

Deformable Shape Retrieval by Learning Diffusion Kernels

Yonathan Aflalo¹, Alexander M. Bronstein²,
Michael M. Bronstein³, and Ron Kimmel¹

¹ Technion, Israel Institute of Technology, Haifa, Israel
`{yaflalo,ron}@cs.technion.ac.il`

² Dept. of Electrical Engineering, Tel Aviv University, Israel
`bron@eng.tau.ac.il`

³ Inst. of Computational Science, Faculty of Informatics,
Università della Svizzera Italiana, Lugano, Switzerland
`michael.bronstein@usi.ch`

Abstract. In classical signal processing, it is common to analyze and process signals in the frequency domain, by representing the signal in the Fourier basis, and filtering it by applying a transfer function on the Fourier coefficients. In some applications, it is possible to design an optimal filter. A classical example is the Wiener filter that achieves a minimum mean squared error estimate for signal denoising. Here, we adopt similar concepts to construct optimal diffusion geometric shape descriptors. The analogy of Fourier basis are the eigenfunctions of the Laplace-Beltrami operator, in which many geometric constructions such as diffusion metrics, can be represented. By designing a filter of the Laplace-Beltrami eigenvalues, it is theoretically possible to achieve invariance to different shape transformations, like scaling. Given a set of shape classes with different transformations, we learn the optimal filter by minimizing the ratio between knowingly similar and knowingly dissimilar diffusion distances it induces. The output of the proposed framework is a filter that is optimally tuned to handle transformations that characterize the training set.

1 Introduction

Recent efforts have shown the importance of *diffusion geometry* in the field of pattern recognition and shape analysis. Such methods based on geometric analysis of diffusion or random walk processes that were first introduced in theoretical geometry [1] have matured into practical applications in the fields of manifold learning [7] and where more recently introduced to shape analysis [9]. In the shape analysis community, diffusion geometry methods were used to define low-dimensional representations for manifolds [7,16], build intrinsic distance metrics and construct shape distribution descriptors [16,10,5], define spectral signatures [15] (shape-DNA), local descriptors [18,6], and bags of features [4]. Diffusion embeddings were used for finding correspondence between shapes [11] and detecting intrinsic symmetries [13].

In many settings, the construction of diffusion geometry boils down to the definition of a *diffusion kernel*, whose choice is problem dependent. Ideally, such an operator should possess certain invariance properties desired in a specific application. For example, the commute time kernel is invariant to scaling transformations of the shape.

In this paper, we propose a framework for *supervised learning* of an optimal diffusion kernel on a training set containing multiple shape classes and multiple transformations of each shape. Considering diffusion kernels related to heat diffusion properties and diagonalized in the eigenbasis of the Laplace-Beltrami operator, we can pose the problem as finding an optimal filter on the Laplace-Beltrami eigenvalues. Optimization criterion is the discriminativity between different shape classes and the invariance to within-class transformations.

The rest of the paper is organized as follows. In Section 2, we review the theoretical foundations of diffusion geometry. Section 3 formulates the problem of optimal kernel learning and its discretization. Section 4 presents experimental results. Finally, Section 5 concludes the paper.

2 Background

2.1 Diffusion Geometry

We model a shape as a Riemannian manifold X embedded into \mathbb{R}^3 . Equipping the manifold with a measure μ (e.g., the standard area measure), we also define an inner product on real functions on X by $\langle f, g \rangle = \int fg d\mu$. A function $k : X \times X \rightarrow \mathbb{R}$ is called a diffusion kernel if it satisfies the following conditions

1. *Non-negativity*: $k(x, x) \geq 0$.
2. *Symmetry*: $k(x, y) = k(y, x)$.
3. *Positive semidefiniteness*: for every bounded f ,

$$\iint k(x, y) f(x) f(y) d(\mu \times \mu) \geq 0.$$
4. *Square integrability*: $\iint k^2(x, y) d(\mu \times \mu) < \infty$.
5. *Conservation*: $\int k(\cdot, y) d\mu = \int k(x, \cdot) d\mu = 1$.

A kernel function can also be considered as a linear operator on all the functions defined on X , $(\mathbf{K}f)(y) = \int k(x, y) f(x) d\mu$. We notice that the operator \mathbf{K} is self-adjoint admitting a discrete eigendecomposition $\mathbf{K}\phi_i = \lambda_i \phi_i$, with $0 \leq \lambda_i \leq 1$ by virtue of the properties of the kernel. Spectral theorem allows us to write

$$k(x, y) = \sum_{i=0}^{\infty} \lambda_i \phi_i(x) \phi_i(y).$$

2.2 Heat Diffusion

There exists a large variety of possibilities to define a diffusion kernel and the related diffusion operator. Here, we restrict our attention to operators describing *heat diffusion*. Heat diffusion on surfaces is governed by the *heat equation*,

$$\left(\Delta_X + \frac{\partial}{\partial t}\right)u(x, t) = 0; \quad u(x, 0) = u_0(x), \quad (1)$$

where $u(x, t)$ is the distribution of heat on the surface at point x in time t , u_0 is the initial heat distribution, and Δ_X is the positive-semidefinite *Laplace-Beltrami operator*, a generalization of the second-order Laplacian differential operator Δ to non-Euclidean domains.

