



HAL
open science

Geodesic Shape Retrieval via Optimal Mass Transport

Julien Rabin, Gabriel Peyré, Laurent D. Cohen

► **To cite this version:**

Julien Rabin, Gabriel Peyré, Laurent D. Cohen. Geodesic Shape Retrieval via Optimal Mass Transport. 11th European Conference on Computer Vision, Sep 2010, Heraklion, Crete, Greece. pp.771-784, 10.1007/978-3-642-15555-0_56 . hal-00498607

HAL Id: hal-00498607

<https://hal.science/hal-00498607>

Submitted on 7 Jul 2010

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Geodesic Shape Retrieval via Optimal Mass Transport ^{*}

Julien Rabin, Gabriel Peyré, and Laurent D. Cohen

CEREMADE, Université Paris-Dauphine
{rabin,peyre,cohen}@ceremade.dauphine.fr

Abstract. This paper presents a new method for 2-D and 3-D shape retrieval based on geodesic signatures. These signatures are high dimensional statistical distributions computed by extracting several features from the set of geodesic distance maps to each point. The resulting high dimensional distributions are matched to perform retrieval using a fast approximate Wasserstein metric. This allows to propose a unifying framework for the compact description of planar shapes and 3-D surfaces.

1 Introduction

Content based 2-D and 3-D shape retrieval is an important problem in computer vision. It requires to design both representations and similarity measures to discriminate shapes from different classes, while being invariant to some deformations.

1.1 Feature-based Shape Retrieval

There is a large amount of literature on content-based retrieval using similarity measures between descriptors. In this section, a brief review is given, focusing on bending and isometric deformations (*i.e.* preserving the topology). We refer the reader to the following review papers devoted to planar shapes [1,2] and 3-D surfaces [3,4] retrieval for a complete review.

Global descriptors. Simple global features are computed using polynomial moments [5,6,7], or Fourier transform [8] (see [9] for review).

The spectrum of the Laplace Beltrami operator defines a descriptor invariant to rigid motion and to simple bendings [10]. Shape distributions [11] compute descriptors as histograms of the distribution of Euclidean distance between points on the surface. This is extended to bending invariant descriptors in [12,13,14] using geodesic distances. It is possible to replace the geodesic distance by a diffusion distance [15] computed by solving a linear Poisson PDE.

^{*} This work has been done with the support of the French “Agence Nationale de la Recherche” (ANR), under grant NatImages (ANR-08-EMER-009).

Local descriptors. Many other shape representations do not make use of a single descriptor. They rather compute similarities by matching points of interest for which local descriptors are defined. Shape context features [16] are local 2-D histograms of contours around points of interest. Geodesic shape context makes use of geodesic curves to gain bending invariance [17]. Local tomographic projections on tangent planes (spin images) [18] define a set of local descriptors.

Similarity measure. Most of the previous approaches make use of Euclidean metric, Kullback-Leibler or χ^2 distance to compare low-dimensional histogram-based descriptors in linear time. When considering high-dimensional descriptors (either histograms or discrete point clouds), another possibility is to use the Wasserstein distance [19], see *e.g.* [20,21,22].

1.2 Contributions

This paper introduces a novel framework for bending invariant recognition of shapes. We use the setting of geodesic distances on Riemannian manifolds, which unifies both planar shape and 3-D surface retrieval problems. This novel framework builds on several already known statistical descriptors, and encompasses them into a single high-dimensional descriptor. This allows us to take advantage of the richness of information available in each separate statistical measure to enhance the retrieval performance. The retrieval method is based on an approximation of the Wasserstein distance, that works directly over discrete point clouds, and can be computed with an iterative algorithm.

2 Geodesic Distances

In the following, we consider shapes as compact 2-D manifolds $\Omega \subset \mathbb{R}^s$, where $s = 2$ (planar shapes) or $s = 3$ (surfaces). Note however that our approach is generic and accommodates for domains of arbitrary dimension.

2.1 Geodesic Distance Definition

The length of a curve $\gamma : [0, 1] \rightarrow \Omega$ traced within the domain is defined as $L(\gamma) = \int_0^1 \|\gamma'(t)\| dt$. The geodesic distance between two points $x_s, x_e \in \Omega$ is the length of the shortest piecewise smooth curve joining the two points

$$d_\Omega(x_s, x_e) = \min_{\gamma(0)=x_s, \gamma(1)=x_e} L(\gamma). \quad (1)$$

The geodesic map $d_\Omega(x_i, x)$ differs significantly from the the Euclidean distance map $\|x_i - x\|$ when the shapes are non convex, as it is illustrated by Fig. 1.

