

Geometric and photometric data fusion in non-rigid shape analysis

Artiom Kovnatsky³, Dan Raviv³, Michael M. Bronstein^{1,*},
Alexander M. Bronstein², and Ron Kimmel³

¹ *Institute of Computational Science, Faculty of Informatics, Università della Svizzera Italiana, Lugano, Switzerland*

² *School of Electrical Engineering, Tel Aviv University, Israel*

³ *Department of Computer Science, Technion, Israel*

Abstract. In this paper, we explore the use of the diffusion geometry framework for the fusion of geometric and photometric information in local and global shape descriptors. Our construction is based on the definition of a diffusion process on the shape manifold embedded into a high-dimensional space where the embedding coordinates represent the photometric information. Experimental results show that such data fusion is useful in coping with different challenges of shape analysis where pure geometric and pure photometric methods fail.

Key words: Laplace-Beltrami operator, diffusion equation, heat kernel descriptors, 3D shape retrieval, deformation invariance

1. Introduction

The birth of the World Wide Web and the explosive growth of text content has brought the need to organize, index, and search text document, which in turn fueled the development of text search engines. In the past decade, the amount of geometric data available in the public-domain repositories such as Google 3D Warehouse, has grown dramatically and created the demand for shape search and retrieval algorithms capable of finding similar shapes in the same way a search engine responds to text queries. While text search methods are sufficiently developed to be ubiquitously used, the search and retrieval of 3D shapes remains a challenging problem. Shape retrieval based on text metadata, like annotations and tags added by the users, is often incapable of providing relevance level required for a reasonable user experience (see Figure 1).

Content-based shape retrieval using the shape itself as a query and based on the comparison of geometric and topological properties of shapes is complicated by the fact that

*Corresponding author.

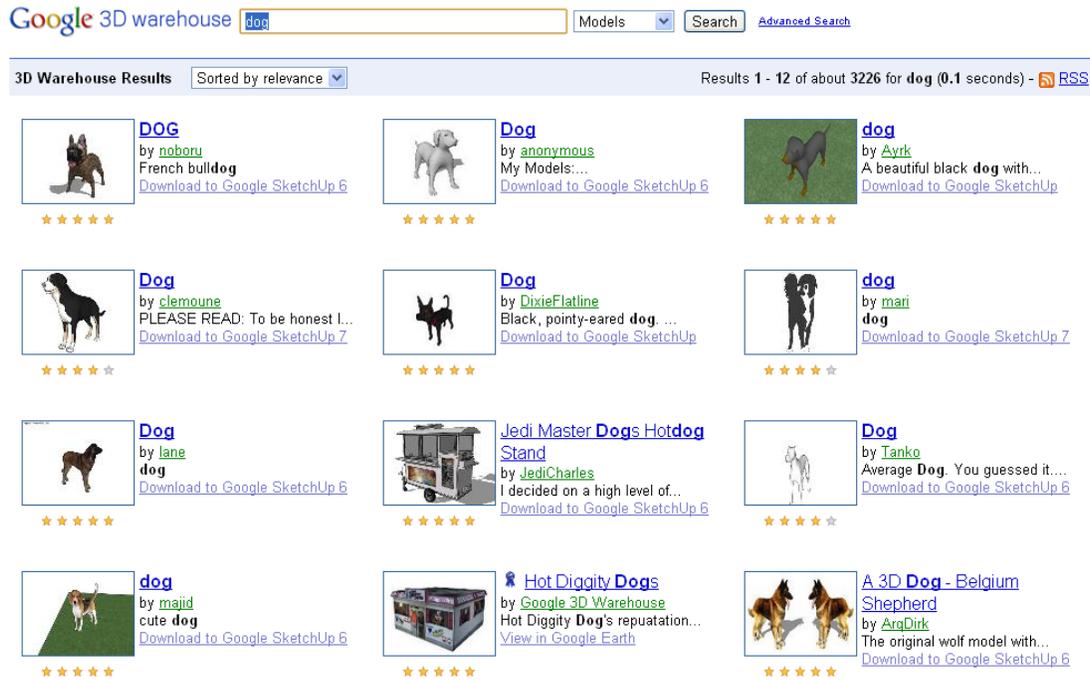


Figure 1: The need for content-based shape retrieval: text-based search engine such as *Google 3D Warehouse* returns shapes of dogs as well as hot-dog cabins in response to the query “dog”. The later is obviously irrelevant.

many 3D objects manifest rich variability, and shape retrieval must often be *invariant* under different classes of transformations. A particularly challenging setting is the case of non-rigid shapes, including a wide range of transformations such as bending and articulated motion, rotation and translation, scaling, non-rigid deformation, and topological changes. The main challenge in shape retrieval algorithms is computing a *shape descriptor*, that would be unique for each shape, simple to compute and store, and invariant under different type of transformations. Shape similarity is determined by comparing the shape descriptors.

1.1. Prior work

Broadly, shape descriptors can be divided into *global* and *local*. The former consider global geometric or topological shape characteristics such as distance distributions [35, 41, 49], geometric moments [28, 59], or spectra [48], whereas the latter describe the local behavior of the shape in a small patch. Popular examples of local descriptors include spin images [3], shape contexts [1], integral volume descriptors [23] and radius-normal histograms [43]. Using the bag of features paradigm common in image analysis [18, 50], a global shape descriptor counting the occurrence of local descriptors in some vocabulary can be computed [15].

Early research on feature descriptors focused mainly on invariance under global Eu-

clidean transformations (rigid motion). However, in the past decade, significant effort has been invested in extending the invariance properties to non-rigid deformations [12,13,38]. Some of the classical rigid descriptors were extended to the non-rigid case by replacing the Euclidean metric with its geodesic counterpart [21,24].

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.* [7] and later by Coifman and Lafon [19] 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 [35,39]. Diffusion geometry offers an intuitive interpretation of many shape properties in terms of spatial 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.

One of the first principled practical uses of these methods in the context of shape processing was explored by Lévy [31]. Several attempts have also been made to construct feature descriptors based on diffusion geometric properties of the shape. Rustamov [49] 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 [19]. 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) related to intrinsic shape symmetry [42,44]. A remedy was proposed by Sun *et al.* who in their influential paper [53] introduced the *heat kernel signature* (HKS), based on the fundamental solutions of the heat equation (heat kernels). Scale invariant [17], affine-invariant [45,46], and volumetric [47] versions of the HKS were subsequently proposed. In [4], another physically-inspired descriptor, the *wave kernel signature* (WKS) was proposed as a remedy to the excessive sensitivity of the HKS to low-frequency information. In [8], a general family of learnable spectral descriptors generalizing the HKS and WKS was introduced.

In [15,57], the use of local diffusion geometric descriptors was proposed to construct global shape descriptors using the bags-of-features paradigm. Global descriptors based on distributions of diffusion metrics were also proposed [16,35]. Spectral constructions were used in [14,39] in the Gromov-Hausdorff framework for intrinsic shape correspondence [12,13,37]. In [42], the properties of the Laplacian eigenfunctions were used for the detection of bilateral intrinsic symmetries in deformable shapes. Diffusion geometry was also exploited for robust intrinsic shape segmentation [51] and the detection of region features [20,33] extending the maximally stable extremal regions popular in image analysis [36]. As of today, diffusion geometric approaches achieve state-of-the-art performance in many deformable shape analysis tasks [9,10].

A major limitation of these methods is that, so far, only *geometric* information has been considered. However, the abundance of textured models in computer graphics and mod-

eling applications, as well as the advance in 3D shape acquisition [64, 65] allowing to obtain textured 3D shapes of even moving objects, bring forth the need for descriptors also taking into consideration *photometric* information. Photometric information plays an important role in a variety of shape analysis applications, such as shape matching and correspondence [56, 62]. Considering 2D views of the 3D shape [40, 61], standard feature detectors and descriptors used in image analysis such as SIFT [34] can be employed. However, such approaches do not take into consideration the invariant geometric properties of the 3D shape. An alternative is to try to create an intrinsic analogy of such descriptors on the surface itself. A recent example of such approach is mesh-HOG, a geometric SIFT-like descriptor for textured shapes defined directly on the surface [66].

1.2. Main contribution

In this paper, we extend the diffusion geometry framework to include photometric information in addition to its geometric counterpart. This way, we incorporate important photometric properties on the one hand, while exploiting a principled and theoretically established approach on the other. The main idea is to define a diffusion process on a manifold in a higher dimensional combined geometric-photometric embedding space, similarly to methods in image processing applications [29, 32]. As a result, we are able to construct local descriptors (heat kernel signatures) and global descriptors (diffusion distance distributions). The proposed data fusion can be useful in coping with different challenges of shape analysis where pure geometric and pure photometric methods fail.

Preliminary results of this study introducing photometric HKS descriptors have been published in [30]. Here, we extend the theoretical part, adding a deeper discussion of the important issues of invariance and showing the usefulness of the affine-invariant diffusion processes. The experimental part has also been extended, including both local and global shape descriptors.

The rest of this paper is organized as follows. In Section 2, we review the mathematical formalism of diffusion processes and their use in shape analysis. In Section 3, we discuss the modified diffusion process, and in Section 4 describe the related local and global shape descriptors. Section 5 addresses numerical implementation details. Section 6 presents experimental results. Finally, Section 7 concludes the paper.

