Scale-Space Spectral Representation of Shape*

Jonathan Bates[1], Xiuwen Liu[2], Washington Mio[1]
[1]*Department of Mathematics,* [2]*Department of Computer Science*
Florida State University, Tallahassee, FL 32306, USA
jbates@math.fsu.edu, liux@cs.fsu.edu, mio@math.fsu.edu

Abstract

We construct a scale space of shape of closed Riemannian manifolds, equipped with metrics derived from spectral representations and the Hausdorff distance. The representation depends only on the intrinsic geometry of the manifolds, making it robust to pose and articulation. The computation of shape distance involves an optimization problem over the 2^p-element group of all p-bit strings, which is approached with Markov chain Monte Carlo techniques. The methods are applied to cluster surfaces in 3D space.

1. Introduction

We employ the heat kernel (HK) associated with the Laplace-Beltrami operator Δ on closed Riemannian manifolds to construct a scale-space representation of shape, with scale controlled by a "time" parameter t. We call this representation a heat-kernel, or HK, representation. HK representations are insensitive to distance-preserving transformations, a property that makes them robust to pose and articulation (cf. [4]). It is difficult to compare the shape of different Riemannian manifolds through the direct use of their HK representations because they are defined on different domains. To address this problem, we assume that the eigenfunctions of Δ (ordered by increasing magnitude of the eigenvalues) are in correspondence. Although this assumption is not always valid [3], *e.g.*, for shapes that exhibit symmetries, the shape distances defined are robust to these mismatches. For a closed Riemannian manifold M and positive integer p, the coordinates of the orthogonal projection of the heat kernel of M onto the orthonormal system $\phi_1, \ldots, \phi_p$, the eigenfunctions of Δ associated with the first p nontrivial eigenvalues, give a shape representation of M in p-dimensional Euclidean space $\mathbb{R}^p$ (Section 2). We use the Hausdorff distance between subsets of $\mathbb{R}^p$ to obtain a family of shape metrics indexed by t (Section 3). Generically, the eigenvalues of Δ have multiplicity 1, so that the orthonormal eigenfunctions are uniquely determined up to sign. This means that, for each t, there are 2^p different p-dimensional representations of a shape. The Hausdorff distance is minimized over all possible choices to eliminate dependence on a specific choice. This poses a non-trivial optimization problem, which we treat with a Markov chain Monte Carlo (MCMC) approach (Section 4). To illustrate the methods and the effectiveness of the HK representation, shape metric and optimization strategies, we carry out a clustering experiment with surfaces in 3D space (Section 5).

*This research was supported in part by NSF grant DMS-0713012 and NIH Roadmap for Medical Research grant U54 RR021813.

2. Scale-Space HK Representation

Let M be a closed Riemannian manifold, for example, a watertight surface in 3D space. The Laplace-Beltrami operator Δ_M on functions $f\colon M \to \mathbb{R}$ is the differential operator $\Delta_M f = \operatorname{div}(\nabla f)$, where div and ∇ are the Riemannian divergence and gradient, respectively. Equivalently, $\Delta f = *\,d\,*\,d$, where d denotes the exterior derivative and $*$ the Hodge star operator. Let $\mathbb{L}^2$ be the Hilbert space of square-integrable functions on M with the inner product $\langle f, g\rangle = \int_M f(x)g(x)\,dV(x)$, where dV is the volume element of M. $\mathbb{L}^2$ admits a complete orthonormal set $\phi_i\colon M \to \mathbb{R}$, $i \geqslant 0$, of eigenfunctions of $-\Delta_M$, with eigenvalues $\lambda_i \geqslant 0$ satisfying $\lim_{i\to\infty} \lambda_i = \infty$. Thus, $\Delta\phi_i = -\lambda_i\phi_i$. As usual, we order the eigenvalues so that $\lambda_i \leqslant \lambda_{i+1}$. Hereafter, we assume that M is connected, in which case $\lambda_0 = 0$ and $\lambda_i > 0$ for $i \geqslant 1$. The eigenfunction ϕ_0 is constant. For symmetric shapes, the non-zero eigenvalues may arise with multiplicity. However, the eigenvalues are generically all distinct. We assume that this is the case since our interest is in shapes extracted from data, and they exhibit no perfect intrinsic

1051-4651/10 $26.00 © 2010 IEEE
DOI 10.1109/ICPR.2010.649

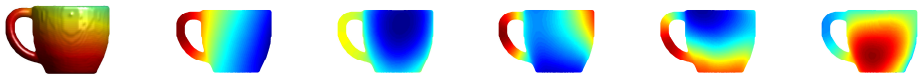

Figure 1. First 6 eigenfunctions of the Laplace-Beltrami operator of the contour of a cup.

self-similarities if dim $M > 1$. Figure 1 shows the first 6 eigenfunctions of the contour surface of a cup.