A curve γ^* satisfying $d_\Omega(x_s, x_e) = L(\gamma^*)$ is called a shortest path, sometimes also referred to as a (globally minimizing) *geodesic*. Figure 1 (on the far right) shows several examples of geodesics, each time computed between a starting point $x_s \in \Omega$ (red dot) and some ending points $\{x_e\}$ lying on the boundary $\partial\Omega$ of the manifold (blue dots).




Fig. 1. Left and center: comparison of Euclidean and geodesic distances inside a 2-D shape. Right: display of geodesic curves.

2.2 Geodesic Distance Computation

Geodesic distance within a planar shape. Given some starting point x_s , the geodesic distance map $U_{x_s}(x) = d_{\Omega}(x_s, x)$ can be shown to be the unique viscosity solution of the following non-linear PDE,

$$\forall x \in \Omega, \quad \|\nabla U_{x_s}(x)\| = 1 \quad \text{and} \quad U_{x_s}(x_s) = 0. \quad (2)$$

where the derivative should be understood in a weak sense at points along the medial axis of x_s where U_{x_s} is not smooth.

The PDE (2) can be discretized with upwind finite difference. The resulting discrete equation can be and solved in $O(N \log(N))$ operations using the Fast Marching algorithm [23,24]. This algorithm performs a front propagation within the shape, as displayed on Fig. 2.




Fig. 2. Fast Marching propagation inside a 2-D shape.

Geodesic distance on a 3-D surface. If the surface Ω is parametrized using $\psi : V \subset [0, 1]^2 \mapsto \Omega$, then one can prove that the distance map

$$\forall x \in V, \quad U_{x_s}(x) = d_{\Omega}(x_s, \psi(x))$$

satisfies an anisotropic Eikonal PDE

$$\forall x \in V, \quad \|\nabla U_{x_s}(x)\|_{T_x^{-1}} = 1, \quad \text{and} \quad U_{x_s}(\psi(x_s)) = 0, \quad (3)$$

$$\text{where } T_x = \left(\left\langle \frac{\partial \psi}{\partial x_i}, \frac{\partial \psi}{\partial x_j} \right\rangle \right)_{0 \leq i, j \leq 1} \quad \text{and} \quad \|x\|_A = \sum_{0 \leq i, j \leq 1} A_{i,j} x_i x_j.$$

This equation (3) extends to surfaces of arbitrary topology using several charts that parametrize locally the surface.

The Eikonal equation (3) can be discretized on 3-D meshes. In the case of mesh with no obtuse angle, the discrete equation can be solved in $O(N \log(N))$ operations [25]. For general meshes, the resolution requires more advanced schemes, see for instance [26]. An example of a Fast Marching propagation from a set of starting points on a 3-D shape is given on Fig. 3.




Fig. 3. Example of Fast Marching propagation on a triangulated mesh.

3 Geodesic Descriptors

Similarity measures between shapes are computed by extracting global or local features $\varphi(\Omega)$, and then performing some comparison between the resulting descriptors.

An important goal in designing a similarity measure is to achieve invariance to some class \mathcal{R} of deformations. This requires that the descriptors are invariant, so that $\varphi(R\Omega) = \varphi(\Omega)$ for any $R \in \mathcal{R}$.

This section details a class of geodesic descriptors that are invariant under geodesic isometries, and quasi-invariant to shape articulations and bendings. This is especially relevant to perform robust retrieval on articulated shapes, such as animal or human with varying poses.

3.1 Local Descriptors

Geodesic distance distributions. To design features invariant to bendings and articulations, we consider, for each point $x \in \mathcal{S} \subset \Omega$, the set $\{d_\Omega(x, y)\}_{y \in \mathcal{E} \subset \Omega} \subset \mathbb{R}^+$ of distances to a subset $\mathcal{E} \subset \Omega$. This set of distances should be thought as being a 1-D distribution of values in \mathbb{R}^+ .

For numerical applications, the set \mathcal{S} is a discrete sub-sampling of the manifold computed as described in Sect. 3.2. The set \mathcal{E} used to compute the distributions can be defined depending on the application. In our numerical examples,

we choose $\mathcal{E} = \partial\Omega$ to be the boundary of the manifold for 2-D shapes, and $\mathcal{E} = \Omega$ for 3-D surfaces.

Figure 4 shows examples of geodesic distance distribution, conveniently displayed using 1-D histograms.

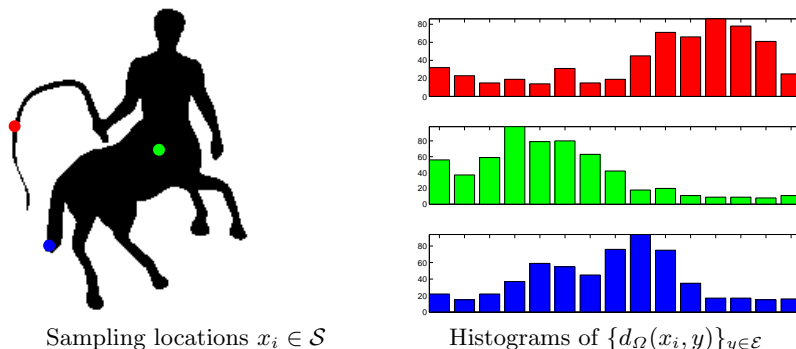


Fig. 4. Histogram of the distribution of the geodesic distance to several points.