2. Background

Throughout the paper, we assume the shape to be modeled as a two-dimensional compact Riemannian manifold X (possibly with a boundary) equipped with a *metric tensor* g . Fixing a system of local coordinates on X , the latter can be expressed as a 2×2 matrix $g_{\mu\nu}$, also known as the *first fundamental form*. The metric tensor allows to express the length of a vector v in the tangent space $T_x X$ at a point x as $g_{\mu\nu} v^\mu v^\nu$, where repeated indices $\mu, \nu = 1, 2$ are summed over following Einstein's convention.

Given a smooth scalar field $f : X \rightarrow \mathbb{R}$ on the manifold, its *gradient* is defined as the vector field ∇f satisfying $f(x + dx) = f(x) + g_x(\nabla f(x), dx)$ for every point x and every

infinitesimal tangent vector $dx \in T_x X$. The metric tensor g defines the *Laplace-Beltrami operator* Δ_g that satisfies

$$\int f \Delta_g h da = - \int g_x(\nabla f, \nabla h) da \quad (2.1)$$

for any pair of smooth scalar fields $f, h : X \rightarrow \mathbb{R}$; here da denotes integration with respect to the standard area measure on X . Such an integral definition is known as the *Stokes identity*. The Laplace-Beltrami operator is positive semi-definite and self-adjoint. Furthermore, it is an *intrinsic* property of X , i.e., it is expressible solely in terms of g . In the case when the metric g is Euclidean, Δ_g becomes the standard Laplacian.

The Laplace-Beltrami operator gives rise to the *heat equation*,

$$\left(\Delta_g + \frac{\partial}{\partial t} \right) u = 0, \quad (2.2)$$

which describes diffusion processes and heat propagation on the manifold. Here, $u(x, t)$ denotes the distribution of heat at time t at point x . The initial condition to the equation is some heat distribution $u(x, 0)$, and if the manifold has a boundary, appropriate boundary conditions (e.g. Neumann or Dirichlet) must be specified. The solution of (2.2) with a point initial heat distribution $u_0(x) = \delta(x, x')$ is called the *heat kernel* and denoted here by $k_t(x, x')$.

By virtue of the spectral theorem, there exists an orthonormal basis on $L_2(X)$ consisting of the eigenfunctions ϕ_0, ϕ_1, \dots of the Laplace-Beltrami operator (i.e., solutions to $\Delta_g \phi_i = \lambda_i \phi_i$, where $0 = \lambda_0 \leq \lambda_1 \leq \dots$ are the corresponding eigenvalues). This basis can be interpreted analogously to the Fourier basis, and the eigenvalues λ_i as frequencies. Consequently, the heat kernel can be represented as [27]

$$k_t(x, x') = \sum_{i \geq 0} e^{-\lambda_i t} \phi_i(x) \phi_i(x'). \quad (2.3)$$

The value of the heat kernel $k_t(x, x')$ can also be interpreted as the transition probability density of a random walk of length t from the point x to the point x' . This allows to construct a family of intrinsic metrics known as *diffusion metrics*,[†]

$$d_t^2(x, x') = \int (k_t(x, \cdot) - k_t(x', \cdot))^2 da = \sum_{i > 0} e^{-2\lambda_i t} (\phi_i(x) - \phi_i(x'))^2. \quad (2.4)$$

These metrics have an inherent multi-scale structure and measure the “connectivity rate” of the two points by paths of length t . We will collectively refer to quantities expressed in terms of the heat kernel or diffusion metrics as to *diffusion geometry*. Since the Laplace-Beltrami operator is intrinsic, the diffusion geometry it induces is invariant under isometric deformations of X (incongruent embeddings of g into \mathbb{R}^3).

[†]Note that here the term *metric* is understood in the sense of metric geometry rather than the Riemannian inner product. To avoid confusion, we refer to the latter as to *metric tensor* throughout the paper.

3. Fusion of geometric and photometric data

The main idea of this paper is to create a modified diffusion operator that combines geometric and photometric properties of the shape by means of definition of a new metric tensor (and hence the Laplace-Beltrami operator). For this purpose, let us further assume that the Riemannian manifold X is a submanifold of some manifold \mathcal{E} ($\dim(\mathcal{E}) = m > 2$) with the Riemannian metric tensor h , embedded by means of a diffeomorphism $\xi : X \rightarrow \mathcal{E}$. A Riemannian metric tensor on X induced by the embedding is the *pullback metric* $(\xi^*h)(r, s) = h(d\xi(r), d\xi(s))$ for $r, s \in T_x X$, where $d\xi : T_x X \rightarrow T_{\xi(x)} \mathcal{E}$ is the differential of ξ . In coordinate notation, the pullback metric is expressed as $\hat{g}_{\mu\nu} = (\xi^*h)_{\mu\nu} = h_{ij} \partial_\mu \xi^i \partial_\nu \xi^j$, where the indices $i, j = 1, \dots, m$ denote the embedding coordinates.

Here, we use the structure of \mathcal{E} to model joint geometric and photometric information. Such an approach has been successfully used in image processing [29]. When considering shapes as geometric object only, we define $\mathcal{E} = \mathbb{R}^3$ and h to be the Euclidean metric. In this case, ξ acts as a *parametrization* of X and the pullback metric becomes simply $(\xi^*h)_{\mu\nu} = \partial_\mu \xi^1 \partial_\nu \xi^1 + \dots + \partial_\mu \xi^3 \partial_\nu \xi^3 = \langle \partial_\mu \xi, \partial_\nu \xi \rangle_{\mathbb{R}^3}$. In the case considered in this paper, the shape is endowed with photometric information given in the form of a field $\alpha : X \rightarrow \mathcal{C}$, where \mathcal{C} denotes some colorspace (e.g., RGB or Lab). In the following, when required, we tacitly assume that α is sufficiently smooth.

This photometric information can be modeled by defining $\mathcal{E} = \mathbb{R}^3 \times \mathcal{C}$ and an embedding $\xi = (\xi_g, \xi_p)$. The embedding coordinates corresponding to geometric information are as before $\xi_g = (\xi^1, \dots, \xi^3)$, and the embedding coordinate corresponding to photometric information are given by $\xi_p = (\xi^4, \dots, \xi^6) = \eta(\alpha^1, \dots, \alpha^3)$, where $\eta \geq 0$ is a scaling constant. In addition to trading off between geometry and photometry parts, the scaling constant η has another role of resolving ambiguities of new isometry group, as discussed later in Section 3.3.

The Laplace-Beltrami operator $\Delta_{\hat{g}}$ associated with such a metric gives rise to diffusion geometry that combines photometric and geometric information (Figure 2).

3.1. Euclidean color metric

The invariance to different classes of photometric transformations is obtained by selecting the structure of the colorspace \mathcal{C} . In the simplest case, we assume \mathcal{C} to have a Euclidean structure with the metric tensor

$$(\hat{g}_{\mu\nu}) = \begin{pmatrix} I_g & 0 \\ 0 & \eta I_p \end{pmatrix} \quad (3.1)$$

where $I_g = I_p$ are identity matrices of size 3×3 .

While being the simplest choice, the Euclidean metric is known to be perceptually meaningful in some colorspace such as the ‘‘color opponent’’ CIE Lab space intended to mimic the nonlinear response of the eye [26]. The photometric coordinates $\xi_p = (L, a, b)$ in this colorspace represent *lightness* and color differences: a varies from green to red, and

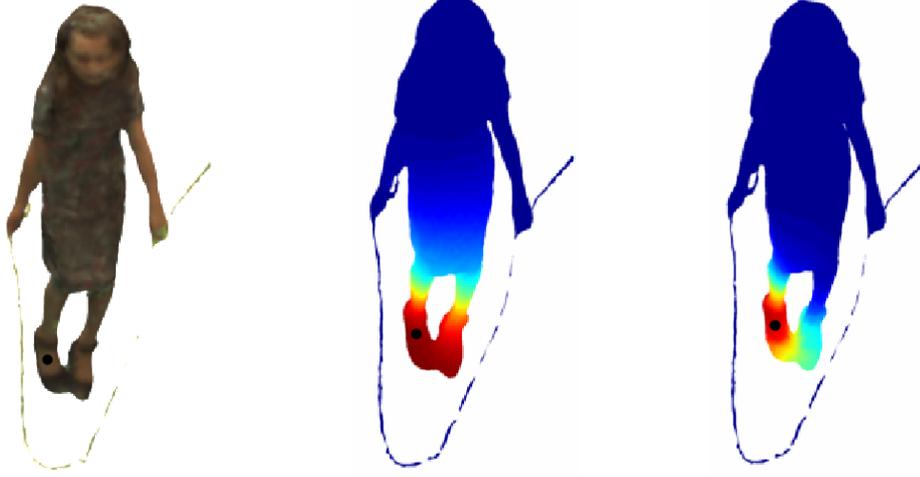


Figure 2: Textured shape (left); values of the heat kernel (x is the black dot marked on the foot, $t = 1024$) arising from regular purely geometric (middle) and mixed photometric-geometric (right) diffusion process.

b varies from blue to yellow. Isometries with respect to the Euclidean metric in the Lab colorspace are shifts (resulting in lightning and hue transformations) and rotations,

$$\xi_p = R\xi_p + c, \quad (3.2)$$

where R denotes a 3×3 rotation matrix, and c is a 3×1 shift vector. Such transformations capture many natural color changes the shape can undergo (Figure 3).

