Computers & Graphics 33 (2009) 381–390

Contents lists available at ScienceDirect

Computers & Graphics

journal homepage: www.elsevier.com/locate/cag

Technical Section

Discrete Laplace–Beltrami operators for shape analysis and segmentation

Martin Reuter [a,b], Silvia Biasotti [c,*], Daniela Giorgi [c], Giuseppe Patanè [c], Michela Spagnuolo [c]

[a] Massachusetts Institute of Technology, Cambridge, MA, USA
[b] A.A. Martinos Center for Biomedical Imaging, Massachusetts General Hospital, Harvard Medical School, Boston, MA, USA
[c] Istituto di Matematica Applicata e Tecnologie Informatiche - Consiglio Nazionale delle Ricerche, Genova, Italy

A R T I C L E I N F O

Article history:
Received 12 December 2008
Received in revised form
2 March 2009
Accepted 3 March 2009

Keywords:
Laplace–Beltrami operator
Eigenfunctions
Nodal sets
Nodal domains
Shape analysis
Shape segmentation

A B S T R A C T

Shape analysis plays a pivotal role in a large number of applications, ranging from traditional geometry processing to more recent 3D content management. In this scenario, spectral methods are extremely promising as they provide a natural library of tools for shape analysis, intrinsically defined by the shape itself. In particular, the eigenfunctions of the Laplace–Beltrami operator yield a set of real-valued functions that provide interesting insights in the structure and morphology of the shape. In this paper, we first analyze different discretizations of the Laplace–Beltrami operator (geometric Laplacians, linear and cubic FEM operators) in terms of the correctness of their eigenfunctions with respect to the continuous case. We then present the family of segmentations induced by the nodal sets of the eigenfunctions, discussing its meaningfulness for shape understanding.

1. Introduction

Shape analysis aims to develop computational tools for reasoning on properties of the objects' shape, and is pivotal in a large number of applications, ranging from traditional geometry processing to more recent 3D content management techniques.

In the recent past, research in shape analysis was boosted by the need to add *semantics* to the geometric description of 3D objects, in order to facilitate the sharing and management of 3D content in many emerging web-based applications. A semantic description of 3D objects is commonly understood as a description of the content by means of terms which are meaningful in some domain of knowledge. For example, a given model can be described as being a *table*, made of four *cylindrical legs* and an *oval top*. Hence, a semantic description calls for segmentation algorithms which capture semantically relevant features in an automatic manner.

Most of the methods developed so far for shape analysis and segmentation do not directly provide any semantically relevant explicit description of the shape, but rather provide a characterization of the geometric and structural properties of the object boundary. Semantic properties are taken into account, at some extent, by cognitive theories supporting part-based decompositions or minima rule-based approaches. *Part-based decomposition* techniques build on Biederman's theory of perception, which characterizes an object as a compound of primitive basic parts (e.g., planes, spheres, cylinders, cubes) [10]. The second class of methods are based on the so-called *minima rule*, which suggests that we perceive relevant parts by focusing our attention on lines of concave discontinuity of the tangent plane [26]. For a recent survey of segmentation methods, we refer the reader to [54].

Our interest is in the development of methods for shape analysis and segmentation able to capture a varied set of morphologically relevant features, possibly at different scales: in other words, we seek for a *library* of tools supporting the semantic annotation of digital shapes. Shape understanding, indeed, is a very complex task and it is now widely accepted that no single segmentation method is capable of capturing relevant features in a broad domain of shapes. In [3], shape understanding is seen as a *multi-segmentation* task driven by the user, who uses in parallel a set of segmentation algorithms and composes the final segmentation with selection and refinement operations on the segments. In that work, the authors push forward the idea of semantic annotation by allowing the user to associate textual tags, defined in an ontology, to the segments.

In this scenario, spectral methods are extremely promising, as they naturally provide a set of tools for shape analysis that are intrinsically defined by the shape itself. Spectral methods have recently gained much interest in computer graphics [62], with applications that include mesh compression [30], parametrization [25,43], segmentation[31,35], remeshing [17], filtering [34,57],

* Corresponding author. Fax: +39 010 6475 696.
E-mail addresses: reuter@mit.edu (M. Reuter), silvia@ge.imati.cnr.it (S. Biasotti), daniela@ge.imati.cnr.it (D. Giorgi), patane@ge.imati.cnr.it (G. Patanè), michi@ge.imati.cnr.it (M. Spagnuolo).

