

Spectral descriptors for deformable shapes

A. M. Bronstein *

February 17, 2022

Abstract

Informative and discriminative feature descriptors play a fundamental role in deformable shape analysis. For example, they have been successfully employed in correspondence, registration, and retrieval tasks. In the recent years, significant attention has been devoted to descriptors obtained from the spectral decomposition of the Laplace-Beltrami operator associated with the shape. Notable examples in this family are the heat kernel signature (HKS) and the wave kernel signature (WKS). Laplacian-based descriptors achieve state-of-the-art performance in numerous shape analysis tasks; they are computationally efficient, isometry-invariant by construction, and can gracefully cope with a variety of transformations. In this paper, we formulate a generic family of parametric spectral descriptors. We argue that in order to be optimal for a specific task, the descriptor should take into account the statistics of the corpus of shapes to which it is applied (the “signal”) and those of the class of transformations to which it is made insensitive (the “noise”). While such statistics are hard to model axiomatically, they can be learned from examples. Following the spirit of the Wiener filter in signal processing, we show a learning scheme for the construction of optimal spectral descriptors and relate it to Mahalanobis metric learning. The superiority of the proposed approach is demonstrated on the SHREC’10 benchmark.

*School of Electrical Engineering, Tel Aviv University, Tel Aviv 69978, Israel. Email: bron@eng.tau.ac.il. This work was supported by the Israeli Science Foundation Grant and the German-Israeli Foundation.

1 Introduction

The notion of a *feature descriptor* is fundamental in shape analysis. A feature descriptor assigns each point on the shape a vector in some single- or multi-dimensional feature space representing the point’s local and global geometric properties relevant for a specific task. This information is subsequently used in higher-level tasks: for example, in shape matching descriptors are used to establish an initial set of potentially corresponding points [1, 2]; in shape retrieval a global shape descriptor is constructed as a bag of “geometric words” expressed in terms of local feature descriptors [3, 4]; segmentation algorithms rely on the similarity or dissimilarity between feature descriptors to partition the shape into stable and meaningful parts [5].

When constructing or choosing a feature descriptor, it is imperative to answer two fundamental questions: which shape properties the descriptor has to capture, and to which transformations of the shape it shall remain invariant.

1.1 Previous work

Early research on feature descriptors focused mainly on invariance under global Euclidean transformations (rigid motion). Classical works in this category include the shape context [6] and spin image [7] descriptors, as well as integral volume descriptors [8, 1] and multiscale local features [9] just to mention a few out of many.

In the past decade, significant effort has been invested in extending the invariance properties to non-rigid deformations. Some of the classical rigid descriptors were extended to the non-rigid case by replacing the Euclidean metric with its geodesic counterpart [10, 11]. Also, the use of conformal factors has been proposed [12]. Being intrinsic properties of a surface, both are independent of the way the surface is embedded into the ambient Euclidean space and depend only on its metric structure. This makes such descriptors invariant to inelastic bending transformations. However, geodesic distances suffer from strong sensitivity to topological noise, while conformal factors, being a local quantity, are influenced by geometric noise. Both types of

noise, virtually inevitable in real applications, limit the usefulness of such descriptors.

Recently, a family of intrinsic geometric properties broadly known as *diffusion geometry* has become growingly popular. The studies of diffusion geometry are based on the theoretical works by Berard *et al.* [13] and later by Coifman and Lafon [14] who suggested to use the eigenvalues and eigenvectors of the Laplace-Beltrami operator associated with the shape to construct invariant metrics known as diffusion distances. These distances as well as other diffusion geometric constructs have been shown significantly more robust compared to their geodesic counterparts [15, 16]. Diffusion geometry offers an intuitive interpretation of many shape properties in terms of spacial frequencies and allows to use standard harmonic analysis tools. Also, recent advances in the discretization of the Laplace-Beltrami operator bring forth efficient and robust numerical and computational tools.

These methods were first explored in the context of shape processing by Lévy [17]. Several attempts have also been made to construct feature descriptors based on diffusion geometric properties of the shape. Rustamov [18] proposed to construct the *global point signature* (GPS) feature descriptors by associating each point with an ℓ^2 sequence based on the eigenfunctions and the eigenvalues of the Laplacian, closely resembling a diffusion map [14]. A major drawback of such a descriptor was its ambiguity to sign flips of each individual eigenfunction (or, in the most general case, to rotations and reflections in the eigenspaces corresponding to each eigenvalue).

A remedy was proposed by Sun *et al.* who in their influential paper [19] introduced the *heat kernel signature* (HKS), based on the fundamental solutions of the heat equation (heat kernels). In [20], another physically-inspired descriptor, the *wave kernel signature* (WKS) was proposed as a solution to the excessive sensitivity of the HKS to low-frequency information. As of today, these descriptors achieve state-of-the-art performance in many deformable shape analysis tasks [21, 22].

1.2 Contribution

In this paper, we remain within the diffusion geometric framework and propose a generic family of spectral feature descriptors that generalize both the HKS and the WKS. We analyze both descriptors within this framework pointing to their advantages and drawbacks, and enumerate a list of desired properties a descriptor should have.

We argue that in order to construct a good task-specific spectral descriptor, one has to be in the position of defining the spectral content of the geometric “signal” (i.e., the properties distinguishing different classes of shapes from each other) and the “noise” (i.e., the changes of the latter properties due to the deformations the shapes undergo). Both are functions of the corpus of data of interest, and the transformations invariance to which is desired. While it is notoriously difficult to characterize such properties analytically, we propose to learn them from examples in a way resembling the construction of a Wiener filter that passes frequencies containing more signal than noise, while attenuating those where the noise covers the signal.

This study was in part inspired by the insightful paper by Auby *et al.* [20], and in part is a continuation of [23] where we attempted to construct optimal diffusion metrics. However, since diffusion metrics are characterized by a single frequency response, the attempt had a modest success. On the other hand, vector-valued feature descriptors allowing for multiple frequency response functions have, in our opinion, more potential. This paper does not intend to exhaust this potential, but merely to explore a part of it.

