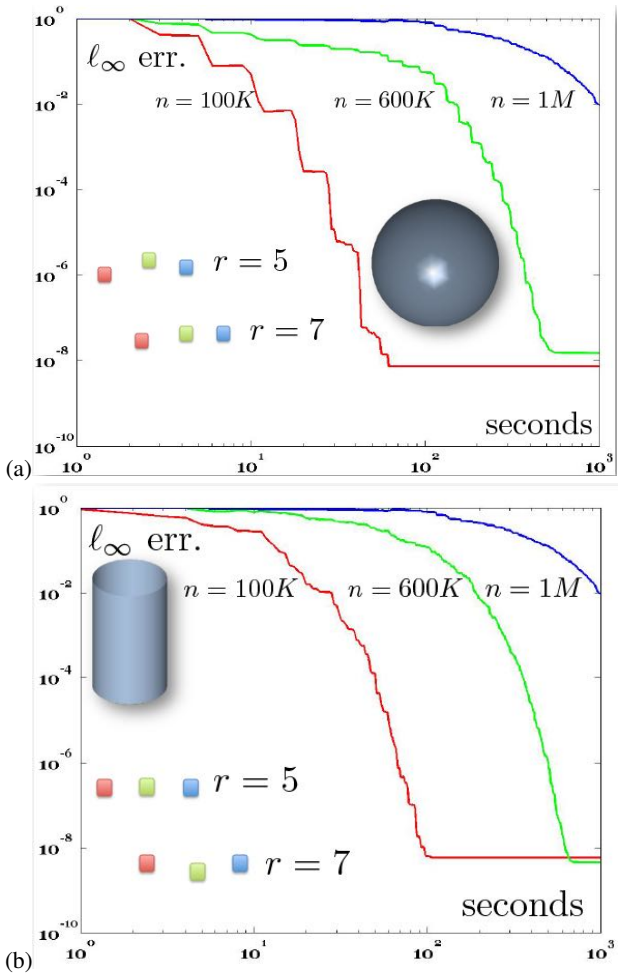

Figure 16: *Trade-off between accuracy (y-axis) and time (x-axis) for the Padé-Chebyshev ($r = 5, 7$) and truncated approximations ($k = 50$ eigenpairs) on the (a) sphere and (b) cylinder.*

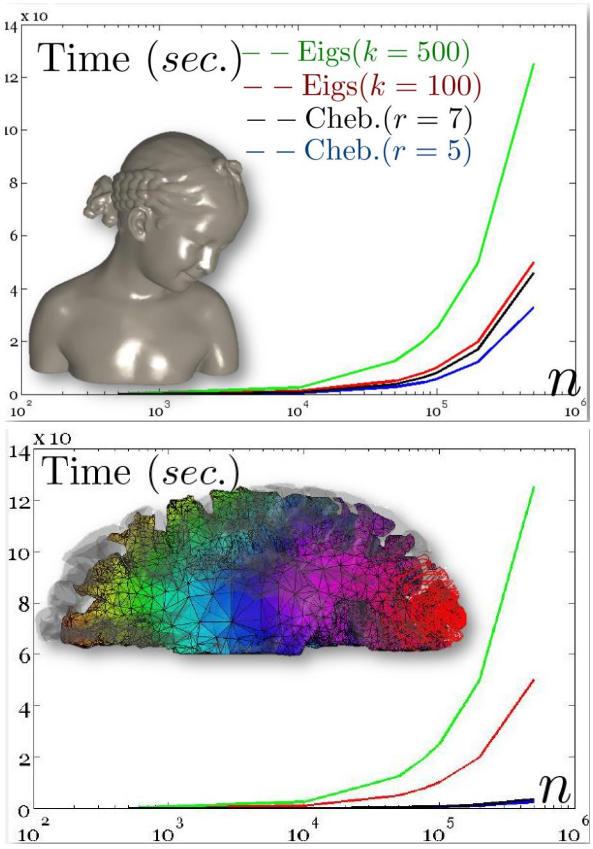

Figure 17: *Timings (in seconds) for the evaluation of the heat kernel on a domain with n points, approximated with $k = 100, 500$ eigenpairs* (Eigs) *and the Padé-Chebyshev approximation ($r = 7$).*

filter function should be chosen in such a way that the filtered Laplacian matrix does not have additional (if any) repeated eigenvalues. This condition is generally satisfied by choosing an injective filter. The selection of periodic filters, the expensive cost of the computation of the Laplacian spectrum, and the sensitiveness of multiple Laplacian eigenvalues to surface discretization are the main motivations for the definition of alternative approaches for the evaluation of the spectral distances and kernels. Among them, we discuss the truncated (Sect. 6.4.1) and spectrum-free (Sect. 6.4.2) approximations.

6.4.1 Truncated approximation

The computational limits for the evaluation of the whole Laplacian spectrum and the decay of the coefficients in Eq. (4) are the main reasons behind the approximation of the solution to the spectral distances as a truncated sum; i.e.,

$$\begin{cases} \Phi_k \mathbf{f} = \sum_{i=1}^{k} \varphi(\lambda_i) \langle \mathbf{f}, \mathbf{x}_i \rangle_{\mathbf{B}} \mathbf{x}_i \\ d^2(\mathbf{p}_i, \mathbf{p}_j) = \sum_{l=1}^{k} \varphi^2(\lambda_l) |\mathbf{x}_l^\top \mathbf{B} \mathbf{e}_i - \mathbf{x}_l^\top \mathbf{B} \mathbf{e}_j|^2, \end{cases}$$

where k is the number of selected eigenpairs. Even though the first k Laplacian eigenpairs are computed in super-linear time [Vallet and Lèvy 2008], the evaluation of the whole Laplacian spectrum is unfeasible for storage and computational cost, which are quadratic in the number of surface samples. Furthermore, the selection of filters that are periodic or do not decrease to zero motivates the need of defining a spectrum-free computation of the corresponding kernels and distances, which cannot be accurately approximated with the contribution of only a subpart of the Laplacian spectrum. The number of selected eigenpairs is heuristically adapted to the decay of the filter function and the approximation accuracy cannot be estimated without computing the whole spectrum.

6.4.2 Spectrum-free approximation

We now introduce the spectrum-free evaluation of the spectral distances, which is based on a polynomial or rational approximation of the filter.

Arbitrary filter: polynomial approximation For an arbitrary filter φ, the matrix $\varphi(\mathbf{A})$ is approximated by selecting a new function g such that the matrix $\tilde{\mathbf{A}} = g(\mathbf{A})$ approximates $\mathbf{A}$ and can be easily calculated. One of the main approaches for the approximation of a matrix function is through the truncated Taylor approximation [Golub and VanLoan 1989]. More precisely, given the power series representation $\varphi(s) = \sum_{n=0}^{+\infty} \alpha_n s^n$ defined on an open disk containing the spectrum of $\mathbf{A}$, we have that $\varphi(\mathbf{A}) = \sum_{n=0}^{+\infty} \alpha_n \mathbf{A}^n$. In this case, it is enough to consider the contribution of the first k terms in the sum and to compute the powers $(\mathbf{A}^i)_{i=1}^k$, through a binary powering [Van Loan 1979].

Let $[0, \lambda]$ be an interval that contains the spectrum of $\tilde{\mathbf{L}}$,

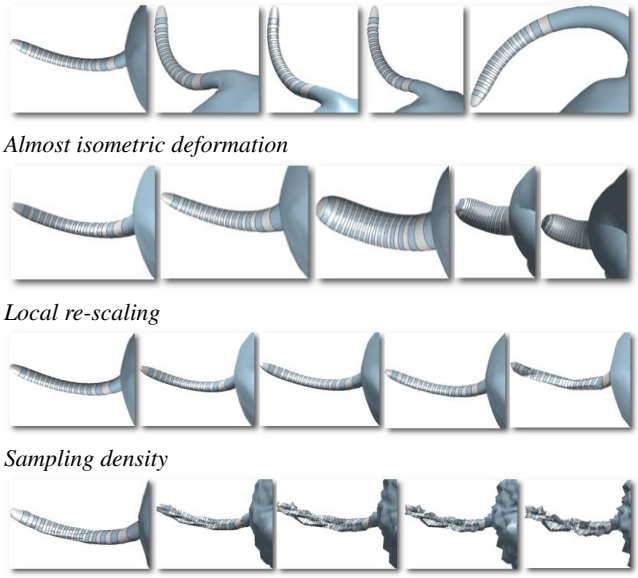

Figure 18: *Robustness of the computation of the linear FEM heat kernel from a seed point placed on the spike of the tail. The transformation strength increases from left to right.*

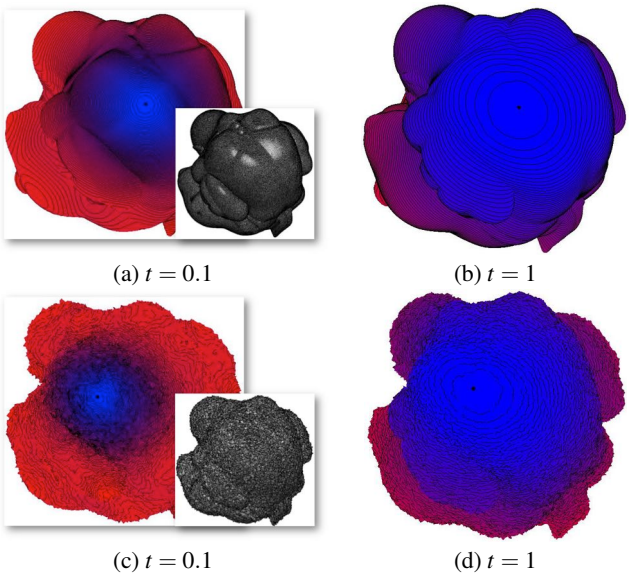

Figure 19: *Level sets of the linear FEM diffusion distance, computed using the Padé-Chebyshev approximation ($r := 7$), from a source point (black dot), with different values of t, on a (a,b) smooth and (c,d) noisy surface.*

where λ is the maximum eigenvalue, which is computed by the Arnoldi method [Golub and VanLoan 1989], or is set equal to the upper bound [Lehoucq and Sorensen 1996; Sorensen 1992] $\lambda_n \leq \min\{\max_i\{\sum_j \tilde{L}(i,j)\}, \max_j\{\sum_i \tilde{L}(i,j)\}\}$. Applying the Taylor approximation $\varphi(s) \approx p_r(s) := \sum_{n=0}^{r} \alpha_n s^n$ to the Laplacian matrix in $[0,\lambda]$, $\mathbf{Ke}_i$ is evaluated as (Algorithm 2)

$$\mathbf{Ke}_i \approx \sum_{n=0}^{r} \alpha_n (\mathbf{B}^{-1}\mathbf{L})^n \mathbf{e}_i = \alpha_0 \mathbf{e}_i + \sum_{n=1}^{r} \alpha_n \mathbf{g}_n, \tag{6}$$

where $\mathbf{g}_n$ satisfies the linear system $\mathbf{Bg}_{n+1} = \mathbf{Lg}_n$, $\mathbf{Bg}_1 = \mathbf{Le}_i$.