The joint metric in this case boils down to $(\xi^*g)_{\mu\nu} = \langle \partial_\mu \xi_g, \partial_\nu \xi_g \rangle_{\mathbb{R}^3} + \eta \langle \partial_\mu \xi_p, \partial_\nu \xi_p \rangle_{\mathbb{R}^3}$.

3.2. Affine-invariant color metric

A more generic class of photometric transformations can be expressed as *affine transformations* in the Lab colorspace $\xi_p = A\xi_p + c$, where A is an invertible 3×3 matrix. In particular, the transformation is called *equi-affine* if $\det(A) = 1$.

Raviv *et al.* [45, 46] showed a construction of a metric that is invariant to equi-affine transformations. In our setting, let us be given some parametrization $\phi(u_1, u_2) : U \subseteq \mathbb{R}^2 \rightarrow X \subset \mathbb{R}^3$ of the shape; the composition of $\alpha \circ \phi$ gives us a parametrization of the texture. First, allowing some relaxed notation, we denote by $g_X(u_1, u_2) = (\langle \partial_{u_\mu}, \partial_{u_\nu} \rangle)$ and $g_{\alpha(X)}(u_1, u_2) = (\langle d\alpha(\partial_{u_\mu}), d\alpha(\partial_{u_\nu}) \rangle)$ the first fundamental forms of X and $\alpha(X)$, respectively in matrix representations at point $\phi(u_1, u_2)$ in our parametrization. Here, $d\alpha$ is the differential of α and $d\alpha(\partial_{u_\mu}) = \frac{\partial}{\partial u_\mu}(\alpha \circ \phi)$. Second, construct an equi-affine *pre-metric tensor* [45, 52]

$$(\bar{g}_M(u_1, u_2))_{\mu\nu} = \tilde{g}_{\mu\nu} \det^{-1/4}(\tilde{g}), \quad (3.3)$$

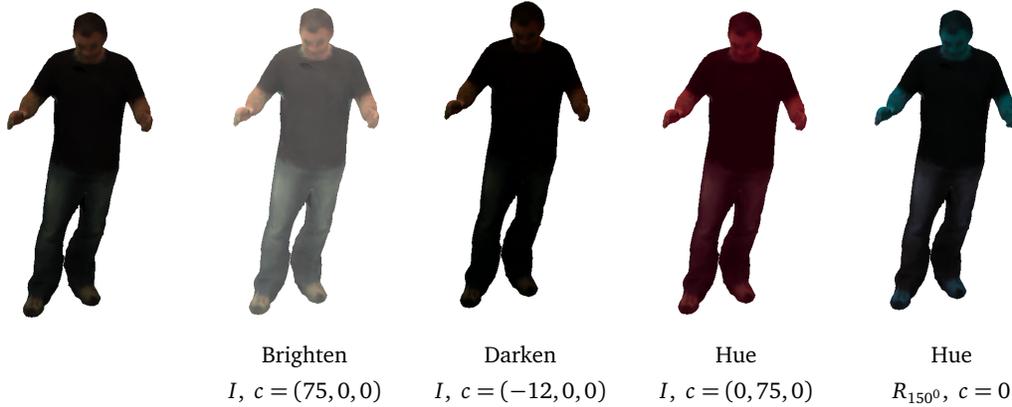


Figure 3: Illustration of photometric isometries in Lab color space equipped with the Euclidean metric corresponding to different choice of R and c .

where $\tilde{g}_{\mu\nu} = \det(d\alpha(\partial_{u_1}), d\alpha(\partial_{u_2}), (\alpha \circ \phi)_{u_\mu u_\nu})$. Note that this is not yet a metric since it might have negative eigenvalues on surfaces which are not strictly convex. Finally, the metric tensor is obtained by forcing \tilde{g} to have positive eigenvalues, $g_{\alpha(X)}(u_1, u_2) = V|\Gamma|V^T$, where V is the matrix of eigenvectors of $(\tilde{g}_{\mu\nu})$ and $|\Gamma|$ is the diagonal matrix of corresponding absolute eigenvalues. Such a normalization tacitly implies that the Gaussian curvature is non-vanishing at every point of the shape, otherwise the metric tensor is not positive-definite.

The modified metric tensor with the equi-affine-invariant photometric component is defined in matrix representation with respect to (u_1, u_2) on X as

$$\hat{g}(u_1, u_2) = g_X(u_1, u_2) + \eta g_{\alpha(X)}(u_1, u_2). \quad (3.4)$$

For additional details about derivation and proof of affine invariance, we refer the reader to [45, 46].

Figure 4 shows a comparison of the heat kernels resulting from the use of three different metrics (Euclidean and equi-affine-invariant) on the colorspace.

3.3. Invariance of the joint diffusion process

The joint metric tensor \hat{g} and the diffusion geometry it induces have inherent ambiguities. Let us denote by $\text{Iso}_g = \text{Iso}((\xi_g^* h)_{\mu\nu})$ and $\text{Iso}_p = \text{Iso}((\xi_p^* h)_{\mu\nu})$ the respective groups of transformation that leave the geometric and the photometric components of the shape unchanged. We will refer to such transformations as geometric and photometric isometries. The diffusion metric induced by \hat{g} is invariant the joint isometry group $\text{Iso}_{\hat{g}} = \text{Iso}((\xi_g^* h)_{\mu\nu} + \eta^2 (\xi_p^* h)_{\mu\nu})$. Ideally, we would like $\text{Iso}_{\hat{g}} = \text{Iso}_g \times \text{Iso}_p$ to hold. In practice, $\text{Iso}_{\hat{g}}$ is bigger: while every composition of a geometric isometry with a photometric isometry is a joint isometry, there exist some joint isometries which cannot be obtained as a composition of geometric and photometric isometries.

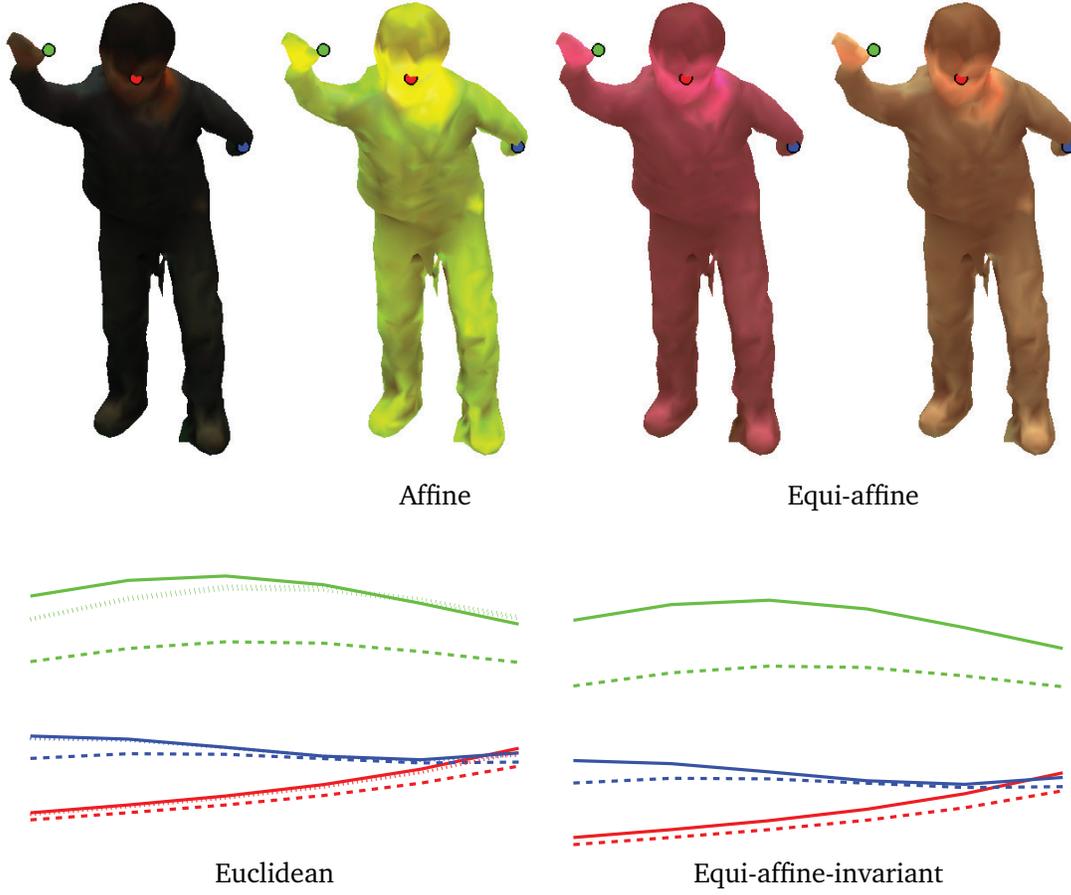


Figure 4: The influence of metric selection. First row: three photometric transformations of a textured shape. Second and third rows: the value of the heat kernel diagonal $h_t(x, x)$ at three different points (marked with red, green, and blue on the shape) for different transformations (solid line: null, dashed: affine, dotted: equi-affine). The heat kernel is constructed using Euclidean and equi-affine-invariant metrics in the colorspace. The more the curves coincide, the better is invariance.