The rest of the paper is organized as follows: In Section 2 we introduce the mathematical notation of the Laplace-Beltrami operator and its spectrum and briefly overview the state-of-the-art descriptors based on its properties. In Section 3, we indicate several drawbacks of these descriptors and analyze the properties a good descriptor should satisfy. We present a spectral descriptor generalizing the heat and the wave kernel signatures, and show an approach for learning its optimal task-specific parameters from examples. Relation to metric learning is highlighted. In Section 4, the superiority of the proposed learnable descriptor over the fixed ones is shown experimentally on the SHREC’10 non-rigid correspondence benchmark. Fi-

nally, Section 5 concludes the paper.

Since the figures visualizing the experiments in Section 4 are relatively self-explanatory, we decided to incorporate them in the flow as illustrations to the phenomena discussed in the paper even before the exact experimental setting are detailed.

2 Spectral descriptors

We model a shape as a compact two-dimensional manifold X , possibly with a boundary ∂X . The manifold is endowed with a Riemannian metric defined as a local inner product $\langle \cdot, \cdot \rangle_x$ on the tangent plane $T_x X$ at each point $x \in X$. Given a smooth scalar field f on the manifold, its gradient $\text{grad } f$ is the vector field satisfying $f(x + dr) = f(x) + \langle \text{grad } f(x), dr \rangle_x$ for every infinitesimal tangent vector $dr \in T_x X$. The inner product $\langle \text{grad } f(x), v \rangle_x$ can be interpreted as the directional derivative of f in the direction v . A directional derivative of f whose direction at every point is defined by a vector field V on the manifold is called the Lie derivative of f along V . The Lie derivative of the manifold volume (area) form along a vector field V is called the divergence of V , $\text{div } V$. The negative divergence of the gradient of a scalar field f , $\Delta f = -\text{div } \text{grad } f$, is called the *Laplacian* of f . The operator Δ is called the *Laplace-Beltrami* operator, and it generalizes the standard notion of the Laplace operator to manifolds. Note that we define the Laplacian with the negative sign to conform to the computer graphics and computational geometry convention.

2.1 Laplacian spectrum and Shape DNA

Being a positive self-adjoint operator, the Laplacian admits an eigendecomposition

$$\Delta \phi = \nu \phi \tag{1}$$

with non-negative eigenvalues ν and corresponding orthonormal eigenfunctions ϕ . Furthermore, due to the assumption that our domain is compact, the spectrum is discrete, $0 = \nu_1 < \nu_2 < \dots$.

In physics, (1) is known as the *Helmholtz equation* representing the spatial component of the wave equation. Thinking of our domain as of a vibrating membrane (with appropriate boundary conditions), the ϕ_k 's can be interpreted as natural vibration modes of the membrane, while the ν_k 's assume the meaning of the corresponding vibration frequencies. In fact, in this setting the eigenvalues have inverse area or squared spatial frequency units.

This physical interpretation leads to a natural question whether the eigenvalues of the Laplace-Beltrami operator fully determine the shape of the domain. The essence of this question was beautifully captured by Mark Kac as “*can one hear the shape of the drum?*” [24]. Unfortunately, the answer to this question is negative as there exist isospectral manifolds that are not isometric. The exact relation between the latter two classes of shapes is unknown, but it is generally believed that most isospectral manifolds are also isometric. Based on this belief, Reuter *et al.* [25] proposed to use truncated sequences of the Laplacian eigenvalues as isometry-invariant shape descriptors, dubbed by the authors as *shape DNA*.

2.2 Heat kernel signature

The Laplace-Beltrami operator plays a central role in the *heat equation* describing diffusion processes on manifolds. In our notation, the heat equation can be written as

$$\left(\Delta + \frac{\partial}{\partial t}\right)u(x, t) = 0 \tag{2}$$

where $u(x, t)$ is the distribution of heat on the manifold at point x at time t . The initial condition is some initial heat distribution $u_0(x)$ at time $t = 0$, and boundary conditions are applied in case the manifold has a boundary.

The solution of the heat equation at time t can be expressed as the application of the *heat operator*

$$u(x, t) = \int h_t(x, y)u_0(y)da(y) \tag{3}$$

to the initial distribution. The kernel $h_t(x, y)$ of this integral operator is called the *heat kernel* and it corresponds to the solution of the heat equation

at point x at time t with the initial distribution being a delta function at point y . From the signal processing perspective, the heat kernel can be interpreted as a non shift-invariant “impulse response”. It also describes the amount of heat transferred from point x to point y after time t , as well as the transition probability density from point x to point y by a random walk of length t .

According to the spectral decomposition theorem, the heat kernel can be expressed as

$$h_t(x, y) = \sum_{k \geq 1} \exp(-\nu_k t) \phi_k(x) \phi_k(y), \quad (4)$$

where $\exp(-\nu t)$ can be interpreted as its “frequency response” (note that with a proper selection of units in (3), the eigenvalues ν_k assume inverse time or frequency units). The bigger is the time parameter, the lower is the cut-off frequency of the low-pass filter described by this response and, consequently, the bigger is the support of h_t on the manifold. The quantity

$$h_t(x, x) = \sum_{k \geq 1} \exp(-\nu_k t) \phi_k^2(x), \quad (5)$$

sometimes referred to as the *autodiffusivity function* [26], describes the amount of heat remaining at point x after time t . Furthermore, for small values of t is it related to the manifold curvature according to

$$h_t(x, x) = \frac{1}{4\pi t} + \frac{K(x)}{12\pi} + \mathcal{O}(t), \quad (6)$$

where $K(x)$ denotes the Gaussian (in general, sectional) curvature at point x .

In [19], Sun *et al.* showed that under mild technical conditions, the sequence $\{h_t(x, x)\}_{t>0}$ contains *full* information about the metric of the manifold. The authors proposed to associate each point x on the manifold with a vector

$$\mathbf{p}(x) = (h_{t_1}(x, x), \dots, h_{t_n}(x, x))^T, \quad (7)$$

of the autodiffusivity functions sampled at some finite set of times t_1, \dots, t_n . The authors dubbed such a feature descriptor as the *heat kernel signature*.