From the upper bound [Patanè 2014]

$$\left\| \varphi(\tilde{\mathbf{L}}) - \sum_{n=0}^{r} \alpha_n \tilde{\mathbf{L}}^n \right\|_2 \leq \frac{n}{(r+1)!} \left[\frac{\lambda_{\max}(\mathbf{L})}{\lambda_{\min}(\mathbf{B})} \right]^{r+1} \|\varphi^{(r+1)}(\tilde{\mathbf{L}})\|_2,$$

it follows that the approximation accuracy is mainly controlled by the degree of the Taylor approximation and the variation of the ratio between the maximum eigenvalue of $\mathbf{L}$ and the minimum eigenvalue of $\mathbf{B}$. If necessary, a higher approximation accuracy is achieved by slightly increasing the degree r. Finally, this computation of both the spectral kernel and distance is independent of the discretization of the input surface as a polygonal mesh or a point cloud. In case of a complex kernel, it is enough to apply the previous discussion to its real and imagery parts; e.g., for the wave kernel we consider the series $\sin(\tilde{\mathbf{L}}) = \sum_{n=0}^{+\infty} (-1)^n \tilde{\mathbf{L}}^{2n+1}/(2n+1)!$ and $\cos(\tilde{\mathbf{L}}) = \sum_{n=0}^{+\infty} (-1)^n \tilde{\mathbf{L}}^{2n}/(2n)!$.

Arbitrary filter: Padé-Chebyshev approximation For an arbitrary filter, we consider the rational Padé-Chebyshev approximation $p_r(s) = \frac{a_r(s)}{b_r(s)}$ of φ [Golub and VanLoan 1989] (Ch. 11) with respect to the $\mathscr{L}_\infty$ norm. Here, $a_r(\cdot)$ and $b_r(\cdot)$ are polynomials of degree equal to or lower than r. Let $p_r(s) = \sum_{i=1}^{r} \alpha_i (1+\beta_i s)^{-1}$ be the partial form of the Padé-Chebyshev approximation, where $(\alpha_i)_{i=1}^r$ are the weights and $(\beta_i)_{i=1}^r$ are the nodes of the r-point Gauss-Legendre quadrature rule [Golub and VanLoan 1989] (Ch. 11). The weights and nodes are precomputed for any degree of the rational polynomial [Carpenter et al. 1984]. Applying this approximation to the spectral kernel, we get that

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

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

$$(\mathbf{B} + \beta_i \mathbf{L})\,\mathbf{g}_i = \mathbf{Bf}, \quad i = 1, \ldots, r. \tag{7}$$

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

Properties According to [Moler and Van Loan 2003], the approximation of the matrix $\varphi(\tilde{\mathbf{L}})$ might be numerically unstable if $\|\tilde{\mathbf{L}}\|_2$ is large. From the bound $\|\mathbf{B}^{-1}\mathbf{L}\|_2 \leq \lambda_{\min}^{-1}(\mathbf{B})\lambda_{\max}(\mathbf{L})$, a well-conditioned mass matrix $\mathbf{B}$ guarantees that $\|\mathbf{B}^{-1}\mathbf{L}\|_2$ is bounded. Recalling that $\mathbf{X}^\top(\mathbf{B}+\beta_i\mathbf{L})\mathbf{X} = (\mathbf{I}+\beta_i\Lambda)$, $\{1+\beta_i\lambda_j\}_{j=1}^n$ are the eigenvalues of $(\mathbf{B}+\beta_i\mathbf{L})$ and its conditioning number is bounded by the constant $(1+\beta_{\max}\lambda_n)$, $\beta_{\max} := \max_{i=1,\ldots,n}|\beta_i|$. Indeed, the coefficient matrices in Eq. (6) are well-conditioned and specialized pre-conditioners [Krishnan et al. 2013] can be applied to further attenuate numerical instabilities.

Approximating an arbitrary filter function with a rational or a polynomial function of degree r, the evaluation of the corresponding spectral distance between two points is reduced to solve r sparse, symmetric, linear systems (c.f., Eq. (6)), whose coefficient matrices have the same structure and sparsity of the connectivity matrix of the input triangle mesh or of the k-nearest neighbor graph for a point set. Applying iterative solvers, such as the Jacobi, Gauss-Seidel, minimum residual methods [Golub and VanLoan 1989], and without extracting the Laplacian spectrum, the computational cost

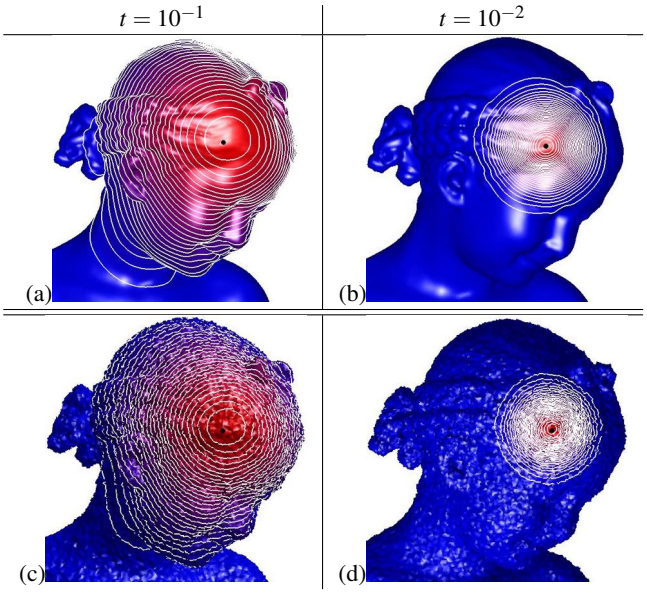

Figure 20: *Robustness of the Padé-Chebyshev approximation ($r = 7$) of the (a,c) diffusion kernel $\mathbf{K}_t\mathbf{e}_i$ and (b,d) distance at $\mathbf{p}_i$ (black dot) on a smooth and noisy triangulated surface.*

is $(\mathcal{O}r\tau(n))$, where $\tau(n)$ is the cost for the solution of a sparse linear system, which varies from $\mathcal{O}(n)$ to $\mathcal{O}(n^2)$, according to the sparsity of the coefficient matrix, and it is $\mathcal{O}(n\log n)$ in the average case.

The spectrum-free computation of the one-to-all distances $\{d(\mathbf{p}_i,\mathbf{p}_j)\}_{j=1}^n$ takes $\mathcal{O}(rn\tau(n))$ time; in fact, we solve the sparse linear system (6) with n different right-hand vectors $(\mathbf{e}_i - \mathbf{e}_j)$, $j = 1,\ldots,n$. Computing a fixed number k of eigenpairs in $\mathcal{O}(kn)$ time, the truncated spectral approximation of the one-to-all distance is evaluated in constant time for any filter. Indeed, the spectrum-free approach is competitive with respect to the truncated spectral approximation with $k(n) \geq r\tau(n)$ Laplacian eigenpairs. In the average case, $\tau(n) \approx n\log n$ and $k(n) \geq k_n$, $k_n = r\log n$. For instance, for a surface with $n = 10^4, 10^5, 10^6$ points and a degree $r = 5$, the number of eigenpairs is $k_n = 46, 58, 69$; in particular, this growth of k_n with respect to n is slow, as a matter of the logarithm in k_n.

6.5 Comparison and discussion

Fig. 13 reports the ℓ_∞ discrepancy (y-axis) between the diffusion distance on the sphere/cylinder and its approximation computed with the Padé-Chebyshev method and the truncated spectral approximation. In this case, the analytical expression of the Laplacian eigenfunctions on the sphere and cylinder has been used to compute the ground-truth distances [Patané 2016]. For small scales (e.g., $t = 10^{-2}, 10^{-3}$), the approximation error remains higher than 10^{-2}, with $k \leq 280$ eigenpairs; in fact, local shape features encoded by the heat kernel are recovered for a small t using the eigenvectors associated with high frequencies, thus requiring the computation of a large part of the Laplacian spectrum. For large scales (e.g., $t = 1, 10^{-1}$), increasing k strongly reduces the approximation error until it becomes almost constant and close to zero. In this case, the behavior of the heat kernel is mainly influenced by the Laplacian eigenvectors related to the smaller eigenvalues. Indeed, the truncated spectral representation generally requires a high number of eigenpairs and does not achieve the approximation accuracy of our approach, which remains lower than 8.9×10^{-6} for all the scales. According to [Vaxman et al. 2010], there are no theoretical guarantees on the approximation accuracy of the heat kernel provided by multi-resolution prolongation operators. Furthermore, a low-resolution sampling of the input surface might affect the resulting accuracy.

Algorithm 2: *Computation of the spectral distances.*

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

For all the scales (Fig. 14), the accuracy of the Padé-Chebyshev method is higher than the truncated approximation with k eigenpairs, $k = 1,\ldots,10^3$, the Euler backward method, and the power method. Reducing the scale, the accuracy of the Padé-Chebyshev remains almost unchanged while the other methods are affected by a larger discrepancy and tend to have an analogous behavior ($t = 10^{-4}$). Finally, the Euler backward method tends to over-smooth the solution, which converges to a constant as $k \to +\infty$, and the selection of the power m is guided by heuristics.