Let $K(x, y, t)$, $x, y \in M$ and $t > 0$, be the heat kernel on M. K has the property that if $f\colon M \to \mathbb{R}$ is an initial distribution of temperature on M, then the distribution for time $t > 0$ is given by

$$\phi(x,t) = \int_M K(x,y,t) f(y)\, dV(y)\,. \tag{1}$$

The heat kernel may be expressed as $K(x,y,t) = \sum_{i=0}^{\infty} e^{-\lambda_i t}\phi_i(x)\phi_i(y)$. Letting $K_{x,t}(y) = K(x,y,t)$, the assignment $x \mapsto K_{x,t}$ gives a scale-space representation of M in function space. For a fixed t, we see that the contributions to $K_{x,t}$ of the eigenmodes decay exponentially with the magnitude of λ_i. Because we expect the most salient shape dissimilarities to be detectable at the lower modes, we truncate $K_{x,t}$ to

$$K^p_{x,t}(y) = \sum_{i=0}^{p} e^{-\lambda_i t}\phi_i(x)\phi_i(y)\,. \tag{2}$$

The coordinates of $K^p_{x,t}$ with respect to the orthonormal set $\phi_1, \ldots, \phi_p$ yield a p-dimensional representation $\psi_t\colon M \to \mathbb{R}^p$, namely,

$$\psi_t(x) = \begin{bmatrix} e^{-\lambda_1 t}\phi_1(x) & \ldots & e^{-\lambda_p t}\phi_p(x) \end{bmatrix}^T\,. \tag{3}$$

We refer to ψ_t as a p-dimensional HK representation of M at time t. As we are assuming that all eigenvalues have multiplicty 1, the choice of ϕ_i is unique up to sign. Thus, there are 2^p distinct p-dimensional heat-kernel representations of a shape.

The shape representation (3) is sensitive to scale. However, in practice, both scale-invariant and scale-sensitive shape models are of interest. Among the many ways of fixing scale, we normalize the eigenvalues as follows. First note that if we scale a shape by a factor $r > 0$, the eigenvalues of the Laplacian divide by r^2. Likewise, for the eigenfunctions to remain orthonormal, they must be divided by $r^{d/2}$, where d is the dimension of the manifold M. If $r = \sqrt{\lambda_1}$, the eigenvalues become λ_i/λ_1 and the HK representation changes to

$$\widetilde{\psi}_t(x) = \frac{1}{{\lambda_1}^{d/4}} \begin{bmatrix} e^{-t}\phi_1(x) & \ldots & e^{-\lambda_p t/\lambda_1}\phi_p(x) \end{bmatrix}^T.$$

Figure 2 shows an example of a 3D HK representation. Point correspondences are color coded.

original shape HK representation

Figure 2. 3D heat-kernel representation.

3. Shape Metric

The sign indeterminacies in the HK representation may be formalized as follows. Let Z_2^p be the mod 2 additive group of all bit strings of length p. That is, binary words of length p, with the bitwise "xor" operation. We introduce an action of Z_2^p on mappings $f\colon M \to \mathbb{R}^p$. For $b = (b_1, \ldots, b_p) \in Z_2^p$, define $b \cdot f\colon M \to \mathbb{R}^p$ to be

$$(b \cdot f)(x) = \begin{bmatrix} (-1)^{b_1} f_1(x) & \ldots & (-1)^{b_p} f_p(x) \end{bmatrix}^T\,. \tag{4}$$

This action simply changes the sign of the ith coordinate of f if $b_i = 1$, and leaves it unchanged, otherwise. The orbit of ψ_t under this action, $O(\psi_t) = \{b \cdot \psi_t \,|\, b \in Z_2^p\}$, is precisely the collection of all p-dimensional HK representations of M for a given t.

Given HK representations φ_t and ψ_t of the Riemannian manifolds M and N, respectively, let $d_H(\varphi_t, \psi_t)$ be the Hausdorff distance between the sets $\varphi_t(M), \psi_t(N) \subset \mathbb{R}^p$, which is given by

$$\max\{\sup_{x\in M} d(\varphi_t(x), \psi_t(N)), \sup_{y\in N} d(\psi_t(y), \varphi_t(M))\}\,. \tag{5}$$

For each fixed $t > 0$, define a shape metric

$$\begin{aligned} D_H(M, N; t) &= \min_{b,c\in Z_2^p} d_H(c \cdot \varphi_t, b \cdot \psi_t) \\ &= \min_{b\in Z_2^p} d_H(\varphi_t, b \cdot \psi_t)\,, \end{aligned} \tag{6}$$

which measures the distance between the orbits of φ_t and ψ_t. The second equality follows from the fact that the action of Z_2^p preserves the Hausdorff distance.