On Euclidean domains ($X = \mathbb{R}^m$), the classical approach to the solution of the heat equation is by representing the solution as a product of temporal and spatial components. The spatial component is expressed in the Fourier domain, based on the observation that the Fourier basis is the eigenbasis of the Laplacian Δ , and the corresponding eigenvalues are the frequencies of the Fourier harmonics. A particular solution for a point initial heat distribution $u_0(x) = \delta(x - y)$ is called the *heat kernel* $h_t(x - y) = \frac{1}{(4\pi t)^{m/2}}e^{-\|x-y\|^2/4t}$, which is shift-invariant in the Euclidean case. A general solution for any initial condition u_0 is given by convolution $\mathbf{H}^t u_0 = \int_{\mathbb{R}^m} h_t(x - y)u_0(y)dy$, where \mathbf{H}^t is referred to as *heat operator*.

In the non-Euclidean case, the eigenfunctions of the Laplace-Beltrami operator $\Delta_X \phi_i = \lambda_i \phi_i$ can be regarded as a “Fourier basis”, and the eigenvalues can be interpreted as the “spectrum”. The heat kernel is not shift-invariant but can be expressed as an explicit short time kernel [17] $h_t(x, y) = \sum_{i=0}^{\infty} e^{-t\lambda_i} \phi_i(x)\phi_i(y)$.

It can be shown that the heat operator is related to the Laplace-Beltrami operator as $\mathbf{H}^t = e^{-t\Delta}$, and as a result, it has the same eigenfunctions ϕ_i and corresponding eigenvalues $e^{-t\lambda_i}$. It can be thus seen as a particular instance of a more general family of diffusion operators \mathbf{K} diagonalized by the eigenbasis of the Laplace-Beltrami operator, namely \mathbf{K} 's as defined in the previous section but restricted to have the eigenfunctions ϕ_i of Δ_X . The corresponding diffusion kernels can be expressed as

$$k(x, y) = \sum_{i=0}^{\infty} K(\lambda_i)\phi_i(x)\phi_i(y), \quad (2)$$

where $K(\lambda)$ is some function (in the case of \mathbf{H}_t , $K(\lambda) = e^{-t\lambda}$) that can be thought of as the *transfer function* of a low-pass filter. Using this signal processing analogy, the kernel $k(x, y)$ can be interpreted as the point spread function at a point y , and the action of the diffusion operator $\mathbf{K}f$ on a function f on X can be thought of as the application of the point spread function by means of a shift-variant version of convolution. In what follows, we will freely interchange between $k(x, y)$ and $K(\lambda)$ referring to both as kernels.

2.3 Diffusion Distances

Since a diffusion kernel $k(x, y)$ measures the degree of proximity between x and y , it can be used to define a metric

$$d^2(x, y) = \|k(x, \cdot) - k(y, \cdot)\|_{L^2(X)}^2, \tag{3}$$

on X , dubbed as the *diffusion distance* by Coifman and Lafon [7]. Another way to interpret the latter distance is by considering the embedding $\Psi : x \mapsto L^2(X)$ by which each point x on X is mapped to the function $\Psi(x) = k(x, \cdot)$. The embedding Ψ is an isometry between X equipped with diffusion distance and $L^2(X)$ equipped with the standard L^2 metric, since $d(x, y) = \|\Psi(x) - \Psi(y)\|_{L^2(X)}$. As a consequence of Parseval’s theorem, the diffusion distance can also be written as

$$d^2(x, y) = \sum_{i=0}^{\infty} K^2(\lambda_i)(\phi_i(x) - \phi_i(y))^2. \tag{4}$$

Here as well we can define an isometric embedding $\Phi : x \mapsto \ell^2$ with $\Phi(x) = \{K(\lambda_i)\phi_i(x)\}_{i=0}^{\infty}$, termed as the *diffusion map* by Lafon. The diffusion distance can be casted as $d(x, y) = \|\Phi(x) - \Phi(y)\|_{\ell^2}$.

2.4 Invariance

The choice of a diffusion operator, or equivalently, the transfer function $K(\lambda)$, is related to the *invariance* of the corresponding diffusion distance.

For example, consider the case of scaling transformation, in which a shape X is uniformly scaled by a factor of α . Abusing the notations we denote by αX the new shape, whose Laplace-Beltrami operator now satisfies $\Delta_{\alpha X} f = \alpha^{-2} \Delta_X f$. Since the eigenbasis is orthonormal ($\|\phi_i\| = 1$), it follows that if ϕ_i is an eigenfunction of Δ_X associated to the eigenvalue λ_i , then $\frac{1}{\alpha}\phi_i$ is an eigenfunction of $\Delta_{\alpha X}$ associated with the eigenvalue $\lambda_i\alpha^{-2}$.