The truncated spectra1 approximation of the diffusion kernel is generally affected by the Gibbs phenomenon; i.e., small negative distance values. This phenomenon is more evident at small cases, which induce diffusion distances that decrease fast to zero and that are largely affected by small negative values. In fact, at small scales the diffusion distances decrease fast to zero and the negative values are no more compensated by the Laplacian eigenvectors related to smaller eigenvalues, as they are not included in the approximation (Fig. 15(e,f)). For the Padé-Chebyshev approximation (Fig. 15(a-d)), the distance values are positive at all the scales; in fact, we approximate the filter function without selecting a sub-part of the Laplacian spectrum.

Results in Figs. 16, 17 confirm that the diffusion distances at small scales generally require a number of eigenpairs that is much higher than the estimated value k_n. All tests have been performed on a 2.7 GHz Intel Core i7 Processor, with 8 GB memory. This case makes our computation of the one-to-all distance competitive with respect to its truncated approximation and useful to evaluate the distances for slowly-increasing (e.g., diffusion distances at small scales) or periodic filters or among seed points, as if happens for the evaluation of shape descriptors [Osada et al. 2002] and bags-of-features [Bronstein and Bronstein 2011b; Bronstein et al. 2011]. Here, the number of seeds is much lower than the number of samples and the higher accuracy of our computation improves the discrimination capabilities of descriptors based on spectral distances.

In our experiments, the analogous behavior of the level-sets of the heat kernel and diffusion distance confirm the robustness of the Padé-Chebyshev of the approximation with respect to sampling, discretization (Figs. 18, 19) and noise (Figs. 20, 21). A higher

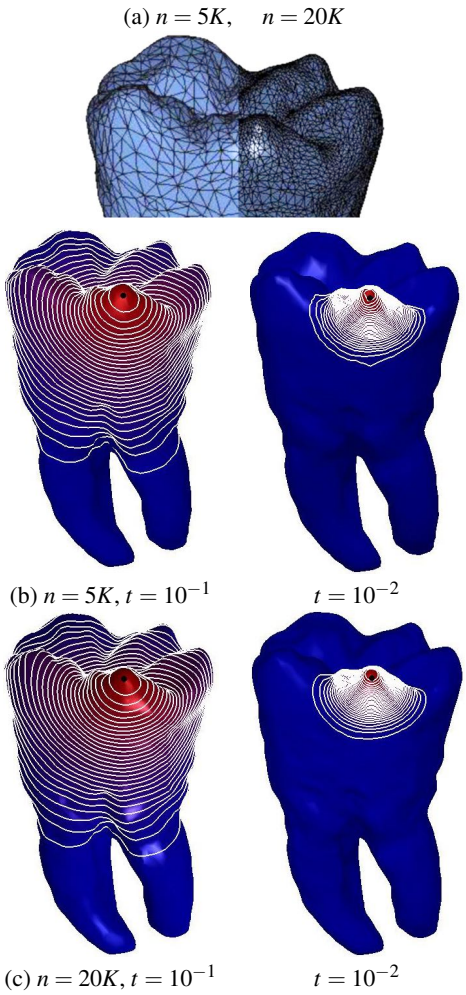

Figure 21: *(b,c) Robustness of the Padé-Chebyshev approximation ($r = 5$) of the heat kernel at different scales ($t = 10^{-1}, 10^{-2}$) with respect to (a) surface sampling ($n = 5K, 20K$).*

resolution of $\mathscr{M}$ improves the quality of the level-sets, which are always uniformly distributed and an increase of the noise magnitude does not affect the shape and distribution of the level sets.

Fixing the number of Laplacian eigenpairs makes the truncated spectral approximation of the one-to-all distances faster than ours but generally provides a lower approximation accuracy. Slowly-increasing filters and small scales for the diffusion distances also require the computation of a large number of Laplacian eigenpairs, thus reducing the gap between the computational cost of the proposed approximation of the one-to-all distances and previous work. An analogous discussion applies to prolongation operators, which compute the truncated spectral approximation on a lower resolution of the input shape. Furthermore, previous work has not addressed methods for the selection of the proper number of eigenpairs with respect to the target approximation accuracy, which cannot be estimated without computing the whole Laplacian spectrum. Finally, Figs. 22, 23 show the robustness of the spectrum-free computation with respect to a different shape discretization, non-manifold and bordered surfaces, topological noise. At large scales only (e.g., $t = 1$), the shape of the level sets of the heat kernel changes in a neighbor of the topological cut.

7 Applications

We now show how the Laplacian spectral properties and kernels have been used for smoothing in geometry processing (Sect. 7.1) and the definition of geometric basis functions (Sect. 7.2).

7.1 Smoothing

In real applications, the noisy component of the input data is due to a low quality of the discrete representations, unstable computations, and numerical approximations. Smoothing typically works in the function space and applies isotropic Laplacian filters [Dong et al. 2006; Ni et al. 2004; Taubin 1995] or bilateral smoothing operators to the function itself [Liu and Zhang 2007]. The isotropy of the Laplacian matrix indiscriminately smooths noise and topological features [Dong et al. 2006; Ni et al. 2004; Taubin 1995] without constraints on their relocations or cancellations. Constrained least-squares techniques [Sorkine et al. 2005] have been efficiently used to define compression schemes based on the selection of a set of anchors. While in [Sorkine et al. 2005] the choice of the constrained vertices is guided by the final approximation accuracy of the reconstructed surface, in [Patanè and Falcidieno 2009] the emphasis is on the preservation of the differential properties of f through the simplification of its critical points.

Unconstrained Laplacian smoothing According to [Patanè and Falcidieno 2009], the smooth approximation $\tilde{f}$ of a noisy scalar function $f : \mathscr{M} \to \mathbb{R}$ can be computed as the compromise between approximation accuracy and smoothness of the solution, we minimize the energy $\mathscr{F}(\tilde{\mathbf{f}}) := \varepsilon\|\tilde{\mathbf{f}} - \mathbf{f}\|^2_{\mathbf{B}} + \|\mathbf{L}\tilde{\mathbf{f}}\|^2_2$, whose normal equation is $(\mathbf{L}^\top\mathbf{L} + \varepsilon\mathbf{B})\tilde{\mathbf{f}} = \varepsilon\mathbf{B}\mathbf{f}$. Since the coefficient matrix is sparse and positive definite, $\tilde{\mathbf{f}}$ is uniquely defined and it is efficiently computed through direct or iterative solvers of sparse linear systems [Golub and VanLoan 1989]. Finally, the spectral representation is $\tilde{\mathbf{f}} = \sum_{i=1}^{n}(\lambda_i^2 + \varepsilon)^{-1}\langle \mathbf{f}, \mathbf{x}_i\rangle_{\mathbf{B}}\mathbf{x}_i$, where the smoothing term $(\lambda_i^2 + \varepsilon)^{-1}$ filters out the contributions to the solution corresponding to the high eigenvalues (Fig. 24).

Laplacian smoothing with interpolating constraints According to [Patanè et al. 2007], we can include a set of constraints on the topological features of f to be preserved (Fig. 25). In this case, we define the smooth scalar function $\tilde{f}$ as the solution of the constrained minimization problem $\min_{\tilde{\mathbf{f}}\in\mathbb{R}^n}\|\mathbf{L}\tilde{\mathbf{f}}\|_2$, $\tilde{f}(\mathbf{p}_i) := f(\mathbf{p}_i)$, $i \in \mathscr{I}$. This problem is equivalent to minimizing the least-squares error $\|\tilde{\mathbf{L}}\mathbf{x} - \mathbf{b}\|_2$, $\mathbf{x} \in \mathbb{R}^{n-k}$. Here, $\tilde{\mathbf{L}}$ is the $(n-k)\times(n-k)$ matrix achieved by removing the i^{th}-row and i^{th}-column of $\mathbf{L}$, $i \in \mathscr{I}$, and the entries of the constant term $\mathbf{b} \in \mathbb{R}^{n-k}$ are $\sum_{j\in N(i)\cap\mathscr{I}} l_{ij} f(\mathbf{p}_j)$, $i \in \mathscr{I}^C$.

Smoothing medical data In medical applications, the heat kernel is central in diffusion filtering and smoothing of images [Alvarez et al. 1992; F. Cattè and Coll 1992; Perona and Malik 1990; Spira et al. 2007; Tasdizen et al. 2002; Witkin 1983], 3D shapes [Bajaj and Xu 2002; Guskov et al. 1999], and anatomical surfaces [Chung et al. 2005; Kim et al. 2012; Wang et al. 2013; Patanè 2015]. Fig. 26 compares the diffusion smoothing of a noisy data set computed with the Padé-Chebyshev approximation of degree $r = 7$ and the truncated approximation with k Laplacian eigenparis. A low number of eigenpairs does not preserve shape details; increasing k reconstructs the surface noise. The ℓ_∞ error between (a) and the smooth approximation of (b) is lower than 1% for (c) the Padé-Chebyshev method and (d) varies from 12% ($k = 100$) up to 13% ($k = 1K$) for the truncated spectral approximation.