In [4], an HKS-based bag-of-features approach was introduced under the name of Shape Google and was shown to achieve state-of-the-art results in deformable shape retrieval. In [27], a scale-invariant version of the HKS was proposed, and [28] extended the descriptor to volumes.

Despite its success, the heat kernel descriptor suffers from several drawbacks. First, being a collection of low-pass filters (Figure 1, top), the descriptor is dominated by low frequencies conveying information mostly about the global structure of the shape. While being important to discriminate between distinct shapes (which usually differ greatly at coarse scales), this emphasize of low frequencies damages the ability of the descriptor to precisely localize features. This phenomenon can be observed in Figure 2 (top). In fact, the distance between HKS computed at a point x and HKS of neighboring points increases slowly, while for good localization a steeper increase is required.

2.3 Wave kernel signature

A remedy to the poor feature localization of the heat kernel descriptor was proposed by Aubry *et al.* [20]. The authors proposed to replace the heat diffusion model that gives rise to the HKS by a different physical model in which one evaluates the probability of a quantum particle with a certain energy distribution to be located at a point x . The behavior of a quantum particle on a surface is governed by the Schrödinger equation

$$\left(i\Delta + \frac{\partial}{\partial t}\right)\psi(x, t) = 0 \quad (8)$$

where $\psi(x, t)$ is the complex wave function. Despite an apparent similarity to the heat equation, the multiplication of the Laplacian by the complex unity in the Schrödinger equation has a dramatic impact on the dynamics of the solution. Instead of representing diffusion, ψ now has oscillatory behavior.

Let us assume that the quantum particle has an initial energy distribution $f(e)$. Since energy is directly related to frequency, we will use $f(\nu)$ instead in order to stick to the previous notation. The solution of the




Figure 1: Examples of (unnormalized) kernels used for the computation of the heat kernel (first row), wave kernel (second row), and trained optimal kernel (last row) descriptors.




Figure 2: Normalized Euclidean distance between the descriptor at a reference point on the left foot (white dot in the leftmost column) and descriptors computed at rest of the points of the same shape (left column), its approximate isometry (middle column), and a distinct shape (right column). Twelve-dimensional descriptors based on the heat kernel (first row), wave kernel (second row), and trained optimal kernel (last row) are shown. Dark blue stands for small distance; red represents large distance.

Schrödinger equation can then be expressed in the spectral domain as [20]

$$\psi(x, t) = \sum_{k \geq 1} \exp(i\nu_k t) f(\nu_k) \phi_k(x) \quad (9)$$

(note the complex unity in the exponential!). The probability to measure the particle at a point x at time t is given by $|\psi(x, t)|^2$. By integrating over all times, the average probability

$$p(x) = \lim_{T \rightarrow \infty} \frac{1}{T} \int_0^T |\psi(x, t)|^2 dt = \sum_{k \geq 1} f^2(\nu_k) \phi_k^2(x). \quad (10)$$

to measure the particle at a point x is obtained. Note that the probability depends on the initial energy distribution f .

Aubry *et al.* considered a family of log-normal energy distributions

$$f_e(\nu) \propto \exp\left(-\frac{(\log e - \log \nu)^2}{2\sigma^2}\right) \quad (11)$$

centered around some mean log energy $\log e$ with variance σ^2 (again, we allow ourselves a certain abuse of the physics and treat energy and frequency as synonyms). This particular choice of distributions is motivated by a perturbation analysis of the Laplacian spectrum [20].

Fixing the family of energy distributions, each point on the surface is associated with a *wave kernel signature* of the form

$$\mathbf{p}(x) = (p_{e_1}(x), \dots, p_{e_n}(x))^T, \quad (12)$$

where $p_e(x)$ is the probability to measure a quantum particle with the initial energy distribution $f_e(\nu)$ at point x . The authors use logarithmically sampled e_1, \dots, e_n .

The WKS descriptor resembles the HKS in the sense that it can also be thought of as an application of a set of filters with the frequency responses $f_e^2(\nu)$. However, unlike the HKS that uses low-pass filters, the responses of the WKS are *band-pass* (Figure 1, middle). This reduces the influence of the low frequencies and allows better separation of frequency bands across the descriptor dimensions. As the result, the wave kernel descriptor exhibits superior feature localization (Figure 2, middle).

3 Spectral descriptor learning

Despite their beautiful physical interpretation, both the heat and wave kernel descriptors suffer from several drawbacks.

The fact that the WKS deemphasizes large-scale features contributes to its higher *sensitivity* (i.e., the ability to identify positives). This property is crucial in matching problems, where a small set of candidate matches on one shape is found for a collection of reference points on the other. The ability to produce a correct match within a small set of best matches (high true positive rate at low false positive rate) greatly increases the performance of correspondence algorithms.

On the other hand, by emphasizing global features HKS has higher *specificity* (i.e., the ability to identify negatives). This property is related to discriminativity, that is, the ability of the descriptor to distinguish between a shape and other classes of distinct shapes. High discriminativity is important in retrieval applications, and the performance of the descriptor at low false negative rates has a big impact on retrieval algorithms based on it. Both phenomena are visualized in Figure 3. While it is impossible to maximize both the sensitivity and the specificity, a good descriptor is expected to have both reasonably high.

Another drawback of both the heat and wave kernel descriptors is the fact that the frequency responses forming their elements have significant overlaps. As the results, the descriptor has redundant dimensions. Finally, both the heat and wave kernel signatures are only invariant to truly isometric deformations of the shape (and can be also made scale-invariant using the scheme proposed in [27]). Deformations that real shapes undergo frequently deviate from this model, and it is unclear how they influence the performance of the HKS and WKS.

We believe that many real-world deformations affect different frequencies differently. At the same time, the geometric features that allow to localize a point on a shape or to distinguish a shape from other shapes also depend differently on different frequencies. Emphasizing information-carrying frequencies while attenuating noise-carrying ones is a classical idea in signal

and is the underlying principle of Wiener filtering [29].