In order to obtain diffusion distance d^2 invariant to scaling transformations, we have to ensure that $K^2(\lambda_i\alpha^{-2})\alpha^{-2} = K^2(\lambda_i)$, which is achieved for $K(\lambda) = \lambda^{-1/2}$. This kernel is known as the *commute-time kernel*, and the associated diffusion distance

$$d^2(x, y) = \sum_{i=0}^{\infty} \frac{1}{\lambda_i} (\phi_i(x) - \phi_i(y))^2. \tag{5}$$

as the *commute-time distance*.

2.5 Distance Distributions

Though diffusion metrics contain significant amount of information about the geometry of the underlying shape, direct comparison of metrics is problematic since it requires computation of correspondence between shapes. A common

way to circumvent the need of correspondence is by representing a metric by its distribution, and measuring the similarity of two shapes by comparing the distributions of the respective metrics.

A metric d on X naturally pushes forward the product measure $\mu \times \mu$ on $X \times X$ (i.e., the measure defined by $d(\mu \times \mu)(x, y) = d\mu(x)d\mu(y)$) to the measure $F = d_*(\mu \times \mu)$ on $[0, \infty)$ defined as $F(I) = (\mu \times \mu)(\{(x, y) : d(x, y) \in I\})$ for every measurable set $I \subset [0, \infty)$. F can be fully described by means of a cumulative distribution function, denoted by

$$F(\delta) = \int_0^\delta dP = \int \chi_{d(x,y) \leq \delta} d\mu(x)d\mu(y) \tag{6}$$

with some abuse of notation (here χ is the indicator function). $F(\delta)$ defined this way is the measure of pairs of points the distance between which is no larger than δ ; $F(\infty) = \mu^2(X)$ is the squared area of the surface X . The density function (empirically approximated as a histogram) can be defined as the derivative $f(\delta) = \frac{d}{d\delta}F(\delta)$. Sometimes, it is convenient to work with normalized distributions, $\hat{F} = F/F(\infty)$ and the corresponding density functions, \hat{f} , which can be interpreted as probabilities.

Using this idea, comparison of two metric measure spaces reduces to the comparison of measures on $[0, \infty)$, or equivalently, comparison of un-normalized or normalized distributions, which is carried out using one of the standard distribution dissimilarity criteria used in statistics, such as L_p or *normalized L_p* , *Kullback-Leibler divergence*, *Bhattacharyya dissimilarity*, χ^2 *dissimilarity*, or *earth mover’s distance* (EMD).

3 Optimal Diffusion Kernels

The main idea of this paper lies in designing an optimal task-specific transfer function $K(\lambda)$ such that the resulting diffusion distance distribution will lead to best discrimination between shapes of a certain class while being insensitive as much as possible to a certain class of transformations.

Let us be given a shape X and some deformation τ such that $Y = \tau(X)$ is also a valid shape. Equipping each of the shapes with its Laplace-Beltrami operator, we define $\Delta_X \phi_i = \lambda_i \phi_i$ on X and $\Delta_{X'} \phi'_i = \lambda'_i \phi'_i$ on Y . A transfer function $K(\lambda)$ defines the diffusion kernel $k(x, x') = \sum_{i \geq 0} K^2(\lambda_i) \phi_i(x) \phi_i(x')$ on X , and $k'(y, y') = \sum_{i \geq 0} K^2(\lambda'_i) \phi'_i(y) \phi'_i(y')$ on Y . We aim at selecting K in such a way that for corresponding pairs of points (x, x') and $(y, y') = (\tau(x), \tau(x'))$ the two kernels coincide as much as possible, while differing as much as possible for non-corresponding points. Denoting by $P = \{((x, x'), (\tau(x), \tau(x'))) : x, x' \in X\}$ the set of all corresponding pairs (positives), and by $N = \{((x, x'), (y, y')) : x, x' \in X, (y, y') \neq (\tau(x), \tau(x'))\}$ the set of all non-corresponding pairs (negatives), we minimize

$$\min_{K(\lambda)} \frac{\sum_{((x,x'),(y,y')) \in P} (k(x, x') - k'(y, y'))^2}{\sum_{((x,x'),(y,y')) \in N} (k(x, x') - k'(y, y'))^2}. \tag{7}$$

We remark that while there is a multitude of reasonable alternative objective functions, in what follows we choose to minimize the above ratio because as it will be shown it lends itself to a simple algebraic problem.

The choice of an appropriate function K can lead to invariance of the kernel under some transformations. For example, the *commute time* kernel $K(\lambda) = \frac{1}{\sqrt{\lambda}}$ is invariant under global scaling. On the other hand, optimal K should be discriminative enough to distinguish between shapes not being one a transformation of the other. This spirit is similar to linear discriminant analysis (LDA) and Wiener filtering and, to the best of our knowledge, has never been proposed before to construct optimal diffusion metrics.