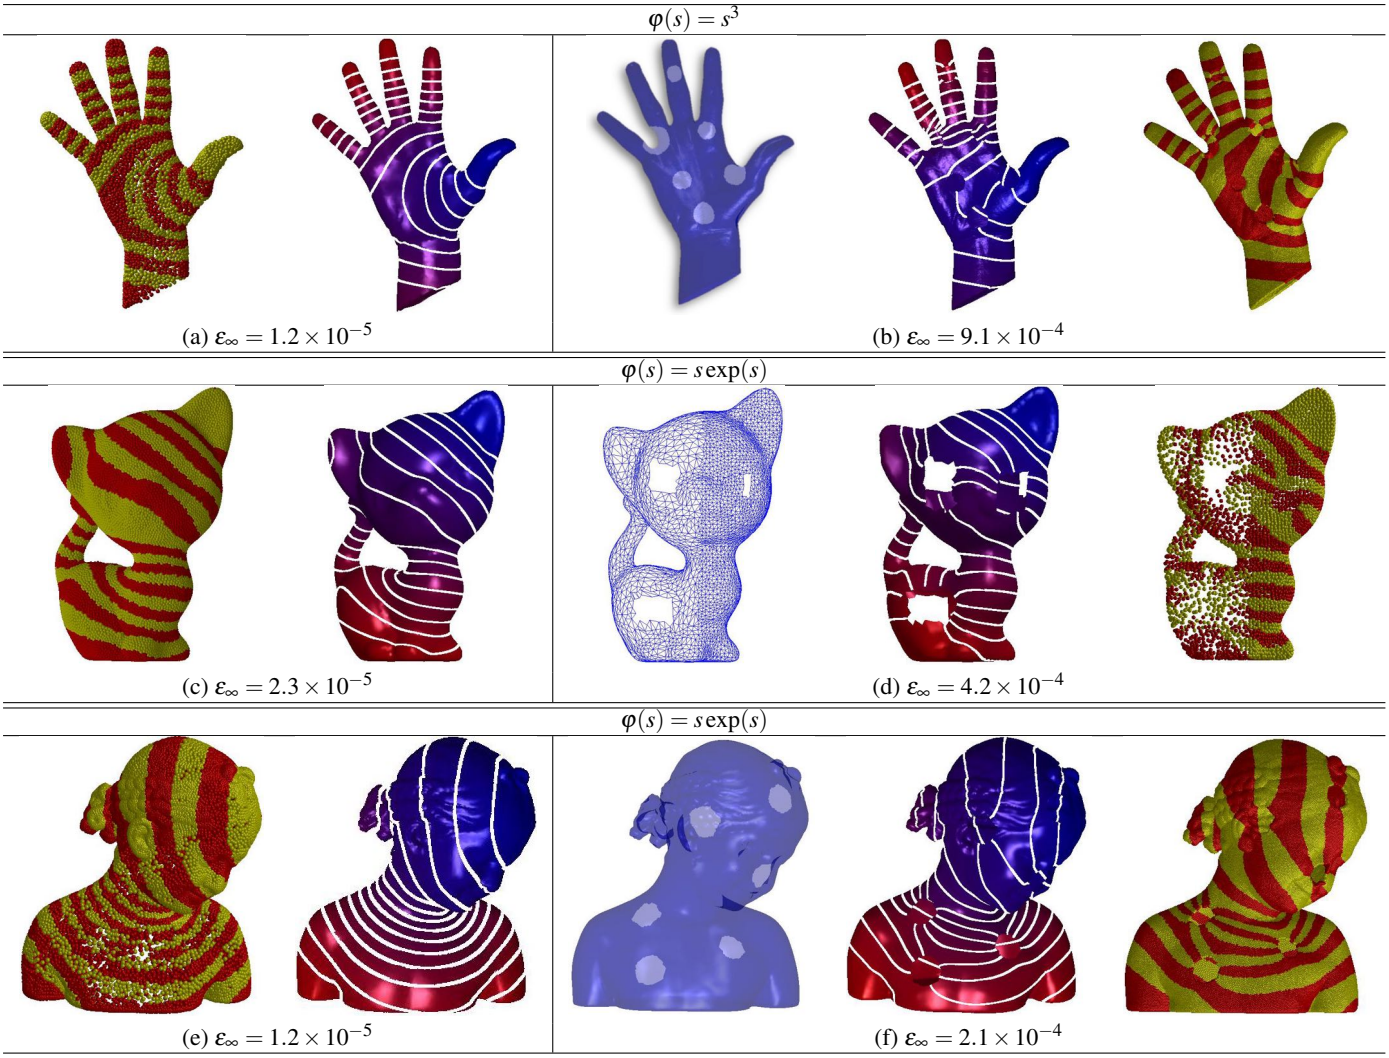

Figure 22: *Distances computed with the Padé-Chebyshev method ($r = 5$) on (a,c,e) regularly-sampled and (b,d,f) irregularly-sampled (left) meshes and (right) point sets with holes. To improve the visualization, points are represented as spheres.*

7.2 Laplacian and diffusion basis functions

Even though the Laplacian eigenvectors are intrinsic to the input surface, they can be computed only for a small set of eigenvalues and do not provide a flexible alignment of the function behavior to specific shape features. The geometry-aware functions [Sorkine et al. 2005] provide a computationally efficient way to encode the local geometric information of $\mathscr{M}$; and a similar approach can be applied to define more general classes of basis functions on a given shape. Applying the heat kernel matrix, we can define the *diffusion basis* $\mathscr{B} := \{\mathbf{K}_t \mathbf{e}_i\}_{i=1}^n$, whose elements have a smooth behavior on $\mathscr{M}$ and are intrinsically defined by $\mathscr{M}$ (Fig. 18). To define a set of *shape-driven* canonical basis functions, as feature points $\{\mathbf{p}_i\}_{i \in \mathscr{A}}$ of a 3D shape we select the maxima and minima of the Laplacian eigenfunctions related to the smallest eigenvalues [Ruggeri et al. 2010] or of the auto-diffusion functions [Gebal et al. 2009]. Finally, the definition of different basis function is also fundamental to define functions between shapes [Gal and Cohen-Or 2006; Gelfand et al. 2005; Hamza and Krim 2003; Li and Guskov 2005; Mèmoli and Sapiro 2005; Osada et al. 2002; Ovsjanikov et al. 2012; Ruggeri et al. 2010].

8 Conclusions

Our survey provides a common background on the definition and numerical computation of Laplacian spectral kernels and distances for geometry processing and shape analysis, as a generalization of the well-known biharmonic, diffusion, and wave distances. To support the reader in the selection of the most appropriate with respect to shape representation, computational resources, and target application, all the reviewed numerical schemes have been discussed and compared in terms of robustness, approximation accuracy, and computational cost.

From the numerical point of view, the evaluation of full shape descriptors (e.g., heat kernel values among all the input points) is partially limited in case of densely sampled shapes, due to the expensive computational time and storage overhead. Indeed, the appropriate selection of seed points on the input domain and the conversion of the spectral descriptor to a sparse approximation are still crucial steps for the evaluation of full shape descriptors. Furthermore, the robustness of the spectrum-free computation to sampling and missing parts suggests the use of the spectral distances and descriptor for partial shape matching.

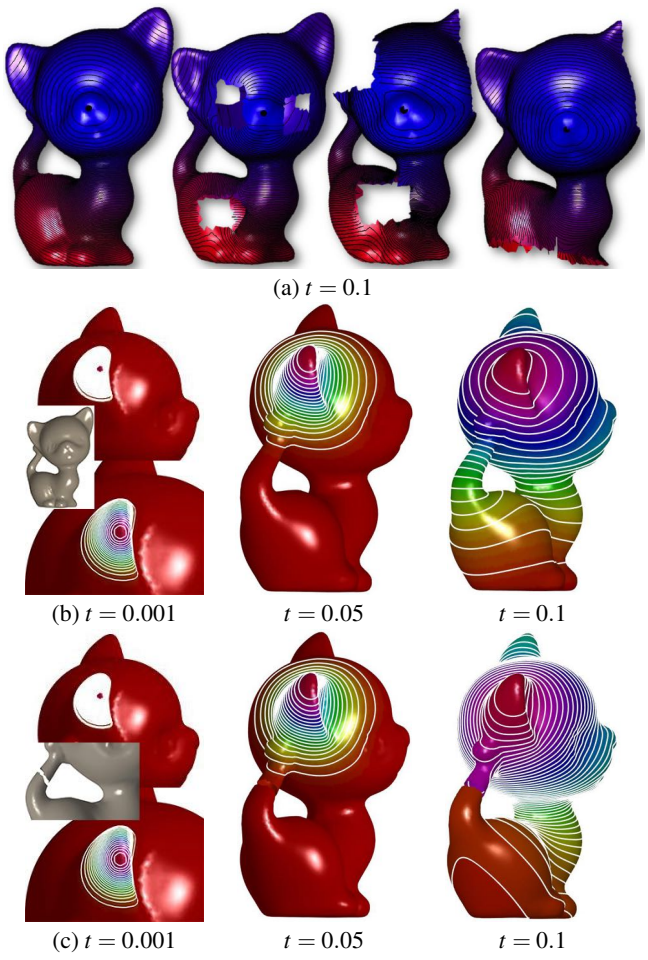

Figure 23: *Robustness of the Padé-Chebyshev approximation of the linear FEM (a) diffusion distance on partially-sampled surfaces and (b,c) heat kernel on smooth and topologically noisy surfaces (cut on the kitten tail), respectively.*

From the point of view of the definition of the spectral kernels and distances, we have discussed the general properties of the filter that guarantee their well-posedness, intrinsic and invariance properties. A deeper analysis of the filter properties with respect to the induced spectral distances would be beneficial to improve current results on surface watermarking and shape comparison. Finally, learning the filter from the geometric properties of a given class of data is an efficient, but partially unexplored, way to address shape segmentation, comparison, and more generally manifold learning.

Acknowledgements Shapes are courtesy of AIM@SHAPE Repository and tet-meshes were generated by the TETGEN software (`http://wias-berlin.de/software/tetgen/`).

References

AFLALO, Y., BRONSTEIN, A. M., BRONSTEIN, M. M., AND KIMMEL, R. 2011. Deformable shape retrieval by learning diffusion kernels. In *Scale space and Variational methods in computer Vision*, 689–700.

ALEXA, M., AND WARDETZKY, M. 2011. Discrete Laplacians on general polygonal meshes. *ACM Trans. on Graphics 30*, 4.