3.1 Desired properties

This observation leads us to the main contribution of this paper: we propose to construct a collection of frequency responses forming an optimal spectral descriptor. In order to be useful, such a descriptor should satisfy the following properties:

1. *Localization*: a small displacement of a point on the manifold should greatly affect the descriptor computed at it.
2. *Sensitivity*: when a point on a shape is queried against another similar shape, a small set of best matches of the descriptor should contain a correct match with high probability.
3. *Discriminativity*: the descriptor should be able to distinguish between shapes belonging to different classes.
4. *Invariance*: the descriptor should be invariant or at least insensitive to a certain class of transformations that the shape may undergo.
5. *Efficiency*: the descriptor should capture as much information as possible within as little number of dimensions as possible.

The localization and sensitivity properties are important for matching tasks, while in order to be useful in shape retrieval tasks, the descriptor should have the discriminativity property. However, discriminativity is data-dependent: a descriptor can be discriminative on one corpus of data, while non-discriminative on another. While it is generally impractical to model classes of shapes axiomatically, machine learning offers an easy alternative of inferring them from training data.

By construction, spectral descriptors are isometry invariant. However, other invariance properties are usually hard to achieve and even harder to model for realistic transformations. We will therefore stick to learning in order to achieve invariance on examples of transformations the training shapes undergo.

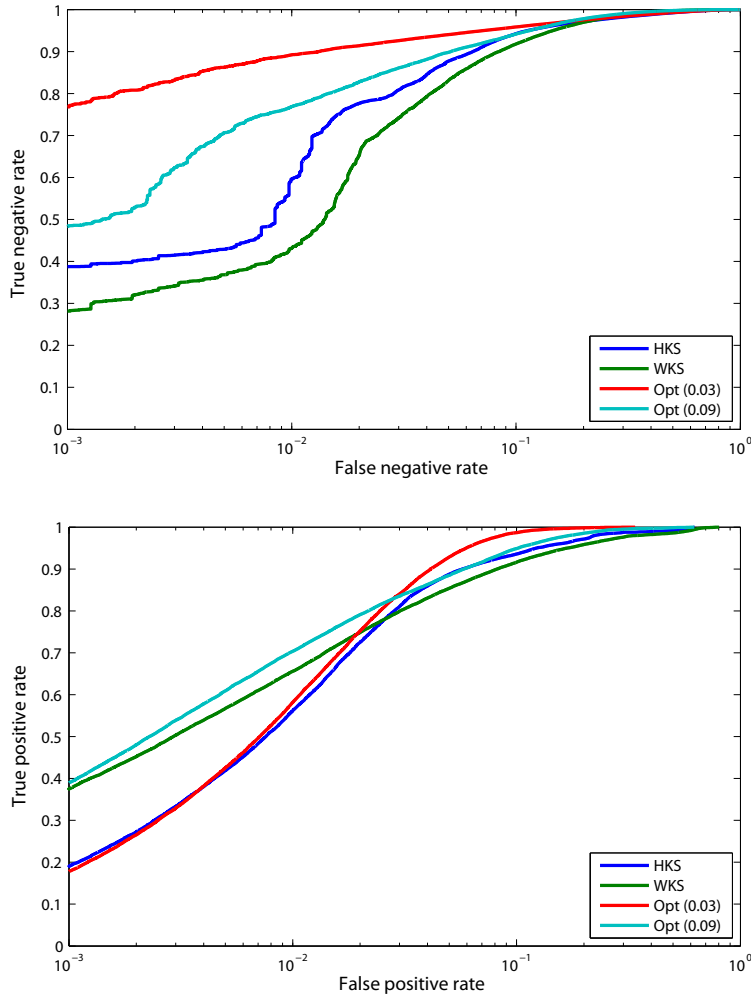


Figure 3: ROC curves of different spectral descriptors when matching points of a shape to itself. A positive match is considered within a geodesic ball of 1% of the shape diameter. Bilaterally symmetric matches are also considered positives. Two regions of the ROC curve are emphasized: the performance of the descriptors for low false negative rate (top), and low false positive rate (bottom). The former case is important to be able to discriminate between different shapes in shape retrieval applications, while the latter is required for establishing an accurate correspondence.

3.2 Parametrization

We are interested in descriptors of the form

$$\mathbf{p}(x) = \sum_{k \geq 1} \mathbf{f}(\nu_k) \phi_k^2(x), \quad (13)$$

parameterized by a vector $\mathbf{f}(\nu) = (f_1(\nu), \dots, f_n(\nu))^T$ of frequency responses. Both the HKS and the WKS are particular cases of this general form. Unlike both heat and wave kernels that are strictly positive, we will allow $\mathbf{f}(\nu)$ assume negative values.

Since the responses $\mathbf{f}(\nu)$ are the design variables of the descriptor, they have to be parametrized with a finite set of parameters. The same parameters have to be compatible with any shape, even though different shapes differ in the set of eigenvalues $\{\nu_k\}$. In order to make the representation independent of a specific shape's eigenvalues, we fix a basis $\{b_1(\nu), \dots, b_m(\nu)\}$, $m > n$, spanning a sufficiently wide interval $[0, \nu_{\max}]$ of frequencies. This allows to express $\mathbf{f}(\nu)$ as

$$\mathbf{f}(\nu) = \mathbf{A}\mathbf{b}(\nu), \quad (14)$$

where \mathbf{A} is the $n \times m$ matrix of coefficients representing the response using the basis functions $\mathbf{b}(\nu) = (b_1(\nu), \dots, b_m(\nu))^T$.

Since the eigenvalues ν_k form a growing progression, we can truncate the series (13) at $\nu_s \geq \nu_{\max}$. Substituting the representation (14), we obtain

$$\mathbf{p}(x) = \mathbf{A}(\mathbf{b}(\nu_1), \dots, \mathbf{b}(\nu_s)) \begin{pmatrix} \phi_1^2(x) \\ \vdots \\ \phi_s^2(x) \end{pmatrix} = \mathbf{A}\mathbf{g}(x) \quad (15)$$

where the $m \times 1$ vector $\mathbf{g}(x)$ with the elements

$$g_j(x) = \sum_{k \geq 1} b_j(\nu_k) \phi_k^2(x) \quad (16)$$

captures all the shape-specific geometric information about the point x . For this reason, we refer to \mathbf{g} as to the *geometry vector* of a point. Note that this representation is no more depends on a specific shape; the matrix of parameters \mathbf{A} describes the same vector of frequency responses on any shape.