0097-8493/$ - see front matter © 2009 Elsevier Ltd. All rights reserved.
doi:10.1016/j.cag.2009.03.005

correspondence [27], matching and retrieval [28,49,52], manifold learning [5], and imaging or medical imaging applications [44,38,50].

In particular, the eigenfunctions of the Laplace–Beltrami operator yield a family of real-valued functions that provide interesting insights in the structure and morphology of shapes. In this paper, we focus on the *nodal sets* of the Laplace–Beltrami eigenfunctions, showing that they induce a shape decomposition which captures features at different scales, generally well aligned with perceptually relevant shape features. The set of decompositions induced by the eigenfunctions yields the sought *library* of intrinsic shape segmentations.

The first contribution of this paper is the analysis of the correctness of the eigenfunctions computed using different discretizations of the Laplace–Beltrami operator (Section 2), evaluated with respect to the exact results known from the theory in the continuous case (Section 3). The second contribution is the introduction of the set of segmentations induced by the nodal sets of the eigenfunctions; the segmentations are discussed in terms of their quality and robustness (Section 4). Finally, we draw some conclusive remarks and highlight possible extensions of this work (Section 5).

2. The Laplace–Beltrami operator

Let f be a $\mathscr{C}^2$ real-valued function defined on a differentiable manifold $\mathscr{M}$ with Riemannian metric [7]. The *Laplace–Beltrami operator* Δ is

$$\Delta f := \mathrm{div}(\mathrm{grad} f),$$

where grad and div are the gradient and divergence on the manifold $\mathscr{M}$ [11]. The *Laplacian eigenvalue problem* is stated as

$$\Delta f = -\lambda f. \tag{1}$$

Since the Laplace–Beltrami operator is self-adjoint and semi-positive definite [51], it admits an *orthonormal eigensystem* $\mathscr{B} := \{(\lambda_i, \psi_i)\}_i$, that is a basis of the space of square integrable function, with $\Delta\psi_i = \lambda_i\psi_i$, $\lambda_0 \leqslant \lambda_1 \leqslant \cdots, \lambda_i \leqslant \lambda_{i+1} \cdots \leqslant +\infty$. For a detailed discussion on the main properties of the Laplace–Beltrami operator, we refer the reader to [46,51,59].

2.1. The discrete case

The solution to (1) on a surface is frequently approximated by a piecewise linear function $f : \mathscr{T} \to \mathbb{R}$ over a triangulation $\mathscr{T}$ with vertices $V := \{\mathbf{p}_i, i = 1, \ldots, n\}$. The function f on $\mathscr{T}$ is defined by linearly interpolating the values $f(\mathbf{p}_i)$ of f at the vertices of $\mathscr{T}$. This is done by choosing a base of piecewise-linear *hat-functions* φ_i, with value 1 at vertex $\mathbf{p}_i$ and 0 at all the other vertices. Then f is given as $f = \sum_{i=1}^{n} f(\mathbf{p}_i)\varphi_i$. Discrete Laplace–Beltrami operators are usually represented as

$$\Delta f(\mathbf{p}_i) := \frac{1}{d_i} \sum_{j \in N(i)} w_{ij}[f(\mathbf{p}_i) - f(\mathbf{p}_j)], \tag{2}$$

where $N(i)$ denotes the index set of the 1-*ring* of the vertex $\mathbf{p}_i$, i.e., the indices of all neighbors connected to $\mathbf{p}_i$ by an edge. The masses d_i are associated to a vertex i and the w_{ij} are the symmetric edge weights. To write (2) in matrix form, we define the vector $\mathbf{f} := (f(\mathbf{p}_1), \ldots, f(\mathbf{p}_n))^T$ of the function values at the vertices, the *weighted adjacency matrix* $W := (w_{ij})$, and the diagonal matrix $V := \mathrm{diag}(v_1, \ldots, v_n)$ containing as diagonal elements $v_i = \sum_{j \in N(i)} w_{ij}$. Then, we can define a *stiffness matrix* $A := V - W$, the *lumped mass matrix* $D := \mathrm{diag}(d_1, \ldots, d_n)$, and finally the *Laplace matrix* $L := D^{-1}A$ (generally not symmetric). Using these matrices, $\Delta f(\mathbf{p}_i)$ is the i-th component of the vector $L\mathbf{f}$. The problem (1) can then be written as $L\mathbf{f} = \lambda\mathbf{f}$ or better as a generalized symmetric problem $A\mathbf{f} = \lambda D\mathbf{f}$. In the following, we distinguish between *geometric operators* and *finite-element operators* on the basis of different edge weights and masses.