An example of such transformations is uniform scaling of $(\xi_g^* h)_{\mu\nu}$ combined with compensating scaling of $(\xi_p^* h)_{\mu\nu}$ (Figure 5). The ambiguity stems from the fact that $\text{Iso}_{\hat{g}}$ is bigger compared to $\text{Iso}_g \times \text{Iso}_p$. Experimental results show that no realistic geometric and photometric transformations lie in $\text{Iso}_{\hat{g}} \setminus (\text{Iso}_g \times \text{Iso}_p)$, however, a formal characterization of the isometry group is an important theoretical question for future research.

It is possible to overcome the ambiguity problem by considering metrics with different values of the scaling factor η . This rules out the compensating scaling situation depicted in Figure 5, and ensures that the shapes appear isometric for all values of η only if their geometric and photometric components are isometric.

The joint metric appears less sensitive to shape connectivity changes and topological noise commonly seen in objects acquired by means of a 3D scanner. An intuitive explana-

tion is shown in Figure 2. Purely geometric diffusion allows heat transfer between connected points, and thus topological noise can greatly affect the diffusion process. On the other hand, joint diffusion also takes into consideration the photometric information, thus allowing heat transfer between connected points with similar color.

4. Shape descriptors

Using the modified diffusion process, it is possible to create diffusion-based shape descriptors that combine both geometric and photometric data and thus are suitable for textured shape comparison. Specifically, we describe shape representations based on local diffusion feature descriptors (heat kernel signatures) and global diffusion metric structures (shape distributions).

4.1. Bags of local features

Sun *et al.* [54] and independently Gebal *et al.* [22] proposed using the heat propagation properties as a local descriptor of the manifold. The diagonal of the heat kernel,

$$k_t(x, x) = \sum_{i \geq 0} e^{-\lambda_i t} \phi_i^2(x), \quad (4.1)$$

referred to as the *heat kernel signature* (HKS), captures the local properties of X at point x and scale t . The descriptor is computed at each point as a vector of the values $p(x) = (k_{t_1}(x, x), \dots, k_{t_n}(x, x))$, where t_1, \dots, t_n are some time values. The resulting n -dimensional descriptor is deformation-invariant, easy to compute, and provably informative[‡] [54].

Ovsjanikov *et al.* [15] employed the HKS local descriptor for large-scale shape retrieval using the *bags of features* paradigm [50]. In this approach, the shape is considered as a collection of “geometric words” from a fixed “vocabulary” $\{p_1, \dots, p_q\} \subset \mathbb{R}^n$ and is described by the distribution of such words, also referred to as a *bag of features* or BOF. The vocabulary is constructed offline by clustering the HKS descriptor space. Then, for each point x on the shape, the HKS $p(x)$ is replaced by the nearest vocabulary word by means of vector quantization,

$$\theta(x) = (\theta_1(x), \dots, \theta_q(x)) = \frac{e^{-\|p(x)-p_i\|_2^2/2\sigma^2}}{\sum_{i=1}^q e^{-\|p(x)-p_i\|_2^2/2\sigma^2}}, \quad (4.2)$$

where σ controls the “softness” of quantization. The BOF is constructed as

$$b_X = \int_X \theta(x) da \quad (4.3)$$

[‡]Sun *et al.* showed that for shapes possessing simple spectra the sequence of functions $\{k_t(x, x)\}_{t>0}$ completely describes the metric tensor.

and can be considered as the frequency of different geometric words. The similarity of two shapes X and Y is then computed as the distance between the corresponding BOFs, $d(X, Y) = \|b_X - b_Y\|$.

Using the proposed approach, we define the *color heat kernel signature* (cHKS), defined in the same way as HKS with the standard Laplace-Beltrami operator replaced by the one resulting from the geometric-photometric embedding. The photometric information is represented in the Lab colorspace with the Euclidean or affine-invariant metric.

As discussed in Section 3.3, in order to avoid ambiguities related to the joint metric, we have to compute the cHKS descriptor with multiple values of the scaling parameter η , each value producing a different set of cHKS descriptors $p_\eta(x)$ and corresponding bags of features $b_{X,\eta}$. This set of BOFs can be compared e.g. as

$$d(X, Y) = \sum_{\eta \in H} \eta \|b_{X,\eta} - b_{Y,\eta}\|. \quad (4.4)$$

4.2. Shape distributions

Recent works [16, 35, 49] showed that global shape descriptors can be constructed considering *distributions* of intrinsic distances. Given some intrinsic distance metric $d_X : X \times X \rightarrow \mathbb{R}_+ \cup \{0\}$, its cumulative distribution is computed as

$$F_X(\delta) = \int \chi_{d_X(x,x') \leq \delta} da(x) da(x'), \quad (4.5)$$

where χ denotes an indicator function. In particular, the choice of diffusion metrics (2.4) as d_X was discussed in depth in [16]. Given two shapes X and Y with the corresponding distance metrics d_X, d_Y , the similarity (referred to as *spectral distance*) is computed as a distance between the corresponding distributions F_X and F_Y .

Following the main idea of this paper, we use joint geometric-photometric diffusion processes to generate families of diffusion metrics of the form $d_{t,\eta}$ for different values of the scaling factor η according to (2.4). To remove the joint metric ambiguity, we compute a product of distances

$$d_t(x, x') = \prod_{\eta \in H} d_{t,\eta}(x, x'). \quad (4.6)$$

where H is some set of scaling factor values.

5. Numerical implementation

Let $\{x_1, \dots, x_N\} \subseteq X$ denote the discrete samples of the shape, and $\xi(x_1), \dots, \xi(x_N)$ be the corresponding embedding coordinates (three-dimensional in the case we consider only geometry, or six-dimensional in the case of geometry-photometry fusion). We further assume to be given a *triangulation* (simplicial complex), consisting of *edges* (i, j) and *faces* (i, j, k) where each (i, j) , (j, k) , and (i, k) is an edge (here $i, j, k = 1, \dots, N$).

5.1. Discrete Laplacian

A function f on the discretized manifold is represented as an N -dimensional vector $(f(x_1), \dots, f(x_N))$. The discrete Laplace-Beltrami operator can be written in the generic form

$$(\hat{\Delta}f)(x_i) = \frac{1}{a_i} \sum_{j \in \mathcal{N}_i} w_{ij} (f(x_i) - f(x_j)), \quad (5.1)$$

where w_{ij} are weights, a_i are normalization coefficients, and \mathcal{N}_i denotes a local neighborhood of point i . Different discretizations of the Laplace-Beltrami operator can be cast into this form by appropriate definition of the above constants. For shapes represented as triangular meshes, a widely-used method is the *cotangent scheme*, which preserves many important properties of the continuous Laplace-Beltrami operator, such as positive semi-definiteness, symmetry, and locality [60]. Yet, in general, the cotangent scheme does not converge to the continuous Laplace-Beltrami operator, in the sense that the solution of the discrete eigenproblem does not converge to the continuous one (pointwise convergence exists if the triangulation and sampling satisfy certain conditions [63]).

Belkin *et al.* [6] proposed a discretization which is convergent without the restrictions on “good” triangulation required by the cotangent scheme. In this scheme, \mathcal{N}_i is chosen to be the entire sampling $\{x_1, \dots, x_N\}$, $a_i = \frac{1}{4\pi\rho^2}$, and $w_{ij} = S_j e^{-\|\xi(x_i) - \xi(x_j)\|^2 / 4\rho}$, where ρ is a parameter. In the case of a Euclidean colorspace, w_{ij} can be written explicitly as

$$w_{ij} = S_j \exp \left\{ -\frac{\|\xi_g(x_i) - \xi_g(x_j)\|^2}{4\rho} - \frac{\|\xi_p(x_i) - \xi_p(x_j)\|^2}{4\sigma} \right\} \quad (5.2)$$

where $\sigma = \rho/\eta^2$, which resembles the weights used in the *bilateral filter* [58]. Experimental results also show that this operator produces accurate approximation of the Laplace-Beltrami operator under various conditions, such as noisy data input and different sampling [6, 55].

5.2. Heat kernel computation

In matrix notation, equation (5.1) can be written as $\hat{\Delta}f = A^{-1}Wf$, where $A = \text{diag}(a_i)$ and $W = \text{diag}(\sum_{l \neq i} w_{il}) - (w_{ij})$. The eigenvalue problem $\hat{\Delta}\Phi = \Lambda\Phi$ is equivalent to the generalized symmetric eigenvalue problem $W\Phi = \Lambda A\Phi$, where $\Lambda = \text{diag}(\lambda_0, \dots, \lambda_K)$ is the diagonal matrix of the first K eigenvalues, and $\Phi = (\phi_0, \dots, \phi_K)$ is the matrix of the eigenvectors stacked as columns. Since typically W is sparse, this problem can be efficiently solved numerically.

Heat kernels can be approximated by taking the first largest eigenvalues and the corresponding eigenfunctions in (2.3). Since the coefficients in the expansion of k_t decay as $\mathcal{O}(e^{-t})$, typically a few eigenvalues (K in the range of 10 to 100) are required.