3.3 Learning

Let $\mathbf{g} = \mathbf{g}(x)$ be the geometry vector representing some point x ; let $\mathbf{g}_+ = \mathbf{g}(x_+)$ be another geometry vector representing a point that is knowingly similar to x (positive); and, finally, let $\mathbf{g}_- = \mathbf{g}(x_-)$ represent a knowingly dissimilar point (negative). We would like to select the matrix of parameters that maximizes the similarity of the descriptors $\mathbf{p} = \mathbf{A}\mathbf{g}$ and $\mathbf{p}_+ = \mathbf{A}\mathbf{g}_+$, and at the same time minimizes the similarity between \mathbf{p} and $\mathbf{p}_- = \mathbf{A}\mathbf{g}_-$. Using the L^2 norm as the similarity criterion, we obtain

$$\begin{aligned} d_{\pm}^2 &= \|\mathbf{p} - \mathbf{p}_{\pm}\|^2 = \|\mathbf{A}(\mathbf{g} - \mathbf{g}_{\pm})\|^2 \\ &= (\mathbf{g} - \mathbf{g}_{\pm})^T \mathbf{A}^T \mathbf{A} (\mathbf{g} - \mathbf{g}_{\pm}). \end{aligned} \quad (17)$$

In other words, the Euclidean distance between the descriptors translates into a Mahalanobis distance between the corresponding geometry vectors. The problem of finding the best positive-definite matrix $\mathbf{A}^T \mathbf{A}$ defining the Mahalanobis metric is known as *metric learning* and has been relatively well explored in the literature [30, 31, 32].

Here, we describe a simple yet efficient learning scheme explicitly addressing the desired properties we required from a good spectral descriptor. We aim at finding a matrix \mathbf{A} minimizing the Mahalanobis distance over the set of positive pairs, while maximizing it over the negative ones. Note that the distance depends only on the differences between positive and negative pairs of vectors. Taking expectation over all positive and negative pairs, we obtain [33]

$$\begin{aligned} \mathbb{E}(d_{\pm}^2) &= \mathbb{E}(\|\mathbf{p} - \mathbf{p}_{\pm}\|^2) = \mathbb{E}(\mathbf{e}_{\pm}^T \mathbf{A}^T \mathbf{A} \mathbf{e}_{\pm}) \\ &= \text{tr}(\mathbf{A} \mathbb{E}(\mathbf{e}_{\pm} \mathbf{e}_{\pm}^T) \mathbf{A}^T) = \text{tr}(\mathbf{A} \mathbf{C}_{\pm} \mathbf{A}^T), \end{aligned} \quad (18)$$

where $\mathbf{e}_{\pm} = \mathbf{g} - \mathbf{g}_{\pm}$, and \mathbf{C}_{\pm} stands for the covariance matrix of the differences of positive and negative pairs of geometry vectors. In practice, the expectations are replaced by averages over a representative set of difference vectors.

Our goal is to minimize $\mathbb{E}(d_-^2)$ simultaneously maximizing $\mathbb{E}(d_+^2)$. This can be achieved by minimizing the ratio $\mathbb{E}(d_-^2)/\mathbb{E}(d_+^2)$, which is solved by

linear discriminant analysis (LDA). However, we unfavor this approach as it does not allow control over the tradeoff between sensitivity and specificity. Instead, we propose to minimize the difference

$$(1 - \alpha)\mathbb{E}(d_+^2) - \alpha\mathbb{E}(d_-^2) = \text{tr}(\mathbf{A}((1 - \alpha)\mathbf{C}_+ - \alpha\mathbf{C}_-)\mathbf{A}^T) = \text{tr}(\mathbf{A}\mathbf{D}_\alpha\mathbf{A}^T), \quad (19)$$

where $0 \leq \alpha \leq 1$ controls the said tradeoff, and \mathbf{D}_α denotes the difference between the positive and the negative covariance matrices.

Note that since the scale of \mathbf{A} is arbitrary, a trivial solution can be obtained. Even when fixing the scale, the solution will be a rank-1 matrix corresponding to the smallest eigenvector of \mathbf{D}_α . While this can be avoided by arbitrarily demanding orthonormality of \mathbf{A} , such a remedy is completely artificial.

Instead, we remind that one of the desired properties of a descriptor was *efficiency*. In an efficient descriptor, each dimension should be statistically independent of the others. Replacing statistical independence by the more tractable lack of correlation, we demand

$$\mathbf{I} = \mathbb{E}(\mathbf{p}\mathbf{p}^T) = \mathbf{A}\mathbb{E}(\mathbf{g}\mathbf{g}^T)\mathbf{A}^T = \mathbf{A}\mathbf{C}\mathbf{A}^T \quad (20)$$

where expectations are taken over all geometry vectors, and \mathbf{C} denotes the covariance matrix of \mathbf{g} .

Combining (19) with (20), we obtain the following minimization problem

$$\min_{\mathbf{A}} \text{tr}(\mathbf{A}\mathbf{D}_\alpha\mathbf{A}^T) \quad \text{s.t.} \quad \mathbf{A}\mathbf{C}\mathbf{A}^T = \mathbf{I}, \quad (21)$$

which we solve for an $n \times m$ matrix \mathbf{A} . The problem has a closed-form algebraic solution, which is easy to derive using variable substitution. Since \mathbf{C} is a positive-definite matrix, we can substitute $\mathbf{B} = \mathbf{A}\mathbf{C}^{1/2}$, obtaining an equivalent minimization problem

$$\min_{\mathbf{B}} \text{tr}(\mathbf{B}\mathbf{C}^{-1/2}\mathbf{D}_\alpha\mathbf{C}^{-1/2}\mathbf{B}^T) \quad \text{s.t.} \quad \mathbf{B}\mathbf{B}^T = \mathbf{I}, \quad (22)$$