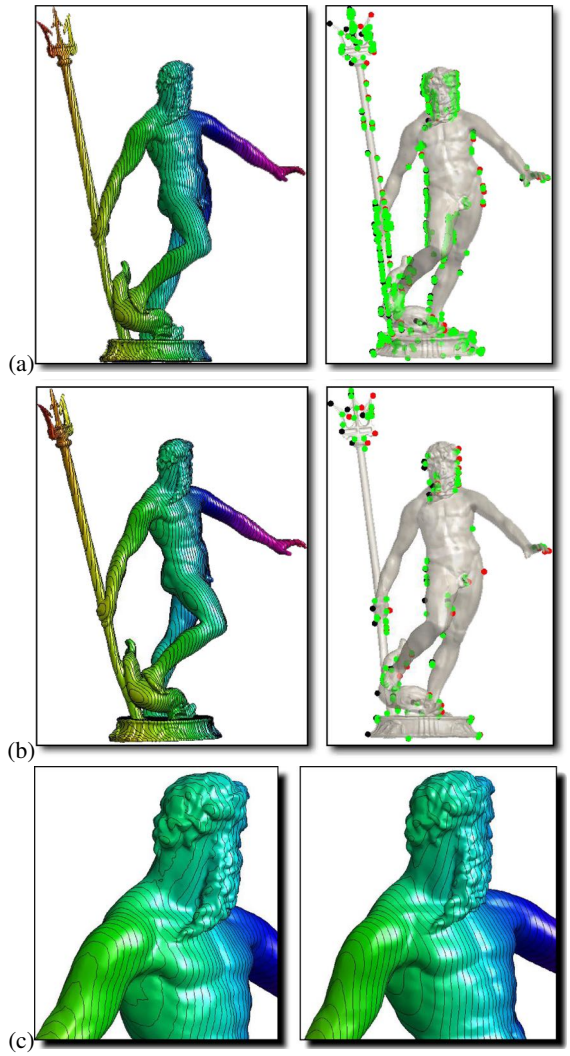

Figure 24: *(a) Noisy ($m = 127$, $M = 57$, $s = 188$) and (b) smoothed scalar function ($m = 12$, $M = 14$, $s = 30$). (c) Zoom-in on the level sets of the input (left) and smoothed (right) function.*

ALLAIRE, G. 2007. *Numerical Analysis and Optimization: An Introduction to Mathematical Modelling and Numerical Simulation*. Numerical Mathematics and Scientific Computation. Oxford University Press, Incorporated.

ALLIEZ, P., COHEN-STEINER, D., YVINEC, M., AND DESBRUN, M. 2005. Variational tetrahedral meshing. *ACM Trans. on Graphics 24*, 3, 617–625.

ALPERT, C. J., KAHNG, A. B., AND YAO, S.-Z. 1999. Spectral partitioning with multiple eigenvectors. *Discrete Applied Mathematics 90*, 1-3, 3–26.

ALVAREZ, L., LIONS, P.-L., AND MOREL, J.-M. 1992. Image selective smoothing and edge detection by nonlinear diffusion. *SIAM Journal on Numerical Analysis 29*, 3, 845–866.

ANDREUX, M., RODOLA, E., AUBRY, M., AND CREMERS, D. 2014. Anisotropic Laplace-Beltrami operators for shape analysis. In *Sixth Workshop on Non-Rigid Shape Analysis and Deformable Image Alignment (NORDIA)*.

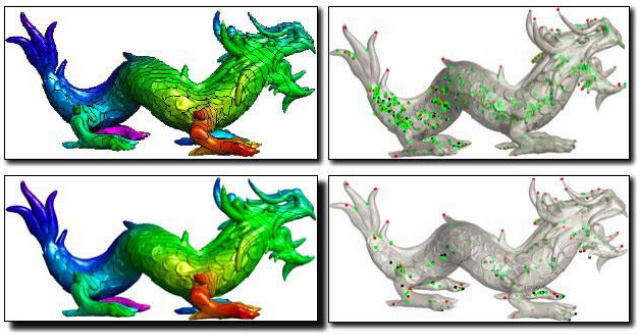

Figure 25: *Smoothing with interpolating constraints. The level sets and critical points of the input and smoothed scalar function are shown in the first and second row. The $\mathscr{L}_\infty$-error is* 0.08.

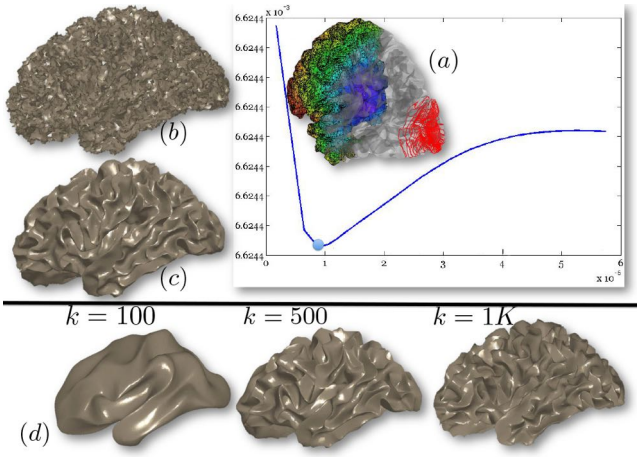

Figure 26: *(a) Input mesh and L-curve of the approximation accuracy (y-axis) versus the solution smoothness (x-axis). (b) Data set achieved by adding a Gaussian noise to (a). Diffusion smoothing computed with (c) the Padé-Chebyshev approximation ($r = 7$) and (d) the truncated approximation with k Laplacian eigenparis. A lower number of eigenpairs smooths local details; increasing k reconstructs the noisy component. The ℓ_∞ error between the ground-truth (a) and the smooth approximation of (b) is lower than* 1% *for the Padé-Chebyshev method (c) and varies from* 12% *($k = 100$) to* 13% *($k = 1K$) for the truncated approximation (d).*

ARYA, S., MOUNT, D. M., NETANYAHU, N. S., SILVERMAN, R., AND WU, A. Y. 1998. An optimal algorithm for approximate nearest neighbor searching fixed dimensions. *Journal of the ACM 45*, 6, 891–923.

AUBRY, M., SCHLICKEWEI, U., AND CREMERS, D. 2011. The wave kernel signature: a quantum mechanical approach to shape analysis. In *IEEE Computer Vision Workshops*, 1626–1633.

BAJAJ, C. L., AND XU, G. 2002. Anisotropic diffusion of subdivision surfaces and functions on surfaces. *ACM Trans. on Graphics 22*, 4–32.

BARNARD, S. T., POTHEN, A., AND SIMON, H. D. 1993. A spectral algorithm for envelope reduction of sparse matrices. In *Proc. of the ACM Supercomputing*, 493–502.

BELKIN, M., AND NIYOGI, P. 2003. Laplacian eigenmaps for dimensionality reduction and data representation. *Neural Computations 15*, 6, 1373–1396.

BELKIN, M., AND NIYOGI, P. 2006. Convergence of Laplacian eigenmaps. In *Neural Information Processing Systems*, 129–136.

BELKIN, M., AND NIYOGI, P. 2008. Towards a theoretical foundation for Laplacian-based manifold methods. *Journal of Computer System Sciences 74*, 8, 1289–1308.

BELKIN, M., SUN, J., AND WANG, Y. 2008. Discrete Laplace operator on meshed surfaces. In *Proc. of the Twenty-fourth Annual Symp. on Computational Geometry*, 278–287.

BELKIN, M., SUN, J., AND WANG, Y. 2009. *Constructing Laplace Operator from Point Clouds in $\mathbb{R}^d$*. ACM Press, ch. 112, 1031–1040.

BENTLEY, J. L. 1975. Multidimensional binary search trees used for associative searching. *Commun. ACM 18*, 9, 509–517.

BERARD, P., BESSON, G., AND GALLOT, S. 1994. Embedding Riemannian manifolds by their heat kernel. *Geometric and Functional Analysis 4*, 4, 373–398.

BOBENKO, A. I., AND SPRINGBORN, B. A. 2007. A discrete Laplace–Beltrami operator for simplicial surfaces. *Discrete & Computational Geometry 38*, 4, 740–756.

BOSCAINI, D., EYNARD, D., KOUROUNIS, D., AND BRONSTEIN, M. M. 2015. Shape-from-operator: recovering shapes from intrinsic operators. *Computer Graphics Forum 34*, 2, 265–274.

BOSCAINI, D., MASCI, J., MELZI, S., BRONSTEIN, M. M., CASTELLANI, U., AND VANDERGHEYNST, P. 2015. Learning class-specific descriptors for deformable shapes using localized spectral convolutional networks. *Computer Graphics Forum 34*, 5, 13–23.

BRONSTEIN, M. M., AND BRONSTEIN, A. M. 2011. Shape recognition with spectral distances. *IEEE Trans. on Pattern Analysis and Machine Intelligence 33*, 5, 1065–1071.

BRONSTEIN, M., AND BRONSTEIN, A. 2011. Shape recognition with spectral distances. *IEEE Trans. on Pattern Analysis and Machine Intelligence 33*, 5, 1065 –1071.

BRONSTEIN, M., AND KOKKINOS, I. 2010. Scale-invariant heat kernel signatures for non-rigid shape recognition. In *IEEE Conf. on Computer Vision and Pattern Recognition*, 1704 –1711.

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

BRONSTEIN, A. M., BRONSTEIN, M. M., BUSTOS, B., CASTELLANI, U., CRISANI, M., FALCIDIENO, B., GUIBAS, L. J., I. KOKKINOS, V. M., ISIPIRAN, I., OVSJANIKOV, M., PATANÈ, G., SPAGNUOLO, M., AND SUN, J. 2010. SHREC 2010: robust feature detection and description benchmark. *Eurographics Workshop on 3D Object Retrieval*.

BRONSTEIN, A. M., BRONSTEIN, M. M., CASTELLANI, U., FALCIDIENO, B., FUSIELLO, A., GODIL, A., GUIBAS, L., KOKKINOS, I., LIAN, Z., OVSJANIKOV, M., PATANÈ, G., SPAGNUOLO, M., AND TOLDO, R. 2010. SHREC 2010: robust large-scale shape retrieval benchmark. *Eurographics Workshop on 3D Object Retrieval*.

BRONSTEIN, A. M., BRONSTEIN, M. M., OVSJANIKOV, M., AND GUIBAS, L. J. 2011. Shape Google: geometric words and expressions for invariant shape retrieval. *ACM Trans. on Graphics 30*, 1.

BRONSTEIN, M., CASTELLANI, U., AND A., B. 2012. Diffusion geometry in shape analysis. *Eurographics Tutorial*.