6. Results

In this section, we show the application of the proposed framework to retrieval of textured shapes. We compare two approaches: bags of local features and distributions of diffusion distances.

6.1. Data

In order to evaluate the proposed method, we used the SHREC 2010 robust large-scale shape retrieval benchmark methodology [11]. The query set consisted of 270 real-world human shapes from 5 classes acquired by a 3D scanner with real geometric transformations and simulated photometric transformations of different types and strengths, totalling in 54 instances per shape (Figure 9). Geometric transformations were divided into *isometry+topology* (real articulations and topological changes due to acquisition imperfections), and *partiality* (occlusions and addition of clutter such as the red ball in Figure 9). Photometric transformations included *contrast* (increase and decrease by scaling of the L channel), *brightness* (brighten and darken by shift of the L channel), *hue* (shift in the a channel), *saturation* (saturation and desaturation by scaling of the a, b channels), and *color noise* (additive Gaussian noise in all channels). *Mixed* transformations included isometry+topology transformations in combination with two randomly selected photometric transformations. In each class, the transformation appeared in five different versions numbered 1–5 corresponding to the transformation strength levels. One shape of each of the five classes was added to the queried corpus in addition to other 75 shapes used as clutter (Figure 10).

6.2. Evaluation methodology

Retrieval was performed by matching 270 transformed queries to the 75 null shapes. Each query had exactly one correct corresponding null shape in the dataset. Performance was evaluated in three ways:

First, we used the precision-recall characteristic. *Precision* $P(r)$ is defined as the percentage of relevant shapes in the first r top-ranked retrieved shapes. *Mean average precision* (mAP), defined as $mAP = \sum_r P(r) \cdot rel(r)$, where $rel(r)$ is the relevance of a given rank, was used as a single measure of performance. Intuitively, mAP is interpreted as the area below the precision-recall curve. Ideal retrieval performance results in first relevant match with mAP=100%. Performance results were broken down according to transformation class and strength.

Second, we used the receiver operation characteristic (ROC), visualizing the tradeoff between the *false positive* (FPR) and *true positive* (TPR) rates. Given a distance d between the query and database shapes and some global threshold τ , FPR is computed as the percentage of shapes belonging to a class different from the query yet having $d < \tau$; TPR is computed as the percentage of same-class shapes with $d < \tau$.

Finally, we also evaluated qualitatively the first matches produced by different methods.

6.3. Local descriptors

In the first experiment, we compared the performance of local geometric and photometric descriptors in the bag of features framework. As a purely geometric local descriptor, we used bags of features based on HKS according to [15]; as the joint photometric-geometric descriptors, we used bags of features computed with the MeshHOG [66] and the proposed cHKS descriptor. For reference, we also show the performance of a purely photometric shape descriptor (color histogram).

For the computation of the bag of features descriptors, we used the Shape Google framework with most of the settings as proposed in [15]. More specifically, HKS were computed at six scales ($t = 1024, 1351.2, 1782.9, 2352.5, \text{ and } 4096$). Soft vector quantization was applied with variance taken as twice the median of all distances between cluster centers. Approximate nearest neighbor method [2] was used for vector quantization. The Laplace-Beltrami operator discretization was computed using the Mesh-Laplace scheme [5] with scale parameter $\rho = 2$. Heat kernels were approximated using the first 200 eigenpairs of the discrete Laplacian. The MeshHOG descriptor was computed at prominent feature points (typically 100-2000 per shape), detected using the MeshDOG detector [66]. The vocabulary size in all the cases was set to 48.

All bags of features were compared using the L_2 distance. The multiscale distance (4.4) with $H = \{0, 0.05, 0.1\}$ was used to compute the similarity between cHKS BOFs in a way that does not introduce ambiguity of the joint metric.

Tables 1–4 and Figure 12 summarize the results of the first experiments. The geometry-only descriptor (HKS) [15] is invariant to photometric transformations, but is somewhat sensitive to topological noise and missing parts (Table 1). On the other hand, the color-only descriptor works well only for geometric transformations that do not change the shape color. Photometric transformations, however, make such a descriptor almost useless (Table 2). MeshHOG is almost invariant to photometric transformations being based on texture gradients, but is sensitive to color noise (Table 3). The fusion of the geometric and photometric data using our approach (Table 4) achieves nearly perfect retrieval for mixed and photometric transformations and outperforms other approaches.

Figure 11 visualizes a few examples of the retrieved shapes in the first experiment ordered by relevance, which is inversely proportional to the distance from the query shape.

6.4. Global descriptors

In the second experiment, we compared shape retrieval approaches that use distance distribution as shape descriptors. Two methods were compared: pure geometric and multiscale joint geometric-photometric distances. In the former, we used average of diffusion distances

$$d(x, x') = \frac{1}{|T|} \sum_{t \in T} d_t(x, x'), \quad (6.1)$$

computed at two scales, $T = \{1024, 4096\}$. In the latter, the joint diffusion distances at each time were also computed at multiple scales η of the photometric component using

Transform.	Strength				
	1	≤ 2	≤ 3	≤ 4	≤ 5
<i>Isom+Topo</i>	100.00	100.00	96.67	95.00	90.00
<i>Partial</i>	66.67	60.42	63.89	63.28	63.63
<i>Contrast</i>	100.00	100.00	100.00	100.00	100.00
<i>Brightness</i>	100.00	100.00	100.00	100.00	100.00
<i>Hue</i>	100.00	100.00	100.00	100.00	100.00
<i>Saturation</i>	100.00	100.00	100.00	100.00	100.00
<i>Noise</i>	100.00	100.00	100.00	100.00	100.00
<i>Mixed</i>	90.00	95.00	93.33	95.00	96.00

Table 1: Performance (mAP in %) of BOFs with purely geometric HKS descriptors.

Transform.	Strength				
	1	≤ 2	≤ 3	≤ 4	≤ 5
<i>Isom+Topo</i>	100.00	100.00	100.00	100.00	100.00
<i>Partial</i>	100.00	100.00	100.00	100.00	100.00
<i>Contrast</i>	100.00	90.83	80.30	71.88	63.95
<i>Brightness</i>	88.33	80.56	65.56	53.21	44.81
<i>Hue</i>	11.35	8.38	6.81	6.05	5.49
<i>Saturation</i>	17.47	14.57	12.18	10.67	9.74
<i>Noise</i>	100.00	100.00	93.33	85.00	74.70
<i>Mixed</i>	28.07	25.99	20.31	17.62	15.38

Table 2: Performance (mAP in %) of color histograms.

Transform.	Strength				
	1	≤ 2	≤ 3	≤ 4	≤ 5
<i>Isom+Topo</i>	100.00	95.00	96.67	94.17	95.33
<i>Partial</i>	75.00	61.15	69.93	68.28	68.79
<i>Contrast</i>	100.00	100.00	100.00	98.33	94.17
<i>Brightness</i>	100.00	100.00	100.00	100.00	99.00
<i>Hue</i>	100.00	100.00	100.00	100.00	100.00
<i>Saturation</i>	100.00	100.00	100.00	98.75	99.00
<i>Noise</i>	100.00	100.00	88.89	83.33	78.33
<i>Mixed</i>	100.00	100.00	100.00	93.33	83.40

Table 3: Performance (mAP in %) of BOFs with MeshHOG descriptor.

equation (4.6) with $H = \{0, 0.1, 0.2\}$. For the computation of distributions, the shapes were subsampled at 2500 points using the farthest point sampling algorithm [25].

Tables 5–6 summarize the retrieval performance. Both descriptors appear insensitive to photometric transformations. The joint distance has superior performance in pure geometric and mixed transformations. We conclude that the use of non-zero weight for the color component adds discriminativity to the distance distribution descriptor, while being

Transform.	Strength				
	1	≤ 2	≤ 3	≤ 4	≤ 5
<i>Isom+Topo</i>	100.00	100.00	96.67	97.50	94.00
<i>Partial</i>	68.75	68.13	69.03	67.40	67.13
<i>Contrast</i>	100.00	100.00	100.00	100.00	100.00
<i>Brightness</i>	100.00	100.00	100.00	100.00	100.00
<i>Hue</i>	100.00	100.00	100.00	100.00	100.00
<i>Saturation</i>	100.00	100.00	100.00	100.00	100.00
<i>Noise</i>	100.00	100.00	100.00	100.00	100.00
<i>Mixed</i>	100.00	100.00	96.67	97.50	98.00

Table 4: Performance (mAP in %) of BOFs with multiscale cHKS descriptors.

still robust under photometric transformations. Figure 13 visualizes a few examples of the retrieved shapes in the second experiment.

Transform.	Strength				
	1	≤ 2	≤ 3	≤ 4	≤ 5
<i>Isom+Topo</i>	80.00	90.00	88.89	86.67	89.33
<i>Partial</i>	56.25	65.62	61.61	58.71	61.13
<i>Contrast</i>	100.00	100.00	100.00	100.00	100.00
<i>Brightness</i>	100.00	100.00	100.00	100.00	100.00
<i>Hue</i>	100.00	100.00	100.00	100.00	100.00
<i>Saturation</i>	100.00	100.00	100.00	100.00	100.00
<i>Noise</i>	100.00	100.00	100.00	100.00	100.00
<i>Mixed</i>	66.67	73.33	78.89	81.67	81.33

Table 5: Performance (mAP in %) of pure geometric spectral shape distance ($\eta = 0$).