(\mathbf{C} is symmetric and so is its root; we therefore keep writing $\mathbf{C}^{-1/2}$ instead of its transpose). Let us denote by $\mathbf{C}^{-1/2}\mathbf{D}_\alpha\mathbf{C}^{-1/2} = \mathbf{U}\mathbf{A}\mathbf{U}^T$ the

eigendecomposition of the scaled covariance difference, with the eigenvalues $\mathbf{\Lambda} = \text{diag}(\lambda_1, \dots, \lambda_m)$ sorted in ascending order, and the corresponding orthonormal eigenvectors $\mathbf{U} = (\mathbf{u}_1, \dots, \mathbf{u}_m)$. The solution to (22) is given by the first n smallest eigenvectors, $\mathbf{B} = \mathbf{U}_n^T = (\mathbf{u}_1, \dots, \mathbf{u}_n)^T$. Note that one must ensure that all the eigenvectors correspond to negative eigenvalues; if this is not the case, n has to be reduced. Finally, the solution to our original problem (21) follows straightforwardly as

$$\mathbf{A} = \mathbf{U}_n^T \mathbf{C}^{-1/2}. \quad (23)$$

3.4 Training set

So far, we have described a learning scheme allowing to construct efficient spectral descriptors with uncorrelated elements based on covariances of geometry vectors describing positive and negative pairs of points. Having no practical possibility to model the statistics of these vectors, their covariance matrices have to be computed empirically from a training set of positive and negative examples. The construction of such a set is therefore crucial for obtaining a good descriptor. In what follows, we describe how to construct the training set in order to achieve each of the desired properties mentioned before.

Localization. Let x be a point on a training shape X . We fix a pair of radii $r < R$ and deem all points $x_+ \in B_r(x)$ positive, while deeming negative all $x_- \in B_R^c(x)$. Here, $B_r(x)$ denotes the geodesic metric ball of radius r centered at x . Points lying in the ring $B_R(x) \setminus B_r(x)$ are excluded from both sets. If the shape possesses an intrinsic symmetry $\varphi : X \rightarrow X$, then $B_r(\varphi(x))$ is also included in the positive set, while $B_R(\varphi(x))$ is excluded from the negative set. The training set is created by sampling many reference points and corresponding positive and negative points on a collection of representative shapes. The selection of r and R gives explicit control over the localization capability of the descriptor.

Discriminativity. Let X and X_- be knowingly dissimilar shapes (i.e., belonging to different classes we would like to tell apart). A random point x on X and a random point x_- on X_- are deemed negative. The training set

is created by sampling many random pairs of points on knowingly dissimilar pairs of shapes.

Invariance. Let X be a shape and X_+ its transformation belonging to a class of transformations invariance under which is desired. We further assume to be given a correspondence $\varphi : X \rightarrow X_+$ between the shapes. A random point x on X and the corresponding point $x_+ = \varphi(x)$ on X_+ are deemed positive. The training set is created by sampling many random points on a collection of null (reference) shapes, paired with corresponding points on the transformed versions of the null shape.

The combination of the positive and negative sets constructed this way allows to train for descriptor localization, discriminativity, and invariance properties.

3.5 Sensitivity-Specificity tradeoff

The proposed learning scheme allows simple control over the tradeoff between the sensitivity and the specificity of the descriptor through the parameter α . The bigger is α , the bigger is the relative influence of \mathbf{C}_- compared to \mathbf{C}_+ . Therefore, for large values of α , the descriptor will emphasize producing large distances on the negative set (low false positive rate), while trying to keep small distances on the positive set (high true positive rate). As the result, high sensitivity is obtained. For small values of α , the converse is observed: the descriptor emphasizes performance on the positive set, resulting in higher specificity.

In order to select the optimal α for a highly-sensitive descriptor, we empirically compute the false negative rate at some small fixed false positive rate (e.g., 1% or 0.1%) and select the α for which it is minimized. For highly-specific descriptors, α is selected to minimize the false positive rate at some small false negative rate. The behavior of the error rates as a function of α is illustrated in Figure 4.




Figure 4: Error rates as a function of the parameter α . Large values of α result in high sensitivity, while for small values high specificity is obtained.

4 Experimental results

The reported experiments were performed on the SHREC’10 robust correspondence benchmark [21]. The benchmark contains three distinct shape classes (human, dog, and horse), each shape undergoing ten different transformations (isometry, topology, sampling, global scaling, local scaling, holes, micro holes, Gaussian noise, and shot noise) with five strengths per transformation (from mild to very strong). Shapes are represented as triangular meshes with about 5×10^4 vertices (except for the sampling transformations, where the meshes are progressively decimated down to about 2.5×10^3 vertices). The benchmark also contains vertex-wise correspondences between the transformed shapes and the reference (null) shapes, including intrinsic bilateral symmetries. In all experiments, training was performed on the isometry, topology, and Gaussian noise transformations of the horse shape. As the negatives, we used five distinct meshes not included in the benchmark. For evaluation, we used the isometry, topology, holes, Gaussian noise and sampling transformations of the human shape, and the dog shape as the negative. All transformation strengths were used both for training and testing.

We used the finite elements scheme [25] to compute the first 300 eigenvalues and eigenvectors of the Laplace-Beltrami operator on each shape. Neumann boundary conditions were used. The range of frequencies ν_{\max} was set to the 95-percentile of ν_{300} over the entire set of training shapes. The interval was evenly divided into $m = 150$ segments and the cubic spline basis was used as $\{b_j(\nu)\}$. The training set containing 2.5×10^6 150-dimensional triplets of the form $(\mathbf{g}, \mathbf{g}_+, \mathbf{g}_-)$ was generated as described in Section 3.4 with 10^4 negative examples per reference point. The radii r and R were set to 2% and 5% of the shape intrinsic diameter, respectively. The parameter α was selected as described in Section 3.5. The values maximizing the descriptor specificity and sensitivity were found to be $\alpha = 0.03$ and 0.09, respectively (Figure 4). Two corresponding 12-dimensional descriptors were trained. Examples of the obtained responses are shown in Figure 1 (bottom).