CAELLI, T., AND K., S. 2004. An eigenspace projection clustering method for inexact graph matching. *IEEE Trans. on Pattern Analysis and Machine Intelligence 26*, 4, 515–519.

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

CHAPELLE, O., WESTON, J., AND SCHÖLKOPF, B. 2003. Cluster kernels for semi-supervised learning. In *Neural Information Processing Systems*, vol. 15, 585–592.

CHUANG, M., LUO, L., BROWN, B. J., RUSINKIEWICZ, S., AND KAZHDAN, M. 2009. Estimating the Laplace-Beltrami operator by restricting 3D functions. In *Proc. of the Symp. on Geometry Processing*, 1475–1484.

CHUNG, M. K., ROBBINS, S. M., DALTON, F. K. M., DAVIDSON, C. R. J., ALEX, A. L., AND EVANS, A. C. 2005. Cortical thickness analysis in autism with heat kernel smoothing. *NeuroImage 25*, 1256–1265.

CHUNG, F. R. K. 1997. *Spectral graph theory*. American Mathematical Society.

CLARENZ, U., DIEWALD, U., AND RUMPF, M. 2000. Anisotropic geometric diffusion in surface processing. In *IEEE Visualization*, 397–405.

CODY, W. J., MEINARDUS, G., AND VARGA, R. S. 1969. Chebyshev rational approximations to $\exp(-z)$ in $(0,+\infty)$ and applications to heat-conduction problems. *Journal of Approximation Theory 2*, 50–65.

COIFMAN, R. R., AND LAFON, S. 2006. Diffusion maps. *Applied and Computational Harmonic Analysis 21*, 1, 5–30.

CRANE, K., PINKALL, U., AND SCHRÖDER, P. 2013. Robust fairing via conformal curvature flow. *ACM Trans. on Graphics 32*, 4, 61.

CRANE, K., WEISCHEDEL, C., AND WARDETZKY, M. 2013. Geodesics in heat: A new approach to computing distance based on heat flow. *ACM Trans. on Graphics 32*, 5 (Oct.), 152:1–152:11.

DE GOES, F., GOLDENSTEIN, S., AND VELHO, L. 2008. A hierarchical segmentation of articulated bodies. *Computer Graphics Forum 27*, 5, 1349–1356.

DESBRUN, M., MEYER, M., SCHRÖDER, P., AND BARR, A. H. 1999. Implicit fairing of irregular meshes using diffusion and curvature flow. In *ACM Siggraph*, 317–324.

DEY, T. K., AND SUN, J. 2005. An adaptive MLS surface for reconstruction with guarantees. In *ACM Symp. on Geometry Processing*, 43–52.

DEY, T. K., LI, K., LUO, C., RANJAN, P., SAFA, I., AND WANG, Y. 2010. Persistent heat signature for pose-oblivious matching of incomplete models. *Computer Graphics Forum 29*, 5, 1545–1554.

DEY, T. K., RANJAN, P., AND WANG, Y. 2010. Convergence, stability, and discrete approximation of Laplace spectra. *ACM Symp. Discrete Algorithms*, 650–663.

DÌAZ, J., PETIT, J., AND SERNA, M. 2002. A survey of graph layout problems. *ACM Computing Surveys 34*, 3, 313–356.

DONG, S., KIRCHER, S., AND GARLAND, M. 2005. Harmonic functions for quadrilateral remeshing of arbitrary manifolds. *Computer Aided Geometric Design 22*, 5, 392–423.

DONG, S., BREMER, P.-T., GARLAND, M., PASCUCCI, V., AND HART, J. C. 2006. Spectral surface quadrangulation. *ACM Siggraph*, 1057–1066.

DYER, R., ZHANG, H., AND MÖLLER, T. 2007. Delaunay mesh construction. In *Proceedings of Eurographics Symposium on Geometry Processing*, 273–282.

F. CATTÈ, P.-L. LIONS, J.-M. M., AND COLL, T. 1992. Image selective smoothing and edge detection by nonlinear diffusion. *SIAM Journal on Numerical Analysis 29*, 1, 182–193.

FALGOUT, R. D. 2006. An introduction to algebraic multigrid. *Computing in Science and Engineering 8*, 6 (Nov.), 24–33.

FENG, P., AND WARREN, J. 2012. Discrete bi-Laplacians and biharmonic b-splines. *ACM Trans. Graphics 31*, 4 (July), 115:1–115:11.

FIEDLER, M. 1973. Algebraic connectivity of graphs. *Czechoslovak Mathematical Journal 23*, 98, 298–305.

FLOATER, M. S. 2003. Mean value coordinates. *Computer Aided Geometric Design 20*, 1, 19–27.

FOUSS, F., PIROTTE, A., AND SAERENS, M. 2005. A novel way of computing similarities between nodes of a graph, with application to collaborative recommendation. In *IEEE/WIC/ACM Intern. Conf. on Web Intelligence*, 550–556.

FOWLKES, C., BELONGIE, S., CHUNG, F., AND MALIK, J. 2004. Spectral grouping using the nystrom method. *IEEE Trans. Pattern Analysis and Machine Intelligence 26*, 2, 214–225.

GAL, R., AND COHEN-OR, D. 2006. Salient geometric features for partial shape matching and similarity. *ACM Trans. on Graphics 25*, 1, 130–150.

GEBAL, K., BÆRENTZEN, J. A., AANÆS, H., AND LARSEN, R. 2009. Shape analysis using the auto diffusion function. *Computer Graphics Forum 28*, 5, 1405–1413.

GELFAND, N., MITRA, N. J., GUIBAS, L. J., AND POTTMANN, H. 2005. Robust global registration. In *Proc. of the Symp. on Geometry Processing*, 197.

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

GOLUB, G., AND VANLOAN, G. 1989. *Matrix Computations*. John Hopkins University Press, 2nd Edition.

GORDON, C. S., AND SZABO, Z. I. 2002. Isospectral deformations of negatively curved riemannian manifolds with boundary which are not locally isometric. *Duke Math. J. 113*, 2 (06), 355–383.

GOTSMAN, C., GU, X., AND SHEFFER, A. 2003. Fundamentals of spherical parameterization for 3D meshes. In *ACM Siggraph 2003*, 358–363.

GRIGORYAN, A. 2006. Heat kernels on weighted manifolds and applications. *Contemporary Mathematics*, 398, 93–191.

GUSKOV, I., SWELDENS, W., AND SCHRÖDER, P. 1999. Multiresolution signal processing for meshes. *ACM Siggraph*, 325–334.

HAMMOND, D. K., VANDERGHEYNST, P., AND GRIBONVAL, R. 2011. Wavelets on graphs via spectral graph theory. *Applied and Computational Harmonic Analysis 30*, 2, 129 – 150.

HAMZA, A. B., AND KRIM, H. 2003. Geodesic object representation and recognition. In *Intern. Conf. on Discrete Geometry for Computer Imagery*, 378–387.

HANSEN, P. C., AND O'LEARY, D. P. 1993. The use of the l-curve in the regularization of discrete ill-posed problems. *SIAM Journal of Scientific Computing 14*, 6, 1487–1503.

HEIN, M., AUDIBERT, J.-Y., AND VON LUXBURG, U. 2005. From graphs to manifolds – weak and strong pointwise consistency of graph Laplacians. In *Learning Theory*, vol. 3559 of *Lecture Notes in Computer Science*. Springer, 470–485.

HERHOLZ, P., KYPRIANIDIS, J. E., AND ALEXA, M. 2015. Perfect Laplacians for polygon meshes. *Computer Graphics Forum 34*, 5, 211–218.

HILDEBRANDT, K., POLTHIER, K., AND WARDETZKY, M. 2006. On the convergence of metric and geometric properties of polyhedral surfaces. *Geometria Dedicata*, 89–112.

HOU, T., AND QIN, H. 2012. Continuous and discrete Mexican hat wavelet transforms on manifolds. *Graphical Models 74*, 4, 221–232.

JACOBSON, A., BARAN, I., POPOVIC, J., AND SORKINE-HORNUNG, O. 2014. Bounded biharmonic weights for real-time deformation. *Commun. ACM 57*, 4, 99–106.

JAIN, V., AND ZHANG, H. 2007. A spectral approach to shape-based retrieval of articulated 3D models. *Computer Aided Design 39*, 398–407.

JAIN, V., ZHANG, H., AND VAN KAICK, O. 2007. Non-rigid spectral correspondence of triangle meshes. *Intern. Journal on Shape Modeling 13*, 1, 101– 124.

JOSHI, P., MEYER, M., DEROSE, T., GREEN, B., AND SANOCKI, T. 2007. Harmonic coordinates for character articulation. *ACM Trans. on Graphics 26*, 3 (July).

KARNI, Z., AND GOTSMAN, C. 2000. Spectral compression of mesh geometry. In *ACM Siggraph 2000*, 279–286.

KAZHDAN, M., SOLOMON, J., AND BEN-CHEN, M. 2012. Can mean-curvature flow be modified to be non-singular? *Computer Graphics Forum 31*, 5, 1745–1754.

KIM, B., AND ROSSIGNAC, J. 2005. Geofilter: Geometric selection of mesh filter parameters. *Computer Graphics Forum 24*, 3, 295–302.

KIM, S.-G., CHUNG, M. K., SEO, S., SCHAEFER, S. M., VAN REEKUM, C. M., AND DAVIDSON, R. J. 2012. Heat kernel smoothing via Laplace-Beltrami eigenfunctions and its application to subcortical structure modeling. In *Proc. of Pacific Conf. on Advances in Image and Video Technology*, 36–47.

KIM, K., TOMPKIN, J., AND THEOBALT, C. 2013. Curvature-aware regularization on riemannian submanifolds. In *IEEE Intern. Conf. on*, 881–888.

KOBBELT, L., CAMPAGNA, S., VORSATZ, J., AND SEIDEL, H.-P. 1998. Interactive multi-resolution modeling on arbitrary meshes. In *ACM Siggraph*, 105–114.

KOREN, Y. 2003. On spectral graph drawing. In *Lecture Notes in Computer Science*, Springer Verlag, vol. 2697, 496–508.