Transform.	Strength				
	1	≤ 2	≤ 3	≤ 4	≤ 5
<i>Isom+Topo</i>	100.00	100.00	100.00	100.00	100.00
<i>Partial</i>	62.50	72.92	65.97	62.50	67.50
<i>Contrast</i>	100.00	100.00	100.00	100.00	100.00
<i>Brightness</i>	100.00	100.00	100.00	100.00	100.00
<i>Hue</i>	100.00	100.00	100.00	100.00	100.00
<i>Saturation</i>	100.00	100.00	100.00	100.00	100.00
<i>Noise</i>	100.00	100.00	100.00	100.00	100.00
<i>Mixed</i>	100.00	93.33	95.56	96.67	93.70

Table 6: Performance of (mAP in %) of the multiscale joint geometric-photometric spectral distance.

7. Conclusions

In this paper, we explored a way to fuse geometric and photometric information in the construction of shape descriptors. Our approach is based on heat propagation on a manifold embedded into a combined geometry-color space. Such diffusion processes capture both geometric and photometric information and give rise to local and global diffusion geometry (heat kernels and diffusion distances), which can be used as informative shape descriptors. The choice of the metric in the joint geometric-photometric space gives rise to different invariance properties both to geometric and photometric transformations. We showed experimentally that the proposed descriptors outperform other geometry-only and photometry-only descriptors, as well as state-of-the-art joint geometric-photometric descriptors. In the future, it would be important to formally characterize the isometry group induced by the joint metric in order to understand the invariant properties of the proposed diffusion geometry, and possibly design application-specific invariant descriptors.

Acknowledgments

The author would like to thank the referees for the helpful suggestions. This work has been supported in part by the Israeli Science Foundation grant 615/11, the German-Israeli Foundation grant 2269/2010, and the Swiss High Performance and High Productivity Computing (HP2C) grant.

References

- [1] J. Amores, N. Sebe, and P. Radeva. Context-based object-class recognition and retrieval by generalized correlograms. *Trans. PAMI*, 29(10):1818–1833, 2007.
- [2] S. Arya, D. M. Mount, N. S. Netanyahu, R. Silverman, and A. Y. Wu. An optimal algorithm for approximate nearest neighbor searching. *J. ACM*, 45:891–923, 1998.
- [3] J. Assfalg, M. Bertini, A.D. Bimbo, and P. Pala. Content-based retrieval of 3-d objects using spin image signatures. *Multimedia, IEEE Transactions on*, 9(3):589–599, apr. 2007.
- [4] M. Aubry, U. Schlickewei, and D. Cremers. The wave kernel signature—a quantum mechanical approach to shape analysis. In *Proc. CVPR*, 2011.
- [5] M. Belkin, J. Sun, and Y. Wang. Constructing Laplace operator from point clouds in Rd. In *Proc. Symp. Discrete Algorithms*, pages 1031–1040, 2009.
- [6] M. Belkin, J. Sun, and Y. Wang. Discrete Laplace operator on meshed surfaces. In *Proc. Symp. Computational Geometry*, pages 278–287, 2009.
- [7] P. Bérard, G. Besson, and S. Gallot. Embedding Riemannian manifolds by their heat kernel. *Geometric and Functional Analysis*, 4(4):373–398, 1994.
- [8] A. M. Bronstein. Spectral descriptors of deformable shapes. *IEEE Trans. Pattern Analysis and Machine Intelligence*.
- [9] A. M. Bronstein, M. M. Bronstein, U. Castellani, A. Dubrovina, L. J. Guibas, R. P. Horaud, R. Kimmel, D. Knossow, E. von Lavante, D. Mateus, et al. SHREC 2010: robust correspondence benchmark. 2010.

- [10] A. M. Bronstein, M. M. Bronstein, U. Castellani, B. Falcidieno, A. Fusiello, A. Godil, L. J. Guibas, I. Kokkinos, Z. Lian, M. Ovsjanikov, et al. SHREC 2010: robust large-scale shape retrieval benchmark. 2010.
- [11] A. M. Bronstein, M. M. Bronstein, U. Castellani, B. Falcidieno, A. Fusiello, A. Godil, L. J. Guibas, I. Kokkinos, Z. Lian, M. Ovsjanikov, G. Patané, M. Spagnuolo, and R. Toldo. Shrec 2010: robust large-scale shape retrieval benchmark. In *Proc. 3DOR*, 2010.
- [12] A. M. Bronstein, M. M. Bronstein, and R. Kimmel. Efficient computation of isometry-invariant distances between surfaces. *SIAM J. Scientific Computing*, 28(5):1812–1836, 2006.
- [13] A. M. Bronstein, M. M. Bronstein, and R. Kimmel. Generalized multidimensional scaling: a framework for isometry-invariant partial surface matching. *Proc. National Academy of Science (PNAS)*, 103(5):1168–1172, 2006.
- [14] A. M. Bronstein, M. M. Bronstein, R. Kimmel, M. Mahmoudi, and G. Sapiro. A Gromov-Hausdorff framework with diffusion geometry for topologically-robust non-rigid shape matching. 89(2-3):266–286, 2010.
- [15] A. M. Bronstein, M. M. Bronstein, M. Ovsjanikov, and L. J. Guibas. Shape google: a computer vision approach to invariant shape retrieval. In *Proc. NORDIA*, 2009.
- [16] M. M. Bronstein and A. M. Bronstein. Shape recognition with spectral distances. *Trans. PAMI*, 33(5):1065, 2011.
- [17] M. M. Bronstein and I. Kokkinos. Scale-invariant heat kernel signatures for non-rigid shape recognition. In *Proc. CVPR*, 2010.
- [18] O. Chum, J. Philbin, J. Sivic, M. Isard, and A. Zisserman. Total recall: Automatic query expansion with a generative feature model for object retrieval. In *Proc. ICCV*, 2007.
- [19] R.R. Coifman and S. Lafon. Diffusion maps. *Applied and Computational Harmonic Analysis*, 21(1):5–30, 2006.
- [20] J. Digne, J.M. Morel, N. Audfray, and C. Mehdi-Souzani. The Level Set Tree on Meshes. 2010.
- [21] A. Elad and R. Kimmel. On bending invariant signatures for surfaces. *IEEE Trans. Pattern Analysis and Machine Intelligence*, pages 1285–1311, 2003.
- [22] K. Gebal, J.A. Bærentzen, H. Aanæs, and R. Larsen. Shape analysis using the auto diffusion function. In *Computer Graphics Forum*, volume 28, pages 1405–1413, 2009.
- [23] N. Gelfand, N. J. Mitra, L. J. Guibas, and H. Pottmann. Robust global registration. In *Proc. SGP*, 2005.
- [24] A.B. Hamza and H. Krim. Geodesic object representation and recognition. In *Discrete Geometry for Computer Imagery*, pages 378–387, 2003.
- [25] D. S. Hochbaum and D. B. Shmoys. A best possible heuristic for the k-center problem. *Mathematics of operations research*, 10(2):180–184, 1985.
- [26] A. K. Jain. *Fundamentals of digital image processing*. Prentice-Hall information and system sciences series. Prentice Hall, 1989.
- [27] P. W. Jones, M. Maggioni, and R. Schul. Manifold parametrizations by eigenfunctions of the Laplacian and heat kernels. *PNAS*, 105(6):1803, 2008.
- [28] M. Kazhdan, T. Funkhouser, and S. Rusinkiewicz. Rotation invariant spherical harmonic representation of 3D shape descriptors. In *Proc. SGP*, pages 156–164, 2003.
- [29] R. Kimmel, R. Malladi, and N. Sochen. Images as embedded maps and minimal surfaces: movies, color, texture, and volumetric medical images. *IJCV*, 39(2):111–129, 2000.
- [30] A. Kovnatsky, A. M. Bronstein, M. M. Bronstein, and R. Kimmel. Photometric heat kernel signatures. In *Proc. Conf. on Scale Space and Variational Methods in Computer Vision (SSVM)*, 2011.
- [31] B. Lévy. Laplace-Beltrami eigenfunctions towards an algorithm that “understands” geometry. In *Proc. Shape Modeling and Applications*, 2006.