4.1 Descriptor performance

Descriptor performance was tested on a distinct set of 2.5×10^6 triplets of points constructed in the same way as the training set but on different shapes. For comparison, we also computed twelve-dimensional HKS and WKS descriptors. The HKS time scales were optimized according to [4]. The WKS energy levels and the variance σ^2 were set as described in [20]. For the fairness of comparison, Euclidean distance was used for all descriptors. Figure 3 shows the ROC curves of the compared descriptors in the low false positive and low false negative work points. As argued before, the HKS is characterized by better performance over the WKS at low false negative rates, while the WKS outperforms the HKS in the low positive rates range. The trained descriptors significantly outperform both the HKS and the WKS in the low false negative rates range, with almost a 40% increase in the true negative rate at $FN = 0.1\%$. The trained high-sensitivity descriptor outperforms WKS by about 6% true positive rate at $FP = 1\%$. The improvement becomes more modest at $FN = 0.1\%$.

4.2 Localization

In order to visualize the localization capability of different descriptors, a reference point was selected on the human shape. The distance between the descriptor at that point was computed to the rest of the points on that shape, to the points of an approximate isometry of the human shape, and to the points on the dog shape. Figure 2 visualizes these normalized distances on a common scale. We observe poor localization capabilities of the HKS along with exceptional localization power of the WKS. The trained high-sensitivity descriptor exhibits even better localization. Both the HKS and the WKS confuse between the reference point on the man’s foot and a region on his hand fingers, which have similar geometric content. On the other hand, our descriptor does not make this confusion. We remind that in the training set, for every reference point all points except its small neighborhood were included as negatives. Even though a different shape was used during the training, the descriptor still seems to be capable of




Figure 5: Normalized Euclidean distance between the descriptor at a reference point on the right hand (white dot) and descriptors computed at rest of the points of the same shape for a twelve-dimensional trained optimal descriptors. Left-to-right: holes, Gaussian noise, and sampling transformations from the SHREC’10 benchmark.

generalizing these relationships.

Finally, both the HKS and the WKS find many points on the dog shape that resemble the reference point on the man’s foot. Our descriptor does not make this confusion as it was trained for discriminativity with numerous negative examples from distinct shapes. Figure 5 shows additional examples of distances computed on other transformations of the human shape using the trained descriptor. In all cases, good localization is observed.

4.3 Correspondence

While evaluation of a particular descriptor-based correspondence algorithms is beyond the scope of this paper, in order to test the performance of the trained high-selectivity descriptor in shape matching tasks, we performed an experiment similar to [20]. 1000 reference points were sampled on the human shape using farthest point sampling in the descriptor space. Such points coincided well with visually “interesting” features. Each reference point was matched to all the points on the transformed versions of the

shape. We computed the probability of finding the correct match (including the symmetric one) within the first k best matches. The CMC curve in Figure 6 depicts the hit rate of different descriptors for up to about first 500 matches (corresponding to 1% of the total points on the shape). The trained descriptor significantly outperforms both the HKS and the WKS. In fact, our descriptor returns the first correct match with over 50% probability, compared to about 25% and 30% in the case of HKS and WKS, respectively.

While the WKS consistently outperforms the HKS on this matching task, we did not notice the dramatic difference reported in [20]. A possible explanation can be the fact that we used only 12 dimensions, while the authors of [20] used a higher-dimensional descriptor. Another, more probable, reason is the fact that in all our experiments Euclidean distance was used as the dissimilarity between the descriptors, while in [20] the authors used WKS with the normalized L^1 distance. We defer to future studies the treatment of distances other than L^2 ; however, we believe that for the fairness of comparison the same distance must be used for all descriptors.

5 Conclusion

We presented a generic framework for the construction of feature descriptors for deformable shapes based on their spectral properties. The proposed descriptor is computed by applying a bank of “filters” to the shape’s geometric features at different “frequencies”, and it generalizes the heat and wave kernel signatures. We also showed a learning approach allowing to construct optimal filters for specific shape analysis tasks, resembling in its spirit optimal signal filtering by means of a Wiener filter.

We formulated the learning approach in terms of the L^2 distance and related it to Mahalanobis metric learning. While the adopted algebraic solution gave good results, other Mahalanobis metric learning approaches, such as the maximum-margin learning [31] can be readily used. Some of these metric learning approaches were designed with a specific task in mind (e.g., ranking), and might be beneficial for the construction of spectral descriptors in some applications. Evidence shows that distances other than the




Figure 6: CMC curve showing the percentage of correct correspondences found in a subset of the first best matches (up to 1% of total points) using different spectral descriptors.

Euclidean one (e.g., the L^1 distance) improve the performance of spectral descriptors. Also, applications where compact and easily searchable descriptors are of importance may benefit from hash learning techniques [34], essentially based on the Hamming distance. We intend to explore alternative learning frameworks and different distances in follow-up studies.

While the main focus of this paper was the construction of the descriptor itself, in future studies we are going to explore its performance in real shape retrieval and matching tasks. Particularly, in retrieval tasks spectral feature descriptors are used to generate global shape descriptors by means of vector quantization or sparse coding, a growingly popular alternative in the computer vision community. Taking this highly non-linear process into account when constructing the feature descriptor will also be a subject of our future research.

References

- [1] N. Gelfand, N. J. Mitra, L. J. Guibas, and H. Pottmann, “Robust global registration,” in *Proc. SGP*, 2005.
- [2] C. Wang, A. M. Bronstein, M. M. Bronstein, and N. Paragios, “Discrete minimum distortion correspondence problems for non-rigid shape matching,” in *Proc. Scale Space and Variational Methods (SSVM)*, 2011.
- [3] N. J. Mitra, L. J. Guibas, J. Giesen, and M. Pauly, “Probabilistic fingerprints for shapes,” in *Proc. SGP*, 2006.
- [4] A. Bronstein, M. Bronstein, L. Guibas, and M. Ovsjanikov, “Shape google: geometric words and expressions for invariant shape retrieval,” *ACM Transactions on Graphics (TOG)*, vol. 30, no. 1, p. 1, 2011.
- [5] P. Skraba, M. Ovsjanikov, F. Chazal, and L. Guibas, “Persistence-based segmentation of deformable shapes,” in *Proc. NORDIA*, 2010, pp. 45–52.