KOVNATSKY, A., BRONSTEIN, M. M., BRONSTEIN, A. M., GLASHOFF, K., AND KIMMEL, R. 2013. Coupled quasi-harmonic bases. *Computer Graphics Forum 32*, 2, 439–448.

KRISHNAN, D., FATTAL, R., AND SZELISKI, R. 2013. Efficient preconditioning of Laplacian matrices for computer graphics. *ACM Trans. on Graphics 32*, 4 (July), 142:1–142:15.

LAFON, S., KELLER, Y., AND COIFMAN, R. R. 2006. Data fusion and multicue data matching by diffusion maps. *IEEE Trans. on Pattern Analysis Machine Intelligence 28*, 11, 1784–1797.

LEE, W.-J., AND DUIN, R. P. 2008. An inexact graph comparison approach in joint eigenspace. In *Proc of the Intern. Workshop on Structural, Syntactic, and Statistical Pattern Recognition*, Springer-Verlag, 35–44.

LEHOUCQ, R., AND SORENSEN, D. C. 1996. Deflation techniques for an implicitly re-started Arnoldi iteration. *SIAM Journal of Matrix Analysis and Applications 17*, 789–821.

LEVY, B. 2006. Laplace-Beltrami eigenfunctions: towards an algorithm that understands geometry. In *Proc. of Shape Modeling and Applications*, 13.

LI, X., AND GUSKOV, I. 2005. Multi-scale features for approximate alignment of point-based surfaces. In *Proc. of the Symp. on Geometry Processing*, 217–226.

LI, X., GUO, X., WANG, H., HE, Y., GU, D., AND QIN, H. 2007. Harmonic volumetric mapping for solid modeling applications. In *Proc. of ACM Symp. on Solid and Physical Modeling*, ACM, 109–120.

LI, X., GU, X., AND QIN, H. 2009. Surface mapping using consistent pants decomposition. *IEEE Trans. on Visualization and Computer Graphics 15*, 4.

LI, X., XU, H., WAN, S., YIN, Z., AND YU, W. 2010. Feature-aligned harmonic volumetric mapping using MFS. *Computers & Graphics 34*, 3, 242–251.

LIAO, S., TONG, R., DONG, J., AND ZHU, F. 2009. Gradient field based inhomogeneous volumetric mesh deformation for maxillofacial surgery simulation. *Computers & Graphics 33*, 3, 424 – 432.

LIPMAN, Y., RUSTAMOV, R. M., AND FUNKHOUSER, T. A. 2010. Biharmonic distance. *ACM Trans. on Graphics 29*, 3.

LITMAN, R., BRONSTEIN, A., AND BRONSTEIN, M. 2011. Diffusion-geometric maximally stable component detection in deformable shapes. *Computers & Graphics 35*, 3, 549–560.

LITMAN, R., BRONSTEIN, A., AND BRONSTEIN, M. 2012. Stable volumetric features in deformable shapes. *Computers & Graphics 36*, 5, 569 – 576.

LIU, R., AND ZHANG, H. 2007. Mesh segmentation via spectral embedding and contour analysis. *Eurographics Tutorial 26*, 385–394.

LIU, Y., PRABHAKARAN, B., AND GUO, X. 2012. Point-based manifold harmonics. *IEEE Trans. on Visualization and Computer Graphics 18*, 10, 1693 –1703.

LIU, Y.-J., XU, C., AND HE, D. Y. 2015. Constructing intrinsic Delaunay triangulations from the dual of geodesic Voronoi diagrams. *ArXiv:1505.05590*.

LIU, Y., XU, C., FAN, D., AND HE, Y. 2015. Efficient construction and simplification of Delaunay meshes. *ACM Trans. on Graphics 34*, 6 (Oct.), 174:1–174:13.

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

MAHMOUDI, M., AND SAPIRO, G. 2009. Three-dimensional point cloud recognition via distributions of geometric distances. *Graphical Models 71*, 1, 22–31.

MALLET, J.-L. 1989. Discrete smooth interpolation. *ACM Trans. on Graphics 8*, 2, 121–144.

MARINI, S., PATANÉ, G., SPAGNUOLO, M., AND FALCIDIENO, B. 2011. Spectral feature selection for shape characterization and classification. *The Visual Computer 27*, 11, 1005–1019.

MARTIN, T., AND COHEN, E. 2010. Volumetric parameterization of complex objects by respecting multiple materials. *Computers & Graphics 34*, 3, 187–197.

MARTIN, T., COHEN, E., AND KIRBY, M. 2008. Volumetric parameterization and trivariate B-spline fitting using harmonic functions. In *Proc. of the ACM Symp. on Solid and Physical Modeling*, 269–280.

MÈMOLI, F., AND SAPIRO, G. 2005. A theoretical and computational framework for isometry invariant recognition of point cloud data. *Foundations of Computational Mathematics 5*, 3, 313–347.

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

MEMOLI, F. 2011. A spectral notion of gromov-wasserstein distance and related methods. *Applied and Computational Harmonic Analysis 30*, 3, 363 – 401.

MITRA, N. J., AND NGUYEN, A. 2003. Estimating surface normals in noisy point cloud data. In *Proc. of Symposium on Computational Geometry*, ACM Press, 322–328.

MOHAR, B., AND POLJAK, S. 1993. Eigenvalues in combinatorial optimization. *Combinatorial and graph-theoretical problems in linear algebra 23*, 98, 107–151.

MOLER, C., AND VAN LOAN, C. 2003. Nineteen dubious ways to compute the exponential of a matrix, twenty-five years later. *SIAM Review 45*, 1, 3–49.

NADIRASHVILI, N. S. 1988. Multiple eigenvalue of the Laplace-Beltrami operator. *Mathematics of the USSR-Sbornik 61*, 1, 225.

NEALEN, A., IGARASHI, T., SORKINE, O., AND ALEXA, M. 2006. Laplacian mesh optimization. In *Proc. of Computer graphics and interactive techniques*, 381–389.

NEUMANN, T., VARANASI, K., THEOBALT, C., MAGNOR, M. A., AND WACKER, M. 2014. Compressed manifold modes for mesh processing. *Computer Graphics Forum 33*, 5, 35–44.

NG, A. Y., JORDAN, M. I., AND WEISS, Y. 2001. On spectral clustering: analysis and an algorithm. In *Advances in Neural Information Processing Systems 14*, MIT Press, 849–856.

NI, X., GARLAND, M., AND HART, J. C. 2004. Fair morse functions for extracting the topological structure of a surface mesh. In *ACM Siggraph 2004*, 613–622.

OHBUCHI, R., TAKAHASHI, S., MIYAZAWA, T., AND MUKAIYAMA, A. 2001. Watermarking 3D polygonal meshes in the mesh spectral domain. In *Graphics Interface 2001*, 9–17.

OHBUCHI, R., MUKAIYAMA, A., AND TAKAHASHI, S. 2002. A frequency-domain approach to watermarking 3D shapes. *Computer Graphics Forum 21*, 3.

ORECCHIA, L., SACHDEVA, S., AND VISHNOI, N. K. 2012. Approximating the exponential, the Lanczos method and an ˜o(m)-time spectral algorithm for balanced separator. In *Proc. of the 44th Symposium on Theory of Computing Conference*, 1141–1160.

OSADA, R., FUNKHOUSER, T., CHAZELLE, B., AND DOBKIN, D. 2002. Shape distributions. *ACM Trans. on Graphics 21*, 4, 807–832.

OVSJANIKOV, M., MÈRIGOT, Q., MÈMOLI, F., AND GUIBAS, L. 2010. One point isometric matching with the heat kernel. *ACM Symp. on Discrete Algorithms*, 650–663.

OVSJANIKOV, M., BEN-CHEN, M., SOLOMON, J., BUTSCHER, A., AND GUIBAS, L. J. 2012. Functional maps: a flexible representation of maps between shapes. *ACM Trans. on Graphics 31*, 4, 30.

PATANÈ, G., AND FALCIDIENO, B. 2009. Computing smooth approximations of scalar functions with constraints. *Computer & Graphics 33*, 3, 399 – 413.

PATANÈ, G., AND FALCIDIENO, B. 2010. Multi-scale feature spaces for shape processing and analysis. In *Proc. of Shape Modeling Intern.*, 113–123.

PATANÉ, G., AND SPAGNUOLO, M. 2012. Local approximation of scalar functions on 3D shapes and volumetric data. *Computers & Graphics 36*, 5, 387 – 397.

PATANÈ, G., AND SPAGNUOLO, M. 2013. Heat diffusion kernel and distance on surface meshes and point sets. *Computers & Graphics 37*, 6, 676 – 686.

PATANÈ, G., AND SPAGNUOLO, M. 2013. An interactive analysis of harmonic and diffusion equations on discrete 3D shapes. *Computer & Graphics*.

PATANÈ, G., SPAGNUOLO, M., AND FALCIDIENO, B. 2007. Families of cut-graphs for bordered meshes with arbitrary genus. *Graphical Models 69(2)*, 119–138.

PATANÈ, G., SPAGNUOLO, M., AND FALCIDIENO, B. 2009. Topology- and error-driven extension of scalar functions from surfaces to volumes. *ACM Trans. on Graphics 29*, 1, 1–20.

PATANÈ, G. 2013. wFEM heat kernel: discretization and applications to shape analysis and retrieval. *Computer Aided Geometric Design 30*, 3, 276–295.

PATANÈ, G. 2014. Laplacian spectral distances and kernels on 3D shapes. *Pattern Recognition Letters 47*, 102–110.

PATANÈ, G. 2015. Diffusive smoothing of 3D segmented medical data. *Journal of Advanced Research 6*, 3, 425 – 431.

PATANÉ, G. 2016. Accurate and efficient computation of Laplacian spectral distances and kernels. *Computer Graphics Forum*, In press.

PERONA, P., AND MALIK, J. 1990. Scale-space and edge detection using anisotropic diffusion. *IEEE Trans. on Pattern Analysis and Machine Intelligence 12*, 7, 629–639.