2.1.1. Discrete geometric Laplacians

A very simple choice of weights w_{ij} for a graph is the adjacency matrix (1 if $\mathbf{p}_i$ and $\mathbf{p}_j$ are connected by an edge, 0 otherwise) and unit masses $d_i = 1$. This operator and simple variations are called graph Laplacians as they usually only consider the connectivity and no geometry. Lévy [33] gives a very good overview and compares this graph Laplacian with a discretization by Desbrun et al. [16] (presented below).

One of the early geometric approaches has been described by Pinkall and Polthier [45], who discretize the Laplace–Beltrami operator using constant masses (i.e., $d_i := 1$) in (2) and weights

$$w_{ij} := \frac{\cot(\alpha_{ij}) + \cot(\beta_{ij})}{2}, \tag{3}$$

where α_{ij} and β_{ij} denote the two angles opposite to the edge (i,j). Because of the lack of a proper mass weighting the cotangent weights alone still depend on mesh sampling.

Desbrun et al. [16] refine the discretization in (3) by using a normalization factor, which takes into account the area $a(i)$ of all triangles at vertex i, i.e.,

$$d_i := a(i)/3. \tag{4}$$

Lévy [33] uses this operator but instead of solving the symmetric generalized problem $A\mathbf{f} = \lambda D\mathbf{f}$ he looks at the non-symmetric matrix $L = D^{-1}A$ and then computes the eigenvalues and eigenfunctions of the symmetric matrix $(L + L^t)/2$ which yields a different spectrum.

Meyer et al. [36] modify the area normalization by Desbrun and propose the mass weighting

$$d_i := a_V(i), \tag{5}$$

with $a_V(i)$ the area obtained by joining the circumcenters of the triangles around vertex i (i.e., the Voronoi region). Founding on discrete exterior calculus, [15,34] reach the same operator. Lévy and Vallet [34] symmetrize the operator by using $1/\sqrt{a_V(i)a_V(j)}$ instead of the inversion of the mass matrix $1/a_V(i)$. This leads to the system $D^{-1/2}AD^{-1/2}\mathbf{y} = \lambda\mathbf{y}$ with the same eigenvalues. The original eigenvectors can be retrieved by $\mathbf{f} = D^{-1/2}\mathbf{y}$.

Belkin et al. [5,6] describe a discretization of the Laplace–Beltrami operator on the k-nearest neighbor graph $\mathscr{T}$ of a point set $\{\mathbf{p}_i\}_{i=1}^{n}$ sampled on an underlying manifold and an extension to meshes by using the heat kernel to construct the weights. The mesh version [6] considers weights not only at the edges of the mesh, but in a larger neighborhood of a vertex (the heat kernel is cut-off thus sparsity is maintained). While the geometric operators in (3)–(5) are not convergent in general and cannot deal well with non-uniform meshes [60], this method exhibits convergence and does not depend much on the shape of the triangles, just on the density of the vertices. However, it can be used to compute eigenfunctions only on closed meshes, as it is unclear how to comply with the Dirichlet or Neumann boundary condition. Another discretization by Floater can be found in [22], but is not a good choice for eigencomputations due to its non-symmetry.

2.1.2. Discrete FEM Laplacians

The solution of the Laplace eigenvalue problem (1) can be computed by imposing that the equation $\Delta f = -\lambda f$ is verified in a *weak sense*, that is,

$$\langle \Delta f, \varphi_i \rangle_{\mathscr{L}^2(\mathscr{M})} = -\lambda \langle f, \varphi_i \rangle_{\mathscr{L}^2(\mathscr{M})}, \quad \forall i. \tag{6}$$