3.1 Discretization

We represent the surface X as triangular mesh with n faces constructed upon the samples $\{\mathbf{x}_1, \dots, \mathbf{x}_n\}$ The computation of discrete diffusion kernels $k(\mathbf{x}_1, \mathbf{x}_2)$ requires computing discrete eigenvalues and eigenfunctions of the discrete Laplace-Beltrami operator. The latter can be computed directly using the finite elements method (FEM) [15], or by discretization of the Laplace operator on the mesh followed by its eigendecomposition. Here, we adopt the second approach according to which the discrete Laplace-Beltrami operator is expressed in the following generic form,

$$(\Delta f)_i = \frac{1}{a_i} \sum_j w_{ij} (f_i - f_j), \tag{8}$$

where $f_i = f(\mathbf{x}_i)$ is a scalar function defined on the mesh, w_{ij} are weights, and a_i are normalization coefficients. In matrix notation, (8) can be written as $\Delta \mathbf{f} = \mathbf{A}^{-1} \mathbf{W} \mathbf{f}$, where \mathbf{f} is an $m \times 1$ vector and $\mathbf{W} = \text{diag} \left\{ \sum_{l \neq i} w_{il} \right\} - w_{ij}$.

The discrete eigenfunctions and eigenvalues are found by solving the *generalized eigendecomposition* [9] $\mathbf{W} \Phi = \mathbf{A} \Phi \Lambda$, where $\Lambda = \text{diag}\{\lambda\}$ is a diagonal matrix of eigenvalues $\lambda = (\lambda_1, \dots, \lambda_n)^T$, and $\Phi = (\phi_l(x_i))$ is the matrix of the corresponding eigenvectors. Similarly, we triangulate the shape Y and get $\mathbf{A}' \Phi' = \text{diag}\{\lambda'\} \mathbf{W}' \Phi'$.

Different choices of \mathbf{W} have been studied, depending on which continuous properties of the Laplace-Beltrami operator one wishes to preserve [8,19]. For triangular meshes, a popular choice adopted in this paper is the *cotangent weight* scheme [14,12], in which

$$w_{ij} = \begin{cases} (\cot \beta_{ij} + \cot \gamma_{ij})/2 & : \mathbf{x}_j \in \mathcal{N}(\mathbf{x}_j); \\ 0 & : \text{else,} \end{cases} \tag{9}$$

where β_{ij} and γ_{ij} are the two angles opposite to the edge between vertices \mathbf{x}_i and \mathbf{x}_j in the two triangles sharing the edge.

We denote by $P = \{(i_m, j_m), (i'_m, j'_m)\}$ the collection of corresponding pairs of vertex indices on X and Y (that is, $i_m \leftrightarrow i'_m$ and $j_m \leftrightarrow j'_m$), and by N the collection of non-corresponding pairs. Denoting by \mathbf{C}_+ and \mathbf{C}'_+ two matrices whose ml -th elements are the products $\phi_l(x_{i_m})\phi_l(x_{j_m})$ and $\phi'_l(y_{i'_m})\phi'_l(y_{j'_m})$, respectively, for $((i_m, j_m), (i'_m, j'_m)) \in P$, we have $\mathbf{k}_+ = \mathbf{C}_+K^2(\boldsymbol{\lambda})$ and $\mathbf{k}'_+ = \mathbf{C}'_+K^2(\boldsymbol{\lambda}')$, where the m -th elements of \mathbf{k}_+ and \mathbf{k}'_+ are $k(x_{i_m}, x_{j_m})$ and $k(x_{i'_m}, x_{j'_m})$, respectively, and $K^2(\boldsymbol{\lambda}) = (K^2(\lambda_1), \dots, K^2(\lambda_n))^T$. Exactly in the same way, the vectors \mathbf{k}_- and \mathbf{k}'_- corresponding to the negative pairs in N are obtained.

In order to make possible the optimization over all functions K , we fix a grid $\gamma = (\gamma_1, \dots, \gamma_r)$ or r points on which $\mathbf{k} = (K^2(\gamma_1), \dots, K^2(\gamma_r))^T$ is evaluated. In this notation, our optimization problem becomes with respect to the elements of \mathbf{k} . Since the grids γ , $\boldsymbol{\lambda}$ and $\boldsymbol{\lambda}'$ are incompatible, we define the interpolation operators \mathbf{I} and \mathbf{I}' transferring a function from the grid γ to the grids $\boldsymbol{\lambda}$ and $\boldsymbol{\lambda}'$: $K^2(\lambda_i) = \mathbf{I}\mathbf{k}$, and $K^2(\lambda'_i) = \mathbf{I}'\mathbf{k}$. This yields $\mathbf{k}_\pm = \mathbf{C}_\pm\mathbf{I}\mathbf{k}$ and $\mathbf{k}'_\pm = \mathbf{C}'_\pm\mathbf{I}'\mathbf{k}$. Substituting the latter result into (7) gives the following minimization problem:

$$\begin{aligned} \mathbf{k}^* &= \arg \min_{\mathbf{k} \geq 0} \frac{\|\mathbf{k}_+ - \mathbf{k}'_+\|^2}{\|\mathbf{k}_- - \mathbf{k}'_-\|^2} = \arg \min_{\mathbf{k} \geq 0} \frac{\|(\mathbf{C}_+\mathbf{I} - \mathbf{C}'_+\mathbf{I}')\mathbf{k}\|^2}{\|(\mathbf{C}_-\mathbf{I} - \mathbf{C}'_-\mathbf{I}')\mathbf{k}\|^2} \\ &= \arg \min_{\mathbf{k} \geq 0} \frac{\mathbf{k}^T \mathbf{P} \mathbf{k}}{\mathbf{k}^T \mathbf{N} \mathbf{k}} = \mathbf{N}^{-\frac{1}{2}} \arg \min_{\substack{\mathbf{k} \geq 0 \\ \|\bar{\mathbf{k}}\|=1}} \bar{\mathbf{k}}^T \mathbf{N}^{-\frac{T}{2}} \mathbf{P} \mathbf{N}^{-\frac{1}{2}} \mathbf{k}, \end{aligned} \quad (10)$$

where $\mathbf{P} = (\mathbf{C}_+\mathbf{I} - \mathbf{C}'_+\mathbf{I}')^T(\mathbf{C}_+\mathbf{I} - \mathbf{C}'_+\mathbf{I}')$ and $\mathbf{N} = (\mathbf{C}_-\mathbf{I} - \mathbf{C}'_-\mathbf{I}')^T(\mathbf{C}_-\mathbf{I} - \mathbf{C}'_-\mathbf{I}')$. Note that the matrices \mathbf{P} and \mathbf{N} are of fixed size $r \times r$ and can be constructed without directly constructing the potentially huge matrices \mathbf{C}_\pm and \mathbf{C}'_\pm . This makes the above problem computationally efficient even on very large training sets.

3.2 Interpolation Operators

Among a plethora of methods for designing the interpolation operations \mathbf{I} and \mathbf{I}' on one-dimensional intervals, we found that regularized spline fitting produced best results. For that purpose, let $\{s_i(\lambda)\}$ be a set of q functions defined on the interval $[\lambda_{\min}, \lambda_{\max}]$. We represent the kernel transfer function as the sum

$$K^2(\lambda) = \sum_{i=1}^q a_i s_i(\lambda) \quad (11)$$

and look for the vector of coefficients $\mathbf{a} = (a_1, \dots, a_q)^T$. Denoting $\mathbf{S} = (s_1(\boldsymbol{\lambda}), \dots, s_q(\boldsymbol{\lambda}))$ with $s_i(\boldsymbol{\lambda}) = (s_i(\lambda_1), \dots, s_i(\lambda_n))^T$, we have $\mathbf{k} = \mathbf{S}\mathbf{a}$. Similarly, for $\mathbf{S}' = (s_1(\boldsymbol{\lambda}'), \dots, s_q(\boldsymbol{\lambda}'))$, we have $\mathbf{k}' = \mathbf{S}'\mathbf{a}$.

To impose the smoothness of the kernel $K(\lambda)$, we add the regularization term

$$R(K) = \int_{\lambda_{\min}}^{\lambda_{\max}} \|\delta K^2(\lambda)\|^2 d\lambda = \int_{\lambda_{\min}}^{\lambda_{\max}} \left(\sum_{i=1}^q a_i \nabla s_i(\lambda) \right)^2 d\lambda = \mathbf{a}^T \mathbf{R} \mathbf{a}, \quad (12)$$

where the ij -th elements of \mathbf{R} are given by $(\mathbf{R})_{ij} = \int_{\lambda_{\min}}^{\lambda_{\max}} \nabla s_i(\lambda) s_j(\lambda) d\lambda$.

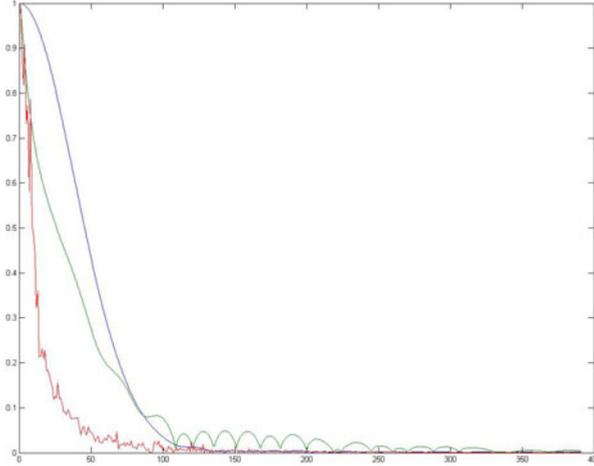


Fig. 1. Optimal kernel designed using straightforward nearest neighbor interpolation (red), splines without smoothness (green), and splines with the smoothness term (blue)