PINKALL, U., AND POLTHIER, K. 1993. Computing discrete minimal surfaces and their conjugates. *Experimental Mathematics 2*, 1, 15–36.

PUSA, M. 2011. Rational approximations to the matrix exponential in burnup calculations. *Nuclear Science and Engineering 169*, 2, 155–167.

RAMANI, K., AND SINHA, A. 2013. Multiscale kernels using random walks. *Computer Graphics Forum 33*, 1, 164–177.

RAVIV, D., BRONSTEIN, M. M., BRONSTEIN, A. M., AND KIMMEL, R. 2010. Volumetric heat kernel signatures. In *Proc. of the ACM Workshop on 3D Object Retrieval*, 3DOR '10, 39–44.

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

REUTER, M., BIASOTTI, S., GIORGI, D., PATANÈ, G., AND SPAGNUOLO, M. 2009. Discrete Laplace-Beltrami operators for shape analysis and segmentation. *Computer & Graphics 33*, 3, 381–390.

REUTER, M., WOLTER, F.-E., SHENTON, M. E., AND NIETHAMMER, M. 2009. Laplace-Beltrami eigenvalues and topological features of eigenfunctions for statistical shape analysis. *Computer-Aided Design 41*, 10, 739–755.

ROSENBERG, S. 1997. *The Laplacian on a Riemannian Manifold*. Cambridge University Press.

ROWEIS, S. T., AND SAUL, L. K. 2000. Nonlinear dimensionality reduction by locally linear embedding. *Science 290*, 2323–2326.

RUDIN, W. 1987. *Real and complex analysis*, second ed. Intern. Series in Pure and Applied Mathematics. McGraw-Hill Inc., New York.

RUGGERI, M., PATANÈ, G., SPAGNUOLO, M., AND SAUPE, D. 2010. Spectral-driven isometry-invariant matching of 3D shapes. *Intern. Journal of Computer Vision 89*, 2-3, 248–265.

RUSTAMOV, R. M. 2007. Laplace-Beltrami eigenfunctions for deformation invariant shape representation. In *Proc. of the Symp. on Geometry Processing*, 225–233.

RUSTAMOV, R. M. 2011. Interpolated eigenfunctions for volumetric shape processing. *The Visual Computer 27*, 11, 951–961.

RUSTAMOV, R. M. 2011. Multiscale biharmonic kernels. *Computer Graphics Forum 30*, 5, 1521–1531.

SAAD, Y. 1992. Analysis of some krylov subspace approximations to the matrix exponential operator. *SIAM Jorunal of Numererical Analysis 29*, 209–228.

SCHOELKOPF, B., AND SMOLA, A. J. 2002. *Learning with Kernels*. The MIT Press.

SHEWCHUK, J. R. 2002. Delaunay refinement algorithms for triangular mesh generation. *Computational Geometry Theory and Applications 22*, 1-3, 21–74.

SHI, J., AND MALIK, J. 2000. Normalized cuts and image segmentation. *IEEE Trans. on Pattern Analysis and Machine Intelligence 22*, 8, 888–905.

SIDJE, R. B. 1998. Expokit: A software package for computing matrix exponentials. *ACM Trans. on Mathematical Software 24*, 1 (Mar.), 130–156.

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

SINKHORN, R. 1964. A relationship between arbitrary positive matrices and doubly stochastic matrices. *Ann. Math. Statist. 35*, 2 (06), 876–879.

SMOLA, A. J., AND KONDOR, R. I. 2003. Kernels and regularization on graphs. In *Conf. on Learning Theory*, 144–158.

SOLOMON, J., DE GOES, F., PEYRÉ, G., CUTURI, M., BUTSCHER, A., NGUYEN, A., DU, T., AND GUIBAS, L. 2015. Convolutional wasserstein distances: Efficient optimal transportation on geometric domains. *ACM Trans. on Graphics 34*, 4 (July), 66:1–66:11.

SORENSEN, D. C. 1992. Implicit application of polynomial filters in a k-step arnoldi method. *SIAM Journal of Matrix Analysis and Applications 13*, 1, 357–385.

SORKINE, O., COHEN-OR, D., AND TOLEDO, S. 2003. High-pass quantization for mesh encoding. In *Proc. of the Symp. on Geometry Processing*, 42–51.

SORKINE, O., LIPMAN, Y., COHEN-OR, D., ALEXA, M., ROESSL, C., AND SEIDEL, H.-P. 2004. Laplacian surface editing. In *Proc. of the Symp. on Geometry Processing*, 179–188.

SORKINE, O., COHEN-OR, D., IRONY, D., AND TOLEDO, S. 2005. Geometry-aware bases for shape approximation. *IEEE Trans. on Visualization and Computer Graphics 11*, 2, 171–180.

SORKINE, O. 2006. Differential representations for mesh processing. *Computer Graphics Forum 25*, 4, 789–807.

SPIELMAN, D. A., AND TENG, S.-H. 2007. Spectral partitioning works: planar graphs and finite element meshes. *Linear Algebra and its Applications 421*, 284–305.

SPIRA, A., KIMMEL, R., AND SOCHEN, N. 2007. A short-time Beltrami kernel for smoothing images and manifolds. *Trans. on Image Processing 16*, 6, 1628–1636.

SUN, J., OVSJANIKOV, M., AND GUIBAS, L. J. 2009. A concise and provably informative multi-scale signature based on heat diffusion. *Computer Graphics Forum 28*, 5, 1383–1392.

TASDIZEN, T., WHITAKER, R., BURCHARD, P., AND OSHER, S. 2002. Geometric surface smoothing via anisotropic diffusion of normals. In *Proc. of the Conf. on Visualization*, 125–132.

TAUBIN, G. 1995. A signal processing approach to fair surface design. In *ACM Siggraph 1995*, 351–358.

TAUBIN, G. 1999. 3D geometry compressionand progressive transmission. In *Eurographics Tutorials*.

TENENBAUM, J. B., SILVA, V., AND LANGFORD, J. C. 2000. A Global Geometric Framework for Nonlinear Dimensionality Reduction. *Science 290*, 5500, 2319–2323.

TONG, Y., LOMBEYDA, S., HIRANI, A. N., AND DESBRUN, M. 2003. Discrete multiscale vector field decomposition. *ACM Trans. on Graphics 22*, 3, 445–452.

UMEYAMA, S. 1988. An eigendecomposition approach to weighted graph matching problems. *IEEE Trans. on Pattern Analysis and Machine Intelligence 10*, 5, 695–703.

VALLET, B., AND LÈVY, B. 2008. Spectral geometry processing with manifold harmonics. *Computer Graphics Forum 27*, 2, 251–260.

VAN LOAN, C. 1979. A note on the evaluation of matrix polynomials. *Automatic Control, IEEE Trans. on 24*, 2 (Apr), 320–321.

VARADHAN, S. R. S. 1967. On the behavior of the fundamental solution of the heat equation with variable coefficients. *Communications on Pure and Applied Mathematics 20*, 2, 431–455.

VAXMAN, A., BEN-CHEN, M., AND GOTSMAN, C. 2010. A multi-resolution approach to heat kernels on discrete surfaces. *ACM Trans. on Graphics 29*, 4, 1–10.

WANG, G., ZHANG, X., SU, Q., CHEN, J., WANG, L., MA, Y., LIU, Q., XU, L., SHI, J., AND WANG, Y. 2013. A heat kernel based cortical thickness estimation algorithm. In *MBIA*, vol. 8159 of *Lecture Notes in Computer Science*, 233–245.

WARDETZKY, M., MATHUR, S., KÄLBERER, F., AND GRINSPUN, E. 2007. Discrete Laplace operators: no free lunch. In *Proc. of the Symp. on Geometry Processing*, 33–37.

WEBER, O., PORANNE, R., AND GOTSMAN, C. 2012. Biharmonic coordinates. *Computer Graphics Forum 31*, 8 (Dec.), 2409–2422.

WITKIN, A. P. 1983. Scale-space filtering. In *Proc. of the Intern. Joint Conf. on Artificial Intelligence*, 1019–1022.

XIAO, B., HANCOCK, E. R., AND WILSONB, R. 2010. Geometric characterization and clustering of graphs using heat kernel embeddings. *Image and Vision Computing 28*, 6, 1003 – 1021.

XU, G. 2007. Discrete Laplace-Beltrami operators and their convergence. *Computer Aided Geometric Design 8*, 21, 398–407.

ZENG, W., GUO, R., L., F., AND GU, X. 2012. Discrete heat kernel determines discrete riemannian metric. *Graphical Models 74*, 4, 121–129.

ZHANG, H., AND FIUME, E. 2003. Butterworth filtering and implicit fairing of irregular meshes. In *Proc. of the Pacific Conf. on Computer Graphics and Applications*, 502.

ZHANG, F., AND HANCOCK, E. R. 2008. Graph spectral image smoothing using the heat kernel. *Pattern Recognition 41*, 11, 3328 – 3342.

ZHANG, H., AND LIU, R. 2005. Mesh segmentation via recursive and visually salient spectral cuts. In *Proc. of Vision, Modeling, and Visualization*, 429–436.

ZHANG, H., VAN KAICK, O., AND DYER, R. 2007. Spectral methods for mesh processing and analysis. In *Eurographics State-of-the-art Report*, 1–22.

ZHENG, Y., TAI, C.-L., ZHANG, E., AND XU, P. 2013. Pairwise harmonics for shape analysis. *IEEE Trans. on Visualization and Computer Graphics 19*, 7 (July), 1172–1184.

ZHOU, K., SYNDER, J., GUO, B., AND SHUM, H.-Y. 2004. Isocharts: stretch-driven mesh parameterization using spectral analysis. In *Proc. of the Symp. on Geometry processing*, 45–54.

ZHU, X., GHAHRAMANI, Z., AND LAFFERTY, J. 2003. Semi-supervised learning using gaussian fields and harmonic functions. In *Intern. Conf. on Machine Learning*, 912–919.

ZIGELMAN, G., KIMMEL, R., AND KIRYATI, N. 2002. Texture mapping using surface flattening via multidimensional scaling. *IEEE Trans. on Visualization and Computer Graphics 8*, 2, 198–207.