- [32] H. Ling and D. W. Jacobs. Deformation invariant image matching. In *In ICCV*, pages 1466–1473, 2005.
- [33] R. Litman, A. M. Bronstein, and M. M. Bronstein. Diffusion-geometric maximally stable component detection in deformable shapes. *Computers & Graphics*, 35(3):549.
- [34] D. Lowe. Distinctive image features from scale-invariant keypoint. *IJCV*, 2004.
- [35] M. Mahmoudi and G. Sapiro. Three-dimensional point cloud recognition via distributions of geometric distances. *Graphical Models*, 71(1):22–31, January 2009.
- [36] J. Matas, O. Chum, M. Urban, and T. Pajdla. Robust wide-baseline stereo from maximally stable extremal regions. *Image and Vision Computing*, 22(10):761–767, 2004.
- [37] F. Mémoli and G. Sapiro. A theoretical and computational framework for isometry invariant recognition of point cloud data. IMA preprint series 1980, University of Minnesota, 2004.
- [38] F. Mémoli and G. Sapiro. A theoretical and computational framework for isometry invariant recognition of point cloud data. *Foundations of Computational Mathematics*, 5:313–346, 2005.
- [39] Facundo Mémoli. Spectral gromov-wasserstein distances for shape matching. In *Proc. NOR-DIA*, 2009.
- [40] R. Ohbuchi, K. Osada, T. Furuya, and T. Banno. Salient local visual features for shape-based 3d model retrieval. pages 93 –102, jun. 2008.
- [41] R. Osada, T. Funkhouser, B. Chazelle, and D. Dobkin. Shape distributions. *TOG*, 21(4):807–832, 2002.
- [42] M. Ovsjanikov, J. Sun, and L. J. Guibas. Global intrinsic symmetries of shapes. In *Computer Graphics Forum*, volume 27, pages 1341–1348, 2008.
- [43] X. Pan, Y. Zhang, S. Zhang, and X. Ye. Radius-normal histogram and hybrid strategy for 3d shape retrieval. pages 372 – 377, jun. 2005.
- [44] D. Raviv, A. M. Bronstein, M. M. Bronstein, and R. Kimmel. Symmetries of non-rigid shapes. In *Proc. Workshop on Non-rigid Registration and Tracking through Learning*, 2007.
- [45] D. Raviv, A. M. Bronstein, M. M. Bronstein, R. Kimmel, and N. Sochen. Affine-invariant geodesic geometry of deformable 3d shapes. *Computers & Graphics*, 35(3):692.
- [46] D. Raviv, A. M. Bronstein, M. M. Bronstein, R. Kimmel, and N. Sochen. Affine-invariant diffusion geometry for the analysis of deformable 3d shapes. In *Proc. CVPR*, 2011.
- [47] D. Raviv, M. M. Bronstein, A. M. Bronstein, and R. Kimmel. Volumetric heat kernel signatures. In *Proc. ACM Multimedia Workshop on 3D Object Retrieval*, 2010.
- [48] M. Reuter, F.-E. Wolter, and N. Peinecke. Laplace-spectra as fingerprints for shape matching. In *Proc. ACM Symp. Solid and Physical Modeling*, pages 101–106, 2005.
- [49] R. M. Rustamov. Laplace-Beltrami eigenfunctions for deformation invariant shape representation. In *Proc. SGP*, pages 225–233, 2007.
- [50] J. Sivic and A. Zisserman. Video google: A text retrieval approach to object matching in videos. In *Proc. CVPR*, 2003.
- [51] P. Skraba, M. Ovsjanikov, F. Chazal, and L. Guibas. Persistence-based segmentation of deformable shapes. In *Proc. NORDIA*, pages 45–52, 2010.
- [52] N. Sochen. On affine invariance in the beltrami framework for vision. *VLSM*, 2001.
- [53] J. Sun, M. Ovsjanikov, and L. Guibas. A Concise and Provably Informative Multi-Scale Signature Based on Heat Diffusion. In *Computer Graphics Forum*, volume 28, pages 1383–1392, 2009.
- [54] J. Sun, M. Ovsjanikov, and L. J. Guibas. A concise and provably informative multi-scale signature based on heat diffusion. In *Proc. SGP*, 2009.
- [55] K. Thangudu. Practicality of Laplace operator, 2009.
- [56] N. Thorstensen and R. Keriven. Non-rigid shape matching using geometry and photometry.

- In *Proc. CVPR*, 2009.
- [57] R. Toldo, U. Castellani, and A. Fusiello. Visual vocabulary signature for 3D object retrieval and partial matching. In *Proc. 3DOR*, 2009.
 - [58] C. Tomasi and R. Manduchi. Bilateral filtering for gray and color images. In *Proc. ICCV*, pages 839–846, 1998.
 - [59] D. V. Vranic, D. Saupe, and J. Richter. Tools for 3D-object retrieval: Karhunen-Loeve transform and spherical harmonics. In *Proc. Workshop Multimedia Signal Processing*, pages 293–298, 2001.
 - [60] M. Wardetzky, S. Mathur, F. Kälberer, and E. Grinspun. Discrete Laplace operators: no free lunch. In *Conf. Computer Graphics and Interactive Techniques*, 2008.
 - [61] C. Wu, B. Clipp, X. Li, J.-M. Frahm, and M. Pollefeys. 3d model matching with viewpoint-invariant patches (vip). pages 1 –8, jun. 2008.
 - [62] J.V. Wyngaerd. Combining texture and shape for automatic crude patch registration. pages 179 – 186, oct. 2003.
 - [63] G. Xu. Convergence of discrete Laplace-Beltrami operators over surfaces. Technical report, Institute of Computational Mathematics and Scientific/Engineering Computing, China, 2004.
 - [64] K.-J. Yoon, E. Prados, and P. Sturm. Joint estimation of shape and reflectance using multiple images with known illumination conditions, 2010.
 - [65] A. Zaharescu, E. Boyer, and R. P. Horaud. Transformesh: a topology-adaptive mesh-based approach to surface evolution, November 2007.
 - [66] A. Zaharescu, E. Boyer, K. Varanasi, and R. Horaud. Surface feature detection and description with applications to mesh matching. In *Proc. CVPR*, 2009.

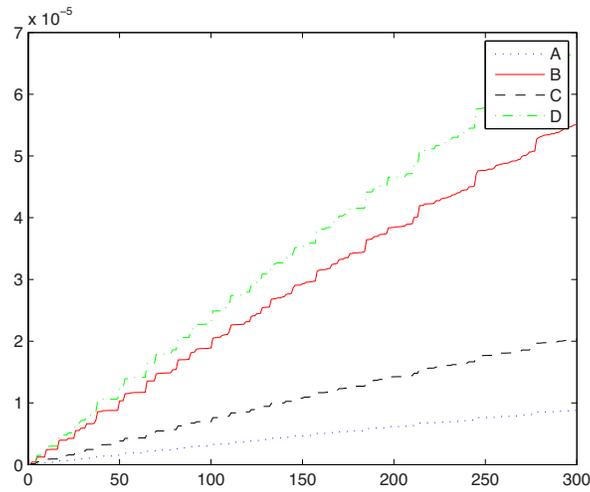
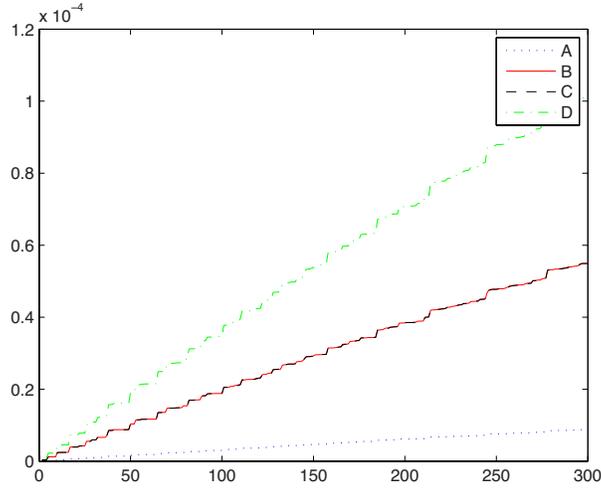
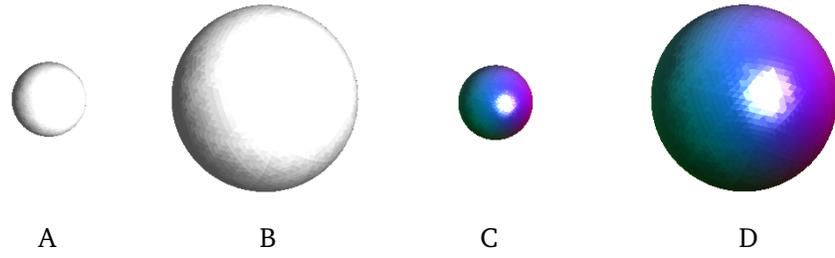


Figure 5: Example of the joint metric \hat{g} ambiguity. Top row: geometric and photometric transformations of a sphere. Sphere B is of radius 0.5 and constant texture $Lab = (1,1,1)$; A is a geometrically-transformed version of B (scaling by factor of 0.4); D is a photometrically-transformed version of B (texture set to $Lab = 0.9165 \cdot (x,y,z)$, where x,z,y are the geometric coordinates of B); C is a combination of the geometric and photometric transformations of B. Second row: first eigenvalues of the modified Laplace-Beltrami operator of all the four spheres using scaling factor $\eta = 1$, which makes the geometric and photometric transformations isometric (spheres B and C are equivalent). Third row: using a different scaling factor $\eta = 0.5$ allows to disambiguate all four shapes.



Figure 6: Top: a shape undergoing near-isometric and topological transformations (hand touches leg in the rightmost shape). Bottom: first components of the cHKS descriptor displayed as RGB ($\eta = 0.1$).