In these terms, the optimization problem (10) becomes

$$\begin{aligned}
 \mathbf{a}^* &= \arg \min_{\mathbf{a}} \frac{\mathbf{a}^T (\mathbf{C}_+ \mathbf{S} - \mathbf{C}'_+ \mathbf{S}')^T (\mathbf{C}_+ \mathbf{S} - \mathbf{C}'_+ \mathbf{S}') \mathbf{a}}{\mathbf{a}^T (\mathbf{C}_- \mathbf{S} - \mathbf{C}'_- \mathbf{S}')^T (\mathbf{C}_- \mathbf{S} - \mathbf{C}'_- \mathbf{S}') \mathbf{a}} + \eta \mathbf{a}^T \mathbf{R} \mathbf{a} \\
 &= \mathbf{N}^{-\frac{1}{2}} \arg \min_{\|\bar{\mathbf{a}}\|=1} \bar{\mathbf{a}}^T \mathbf{N}^{-\frac{T}{2}} (\mathbf{P} + \eta \mathbf{R}) \mathbf{N}^{-\frac{1}{2}} \bar{\mathbf{a}}, \tag{13}
 \end{aligned}$$

where now $\mathbf{P} = (\mathbf{C}_+ \mathbf{S} - \mathbf{C}'_+ \mathbf{S}')^T (\mathbf{C}_+ \mathbf{S} - \mathbf{C}'_+ \mathbf{S}')$, $\mathbf{N} = (\mathbf{C}_- \mathbf{S} - \mathbf{C}'_- \mathbf{S}')^T (\mathbf{C}_- \mathbf{S} - \mathbf{C}'_- \mathbf{S}')$, and η is a parameter controlling the smoothness of the obtained kernel. The effect of the smoothness term is illustrated in Figure 1.

4 Results

In our experiments, to build the training set, we used the SHREC'10 correspondence benchmark [2]. The dataset contained high-resolution shapes (10,000 – 30,000 vertices) organized in seven shape classes with 55 simulated transformations of varying strength in each class (Figure 2) Testing was performed on the SHREC'10 shape retrieval benchmark [3], containing a total of 1184 shapes. Retrieval performance was evaluated using precision/recall characteristic. *Precision* $P(r)$ is defined as the percentage of relevant shapes in the first r top-ranked retrieved shapes (in the used benchmark, transformed shapes were used as queries, while a single relevant null shape existed in the database for each query). *Mean average precision* (mAP), defined as

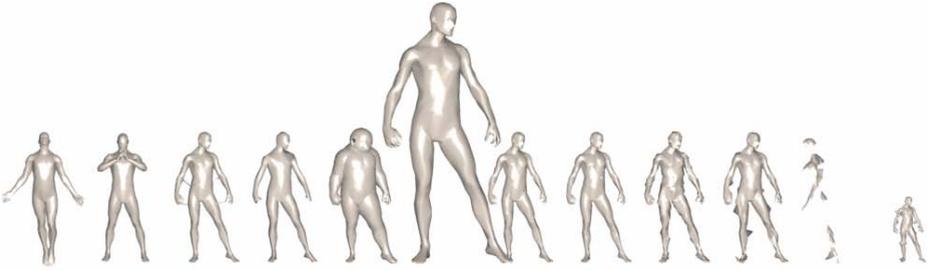


Fig. 2. Transformations of the human shape used as queries (shown in strength 5, left to right): null, isometry, topology, sampling, local scale, scale, holes, micro holes, noise, shot noise, partial, all

$$mAP = \sum_r P(r) \cdot rel(r),$$

where $rel(r)$ is the relevance of a given rank, was used as a single measure of performance. Ideal performance retrieval performance results in first relevant match with $mAP=100\%$. Discretization of the Laplace-Beltrami was based on the cotangent weight formula (9).

In the first experiment, we used our approach to learn a scale-invariant diffusion kernel. We used a training set containing only scaling transformations of the shapes. As can be seen from Figure 3, the learned diffusion kernel is very close to the theoretically-optimal commute-time kernel $K(\lambda) = \lambda^{-1/2}$.

In the second experiment, we extended the training set to include all the shape transformations, resulting in a kernel shown in Figure 4 (red). The learned kernel was used to compute diffusion distance distributions, which were compared to compute the shape similarity, following the spectral distance framework [5]. The performance results with this kernel are summarized in Table 1 (fifth column). For comparison, performance using the commute time kernel is shown (Table 1, sixth column).

In the third experiment, instead of designing a kernel with a discretization of $K(\lambda)$, we used a parametric kernel of the form $K(\lambda) = \exp(-t\lambda)$ and optimized our criterion for the time scale t . The optimal scale was found to be $t^* = 1011$; the performance results with this kernel are summarized in Table 1 (fourth column). For comparison, we show the performance of the same kernel with two other values of the parameter, $t = 700$, and 1700 (Table 1, second and third columns).