The main computational problem that arises in calculating d_H is the minimization over Z_2^p, which we approach with a Markov chain Monte Carlo (MCMC) algorithm, as explained in Section 4.

In practice, a surface M in $\mathbb{R}^3$ is discretized as a triangle mesh K. We use the discrete Laplace-Beltrami operator on K as an approximation. The details of this discretization and the computation of eigenvalues and eigenfunctions are described in [1].

Functions on M are discretized over the vertex set of K. Let W be the $n \times p$ matrix whose jth column is the jth normalized eigenvector of $-\Delta_K$, where n is the number of vertices of K. Then, the value of the discrete version of (3) on the ith vertex may be expressed as $\left[e^{-\lambda_1 t} w_{i1} \quad \ldots \quad e^{-\lambda_p t} w_{ip}\right]^T$. These vectors are the columns of the $p \times n$ matrix $P(t) = e^{-\Lambda t} W^T$, where Λ is the diagonal matrix having λ_i as the ith diagonal entry. $P(t)$ is the discrete version of (3) in matrix form. The columns of $P(t)$ may be viewed as the point cloud $\{p_1(t), \ldots, p_n(t)\} \subset \mathbb{R}^p$. We abuse notation and denote both the matrix and the point cloud as $P(t)$. It is simple to check that the action of a bit string $b = (b_1, \ldots, b_p) \in Z_2^p$ on $P(t)$ can be expressed as

$$b \cdot P(t) = (-1)^B P(t) = (-1)^B e^{-\Lambda t} W^T , \quad (7)$$

where $(-1)^B$ is the diagonal matrix with ith diagonal entry $(-1)^{b_i}$. For $1 \leqslant i \leqslant p$, this action simply flips the signs of the entries on the ith row of $P(t)$ if $b_i = 1$. The scale-invariant case is similar.

4. Computation of Shape Distance

The main computational problem that arises in calculating d_H is the minimization over Z_2^p, which we approach with a Markov chain Monte Carlo (MCMC) algorithm (cf. [2]). We drop t from the notation because it is held fixed throughout the discussion. With this convention, if K and L are meshes, the problem is to calculate

$$D_H(K, L) = \min_{b \in Z_2^p} d_H(P, b \cdot Q) , \quad (8)$$

where P and Q are p-dimensional HK representations of K and L, respectively. Define the plausibility of a bit string b as $F(b) = \exp\left(-d_H^2(P, b \cdot Q)/2\sigma^2\right)$, where $\sigma > 0$ is a constant. The associated probability distribution on Z_2^p is $\pi(b) = F(b)/\sum_{c \in Z_2^p} F(c)$, but the normalizing constant typically is unknown because it involves a computation of exponential cost. The goal is to find b that maximizes $F(b)$.

To describe the MCMC algorithm, let $\epsilon_i \in Z_2^p$ be the p-bit string with ith coordinate 1 and other coordinates 0. Note that $b, c \in Z_2^p$ differ on a single coordinate if and only if $c = b + \epsilon_i$ for some i.

MCMC Algorithm:

1. Initialize the search with an arbitrary bit string, say, $b = (0, \ldots, 0)$.

2. Calculate the plausibility $F(b)$.

3. Randomly choose an integer $i \in \{1, \ldots, p\}$ with uniform probability $1/p$ and set $b^* = b + \epsilon_i$.

4. Calculate $F(b^*)$ and let $q = \min\{1, F(b^*)/F(b)\}$. Replace b with b^* with probability q.

5. Repeat 3, 4, 5.

Figure 3 shows results for the first 7 eigenfunctions of two human silhouettes. We used 6 runs of the chain with $t = 0.01$ and selected the most plausible bit string visited. Rows (a) and (b) show the initial arbitrary choices of eigenfunctions, and row (c) shows the choices prescribed by the MCMC algorithm. Note that all but the third align correctly. This is caused by the fact that ϕ_3 captures a sagittal plane (near) symmetry of the shape. However, this does not have a noticeable effect in the calculation of shape distance because it is robust to ambiguities associated with near self-isometries of a shape.