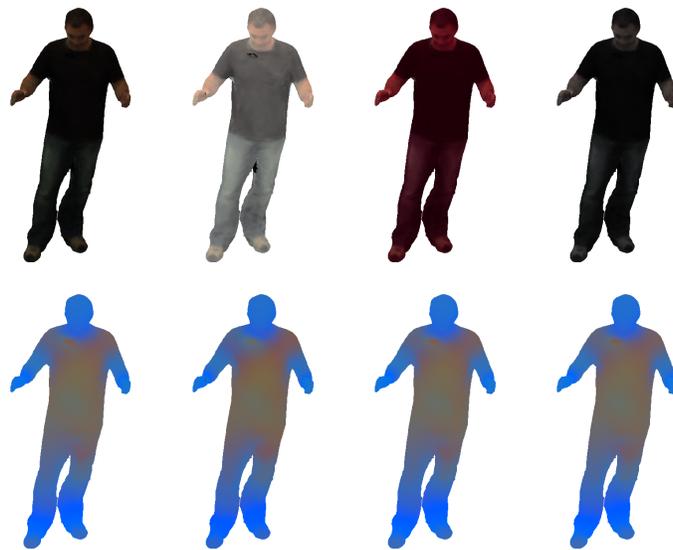


Figure 7: Top: a shape undergoing photometric transformations. Bottom: first components of the cHKS descriptor displayed as RGB ($\eta = 0.1$).

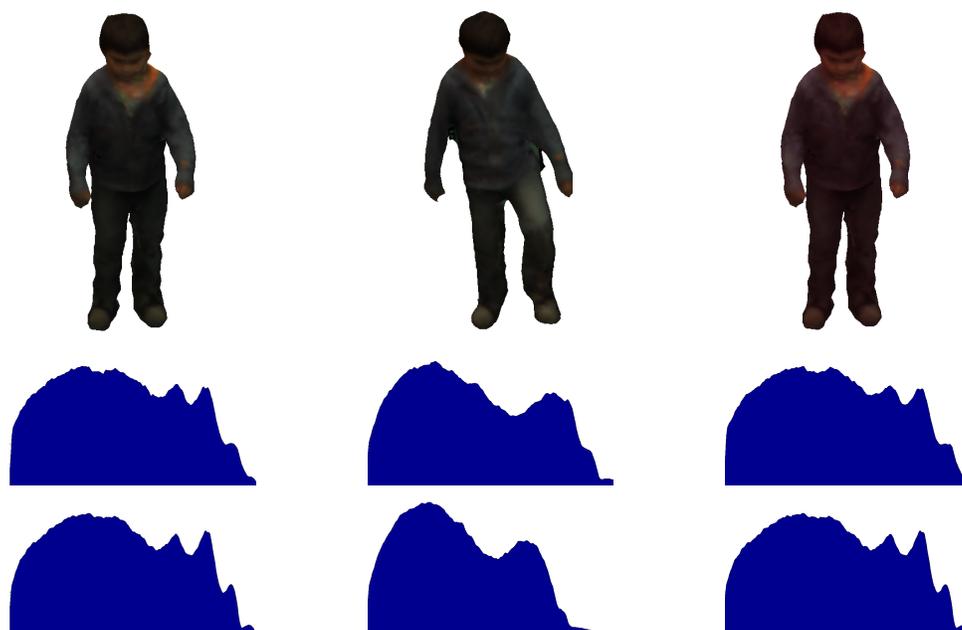


Figure 8: First row: shape undergoing geometric (second column) and photometric (third column) transformations. Second row: corresponding purely geometric diffusion distance distributions ($\eta = 0$). Third row: corresponding joint diffusion distance distributions for $\eta = 0.1$.

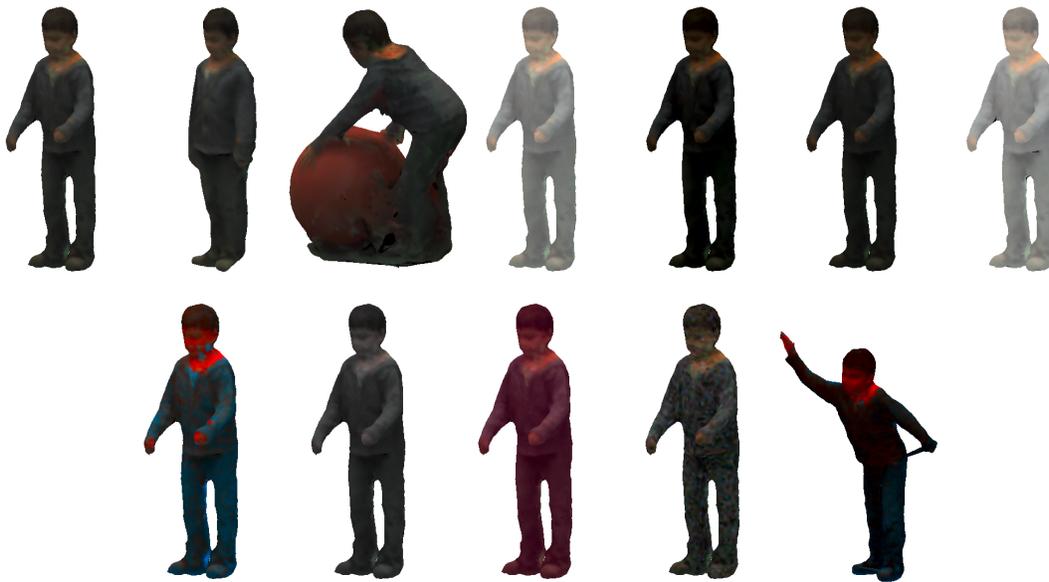


Figure 9: Examples of geometric and photometric shape transformations used as queries (shown at strength 5). First row, left to right: null, isometry+topology, partiality, two brightness transformations (brighten and darken), two contrast transformations (increase and decrease contrast). Second row, left to right: two saturation transformations (saturate and desaturate), hue, color noise, mixed.



Figure 11: Retrieval results in the first experiment. Shown are the first three matches for each method; first column shows the query shapes. Shape annotation follows the convention *shapeid.transformation.strength*; numbers below show distance from query. Only a single correct match exists in the database (marked in green), and ideally, it should be the first one.

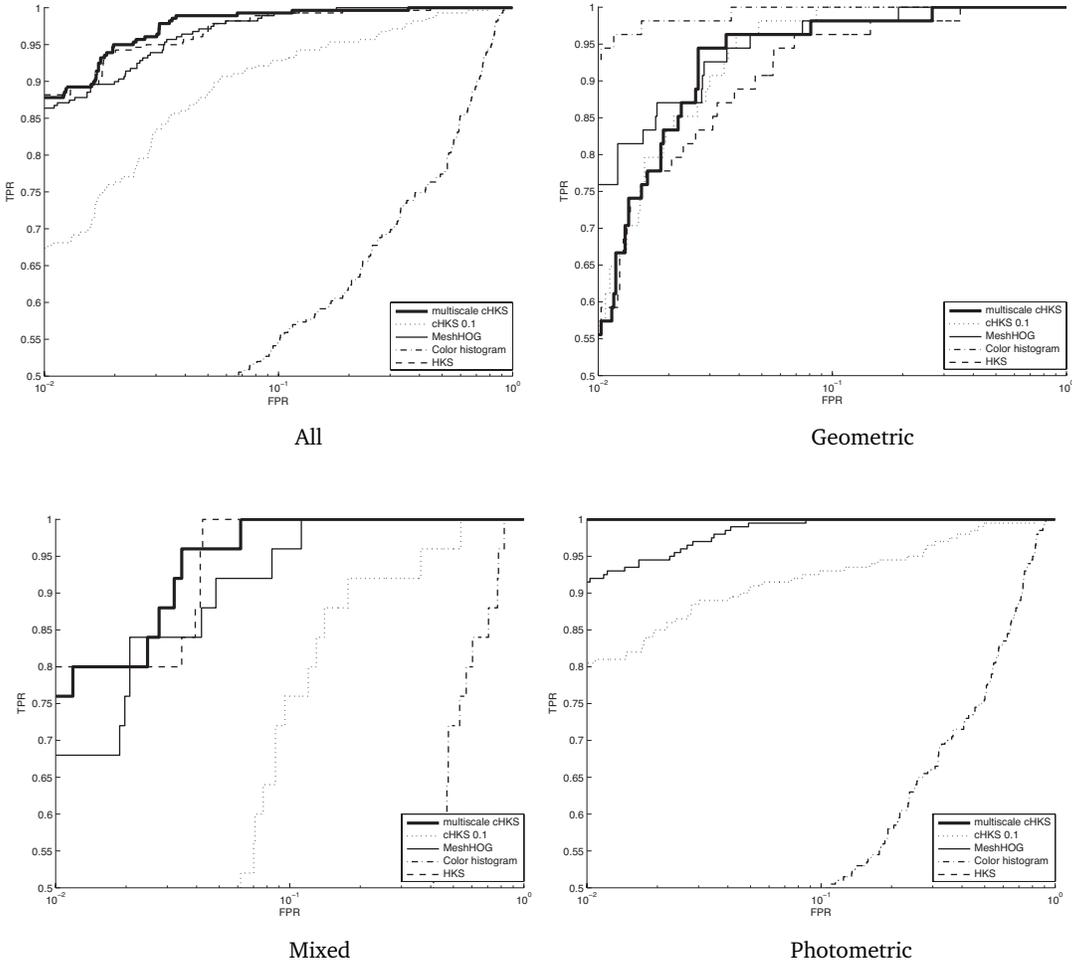


Figure 12: ROC curves comparing the performance of different methods in the first experiment. TPR is the recognition rate (the higher the better).