In the fourth experiment, we used the diagonal our optimal non-parametric diffusion kernel $k(x, x)$ as a local scalar shape descriptor at each point, similar to the heat kernel signature [18]. A global descriptor was constructed as the histogram of the values of $k(x, x)$ on the entire shapes. We notice that both the local descriptors (Figure 5, top) and the global descriptors (Figure 5, bottom) resulting from our learned diffusion kernel signature computed on two different transformations of a shape are very close one to the other.

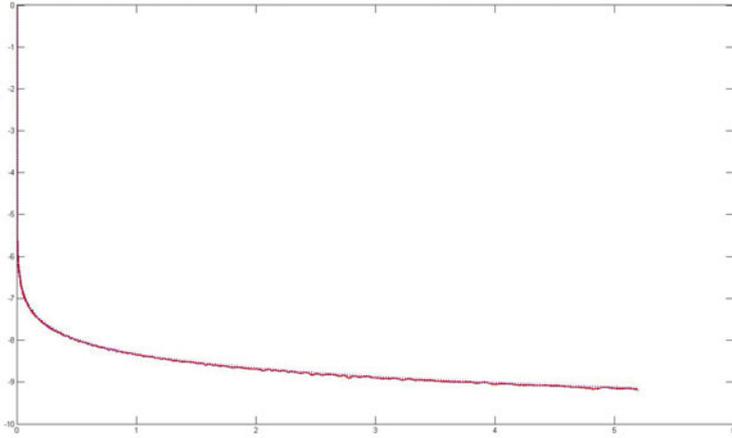


Fig. 3. Theoretical scale invariant (commute time) kernel (blue) and the learned kernel on examples of scaling transformations (red)

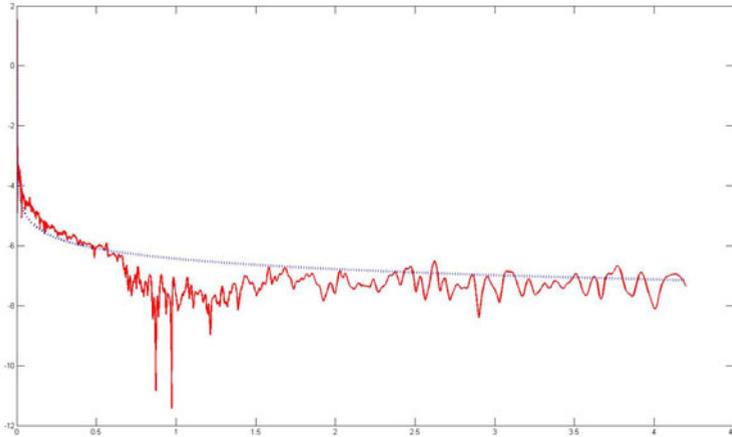
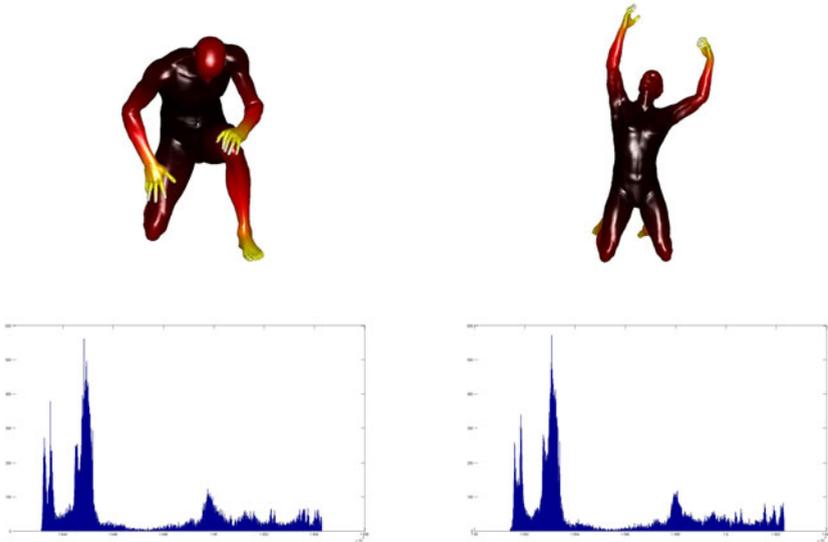


Fig. 4. The lerned kernel using all transformations (red). For comparison, the commute time kernel is shown (blue).