5. Clustering

We carried out a clustering experiment to illustrate the ability of the metric to classify shapes. The data set comprises 24 meshes, each on the order of 10^3 vertices, representing the contour surfaces of 6 caudate nuclei (C), 6 hippocampi (H), 6 putamina (P), and 6 thalami (T). The data was provided by the Center for Morphometric Analysis at Massachusetts General Hospital and is available at http://www.cma.mgh.harvard.edu/ibsr/. Figure 4 shows four samples from each group. We used a scale-invariant 12-dimensional HK representation with $t = 1$. The parameter in the expression of F was set to $\sigma = 0.1$. We calculated the pairwise D_H distance between all shapes and the dendrogram showed that all shapes cluster correctly. To visualize the results, Figure 5 shows a 2D multi-dimensional scaling (MDS) representation of the distance data. The colors represent: red (H), green (P), purple (T), cyan (C). The clusters that lie closest together with respect to D_H are those formed by putamina and thalami, but the cluster separation is sharp. The shapes in these two groups are visually the most similar.

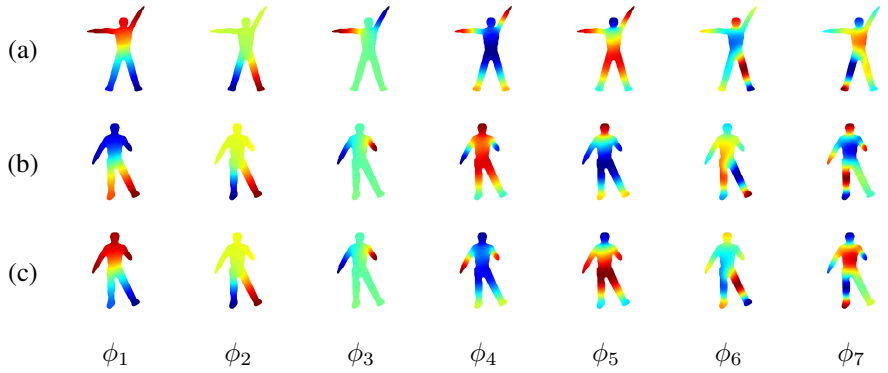

Figure 3. Matching eigenfunctions: (a) and (b) show the original choices of eigenfunctions and (c) shows the final choices for the shape of (b), calculated with MCMC.

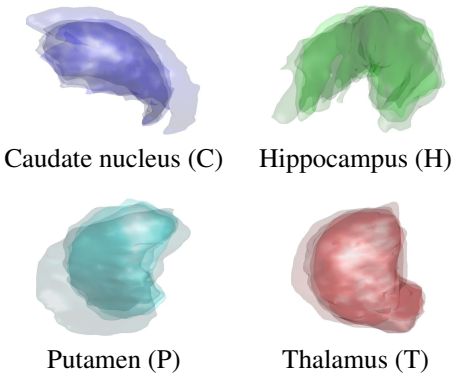

Figure 4. Sixteen samples of the shapes used in the clustering experiment.

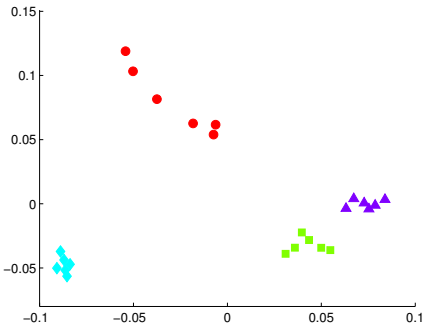

Figure 5. 2D MDS realization of the distance data.

6. Summary and Discussion

We used the heat kernel to construct a scale-space representation of shape that depends only on its intrinsic geometry. Shape dissimilarity was quantified using metrics derived from the Hausdorff distance of subsets of Euclidean space and an MCMC algorithm was developed for the computation of distances. Both scale-invariant and scale-sensitive models were discussed, and clustering experiments were carried out to illustrate the applicability of the methods. In future work, we will compare the relative merits of our HK representation to other representations based on Laplacian eigenfunctions, e.g. [5, 6].

References

[1] J. Bates, Y. Wang, X. Liu, and W. Mio. Registration of contours of brain structures through a heat-kernel representation of shape. In *ISBI*, 2009.

[2] P. Diaconis. The Markov chain Monte Carlo revolution. *Bull. Amer. Math. Soc.*, 46(2):179–205, 2009.

[3] V. Jain, H. Zhang, and O. van Kaick. Non-rigid spectral correspondence of triangle meshes. *Internat. J. Shape Modeling*, 13(1):101–124, 2007.

[4] F. Mémoli and G. Sapiro. A theoretical and computational framework for isometry invariant recognition of point cloud data. *Found. Comp. Maths.*, 5(4):313–347, 2005.

[5] R. M. Rustamov. Laplace-Beltrami eigenfunctions for deformation invariant shape representation. In *Symposium on Geometry Processing*, pages 225–233, 2007.

[6] J. Sun, M. Ovsjanikov, and L. J. Guibas. A concise and provably informative multi-scale signature based on heat diffusion. *Comp. Graph. Forum*, 28(5):1383–1392, 2009.