- [6] S. Belongie, J. Malik, and J. Puzicha, “Shape context: A new descriptor for shape matching and object recognition,” in *Proc. NIPS*, 2000.
- [7] A. E. Johnson and M. Hebert, “Using spin images for efficient object recognition in cluttered 3D scenes,” *Trans. PAMI*, vol. 21, no. 5, pp. 433–449, 1999.
- [8] S. Manay, B. Hong, A. Yezzi, and S. Soatto, “Integral invariant signatures,” *Lecture Notes in Computer Science*, pp. 87–99, 2004.
- [9] M. Pauly, R. Keiser, and M. Gross, “Multi-scale feature extraction on point-sampled surfaces,” in *Computer Graphics Forum*, vol. 22, no. 3, 2003, pp. 281–289.
- [10] A. Hamza and H. Krim, “Geodesic object representation and recognition,” in *Discrete Geometry for Computer Imagery*, 2003, pp. 378–387.
- [11] A. Elad and R. Kimmel, “On bending invariant signatures for surfaces,” *IEEE Trans. Pattern Analysis and Machine Intelligence*, pp. 1285–1311, 2003.
- [12] Y. Lipman and T. Funkhouser, “Möbius voting for surface correspondence,” in *ACM Trans. on Graphics*, vol. 28, no. 3, 2009, p. 72.
- [13] P. Bérard, G. Besson, and S. Gallot, “Embedding Riemannian manifolds by their heat kernel,” *Geometric and Functional Analysis*, vol. 4, no. 4, pp. 373–398, 1994.
- [14] R. Coifman and S. Lafon, “Diffusion maps,” *Applied and Computational Harmonic Analysis*, vol. 21, no. 1, pp. 5–30, 2006.
- [15] F. Méholi, “Spectral Gromov-Wasserstein distances for shape matching,” in *Proc. ICCV Workshops*, 2009, pp. 256–263.
- [16] A. Bronstein, M. Bronstein, R. Kimmel, M. Mahmoudi, and G. Sapiro, “A Gromov-Hausdorff framework with diffusion geometry for topologically-robust non-rigid shape matching,” *Int’l Journal of Computer Vision*, vol. 89, no. 2, pp. 266–286, 2010.

- [17] B. Lévy, “Laplace-Beltrami eigenfunctions towards an algorithm that understands geometry,” in *Proc. SMI*, 2006, pp. 13–13.
- [18] R. Rustamov, “Laplace-Beltrami eigenfunctions for deformation invariant shape representation,” in *Proc. Symp. on Geometry Processing (SGP)*, 2007, pp. 225–233.
- [19] J. Sun, M. Ovsjanikov, and L. Guibas, “A Concise and Provably Informative Multi-Scale Signature Based on Heat Diffusion,” in *Computer Graphics Forum*, vol. 28, no. 5, 2009, pp. 1383–1392.
- [20] M. Aubry, U. Schlickewei, and D. Cremers, “The wave kernel signature—a quantum mechanical approach to shape analysis,” in *Proc. CVPR*, 2011.
- [21] A. Bronstein, M. Bronstein, U. Castellani, A. Dubrovina, L. Guibas, R. Horaud, R. Kimmel, D. Knossow, E. von Lavante, D. Mateus *et al.*, “SHREC 2010: robust correspondence benchmark,” 2010.
- [22] A. Bronstein, M. Bronstein, U. Castellani, B. Falcidieno, A. Fusiello, A. Godil, L. Guibas, I. Kokkinos, Z. Lian, M. Ovsjanikov *et al.*, “SHREC 2010: robust large-scale shape retrieval benchmark,” 2010.
- [23] J. Aflalo, A. M. Bronstein, M. M. Bronstein, and R. Kimmel, “Deformable shape retrieval by learning diffusion kernels,” in *Proc. Scale Space and Variational Methods (SSVM)*, 2011.
- [24] M. Kac, “Can one hear the shape of a drum?” *The American Mathematical Monthly*, vol. 73, no. 4, pp. 1–23, 1966.
- [25] M. Reuter, F. Wolter, and N. Peinecke, “Laplace-Beltrami spectra as “Shape-DNA” of surfaces and solids,” *Computer-Aided Design*, vol. 38, no. 4, pp. 342–366, 2006.
- [26] A. Sharma and R. Horaud, “Shape matching based on diffusion embedding and on mutual isometric consistency,” in *Proc. CVPR Workshops*, 2010, pp. 29–36.

- [27] M. M. Bronstein and I. Kokkinos, “Scale-invariant heat kernel signatures for non-rigid shape recognition,” in *Proc. CVPR*, 2010.
- [28] D. Raviv, M. Bronstein, A. Bronstein, and R. Kimmel, “Volumetric heat kernel signatures,” in *Proc. ACM Workshop on 3D Object Retrieval*, 2010, pp. 39–44.
- [29] N. Wiener, *Extrapolation, interpolation, and smoothing of stationary time series*. Wiley.
- [30] L. Yang and R. Jin, “Distance metric learning: A comprehensive survey,” *Michigan State University*, pp. 1–51, 2006.
- [31] K. Weinberger, J. Blitzer, and L. Saul, “Distance metric learning for large margin nearest neighbor classification,” in *Proc. NIPS*, 2006.
- [32] J. Davis, B. Kulis, P. Jain, S. Sra, and I. Dhillon, “Information-theoretic metric learning,” in *Proc. ICML*, 2007, pp. 209–216.
- [33] C. Strecha, A. Bronstein, M. Bronstein, and P. Fua, “LDAHash: Improved matching with smaller descriptors,” *IEEE Trans. Pattern Analysis and Machine Intelligence*, 2011.
- [34] Y. Weiss, A. Torralba, and R. Fergus, “Spectral hashing,” in *Proc. NIPS*, 2008.