Table 1. Shape retrieval performance (mAP in %) using the spectral distance with different diffusion kernels

Transformation	Heat ($t = 700$)	Heat ($t = 1700$)	Optimal param. ($t^* = 1011$)	Optimal non-param.	Commute time
Isometry	100	99.23	100	98.21	97.95
Topology	91.28	80.36	86.79	77.33	79.16
Holes	82.3	90.33	87.81	72.48	73.94
Micro holes	100	100	100	100	100
Scale	30.97	32.66	32.44	100	100
Local scale	65.64	70.79	70.73	67.92	68.22
Sampling	100	99.23	100	98.21	98.21
Noise	99.23	100	98.46	100	100
Shot noise	99.23	100	99.23	99.23	98.65
Partial	5.54	7.37	6.01	8.06	31.03
All	64.15	64.36	69.56	64.66	64.51

**Fig. 5.** Top: diagonal of the diffusion kernel $k(x, x)$ used as a local descriptor. Bottom: histogram of the local descriptors.

5 Conclusions

We provided a design framework for kernels that optimize for the ratio between the within class and between classes required for shape recognition under typical type of deformations. So far, our experiments show that the commute time distance is dominating as an optimal filter for the mix of distortions we used. In our future experiments we will investigate the deviation from that type of a filter and try to come up with design framework for specific types of distortions.

References

1. Bérard, P., Besson, G., Gallot, S.: Embedding riemannian manifolds by their heat kernel. *Geometric and Functional Analysis* 4(4), 373–398 (1994)
2. Bronstein, A.M., Bronstein, M.M., Bustos, B., Castellani, U., Crisani, M., Falcidieno, B., Guibas, L.J., Sipiran, I., Kokkinos, I., Murino, V., Ovsjanikov, M., Patané, G., Spagnuolo, M., Sun, J.: SHREC 2010: robust feature detection and description benchmark. In: *Proc. 3DOR* (2010)
3. Bronstein, A.M., Bronstein, M.M., Castellani, U., Falcidieno, B., Fusiello, A., Godil, A., Guibas, L.J., Kokkinos, I., Lian, Z., Ovsjanikov, M., Patané, G., Spagnuolo, M., Toldo, R.: Shrec 2010: robust large-scale shape retrieval benchmark. In: *Proc. 3DOR* (2010)
4. Bronstein, A.M., Bronstein, M.M., Ovsjanikov, M., Guibas, L.J.: Shape google: a computer vision approach to invariant shape retrieval. In: *Proc. NORDIA* (2009)
5. Bronstein, M.M., Bronstein, A.M.: Shape recognition with spectral distances. *Trans. PAMI* (2010) (to appear)
6. Bronstein, M.M., Kokkinos, I.: Scale-invariant heat kernel signatures for non-rigid shape recognition. In: *Proc. CVPR* (2010)
7. Coifman, R.R., Lafon, S.: Diffusion maps. *Applied and Computational Harmonic Analysis* 21, 5–30 (2006)
8. Floater, M.S., Hormann, K.: Surface parameterization: a tutorial and survey. In: *Advances in Multiresolution for Geometric Modelling*, vol. 1 (2005)
9. Lévy, B.: Laplace-Beltrami eigenfunctions towards an algorithm that “understands” geometry. In: *Proc. Shape Modeling and Applications* (2006)
10. Mahmoudi, M., Sapiro, G.: Three-dimensional point cloud recognition via distributions of geometric distances. *Graphical Models* 71(1), 22–31 (2009)
11. Mateus, D., Horaud, R.P., Knossow, D., Cuzzolin, F., Boyer, E.: Articulated shape matching using laplacian eigenfunctions and unsupervised point registration. In: *Proc. CVPR* (June 2008)
12. Meyer, M., Desbrun, M., Schroder, P., Barr, A.H.: Discrete differential-geometry operators for triangulated 2-manifolds. In: *Visualization and Mathematics III*, pp. 35–57 (2003)
13. Ovsjanikov, M., Sun, J., Guibas, L.J.: Global intrinsic symmetries of shapes. *Computer Graphics Forum* 27, 1341–1348 (2008)
14. Pinkall, U., Polthier, K.: Computing discrete minimal surfaces and their conjugates. *Experimental Mathematics* 2(1), 15–36 (1993)
15. Reuter, M., Wolter, F.-E., Peinecke N.: Laplace-spectra as fingerprints for shape matching. In: *Proc. ACM Symp. Solid and Physical Modeling*, pp. 101–106 (2005)
16. Rustamov, R.M.: Laplace-Beltrami eigenfunctions for deformation invariant shape representation. In: *Proc. SGP*, pp. 225–233 (2007)
17. Spira, A., Sochen, N., Kimmel, R.: Geometric filters, diffusion flows, and kernels in image processing. In: *Handbook of Computational Geometry for Pattern Recognition, Computer Vision, Neurocomputing and Robotics*. Springer, Heidelberg (2005)
18. Sun, J., Ovsjanikov, M., Guibas, L.J.: A concise and provably informative multi-scale signature based on heat diffusion. In: *Proc. SGP* (2009)
19. Wardetzky, M., Mathur, S., Kälberer, F., Grinspun, E.: Discrete Laplace operators: no free lunch. In: *Conf. Computer Graphics and Interactive Techniques* (2008)