Geodesic quantile measures. The whole set of distances $\{d_\Omega(x, y), y \in \mathcal{E}, x \in \mathcal{S}\}$ is too large to be used for retrieval applications. To achieve dimensionality reduction, we retain only a few statistical measures out of this distribution of values. This article considers quantiles statistical measures $Q_x(\alpha)$ defined as, for all $\alpha \in [0, 1]$,

$$\forall x \in \mathcal{S}, \quad Q_x(\alpha) = F_x^{-1}(\alpha) = \max\{\delta \in \mathbb{R}^+, F_x(\delta) \leq \alpha\} \quad (4)$$

where F_x is the cumulative distribution function of the set $\{d_\Omega(x, y), y \in \mathcal{E}\}$, and F_x^{-1} is its pseudo-inverse.

Observe that $Q_x(0)$ is the minimum geodesic distance between x and \mathcal{E} , while $Q_x(1/2)$ is the median distance. The maximum distance $Q_x(1)$ is also known as the *eccentricity*, and has been used for 2-D shapes [13] and surfaces [14] retrieval. Other statistical measures can be retained as well. For instance, the mean of the distance $\int_{\mathcal{E}} d_\Omega(x, y) dy$ is used in [12] to perform surface retrieval.

Geodesic local descriptors. At each location $x \in \mathcal{S}$, the local descriptor $p_x \in \mathbb{R}^d$ is a vector of d quantiles

$$p_x = (Q_x(\alpha_\ell))_{1 \leq \ell \leq d} \in \mathbb{R}^d,$$

where $0 \leq \alpha_\ell \leq 1$ are equi-spaced values. Figure 5 displays each of the $d = 3$ components $Q_x(\alpha_\ell)$ of p_x as a function of $x \in \Omega$, with $\alpha_\ell \in \{0, \frac{1}{2}, 1\}$.

The key feature of our approach, that makes it significantly different from these previous works is that it uses several statistical measures, and thus builds high dimensional descriptors.

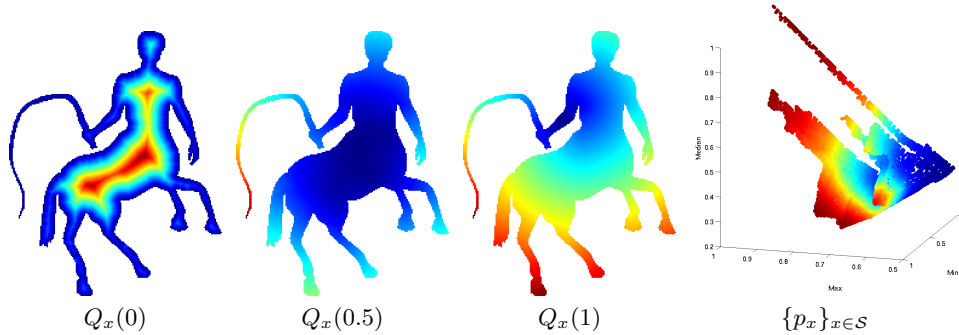


Fig. 5. Display of $x \mapsto Q_x(\alpha_\ell)$ for several $\alpha_\ell \in \{0, \frac{1}{2}, 1\}$ and of the corresponding 3-D distribution $\{p_x = (Q_x(\alpha_\ell))_{1 \leq \ell \leq 3}\}_{x \in S} \subset \mathbb{R}^3$.

3.2 Global Descriptors

The local descriptors p_x are sampled on a set $\mathcal{S} \subset \Omega$ to obtain a global descriptor that characterizes the shape.

Farthest Point Sampling. Estimating the full set $\{p_x\}_{x \in \Omega}$ of descriptors is computationally intractable, and one thus needs to compute a sub-sampling $\mathcal{S} = \{x_i\}_{i \in I}$ of n points on the manifold, where $I = \{0, \dots, n-1\}$. To perform a uniform sampling of the manifold, we use the farthest point sampling strategy. It corresponds to a greedy scheme, originally introduced in [27], and extended to geodesic distances on manifolds for surface remeshing in [28].

The initial point $x_0 \in \Omega$ is sampled at random. Given a set of k points $\{x_0, \dots, x_{k-1}\}$, the next point is computed as

$$x_k = \operatorname{argmax}_{x \in \Omega} \min_{0 \leq i < k} d_\Omega(x_i, x).$$

Once this new point x_k has been computed, the set $\{d_\Omega(x_k, x)\}_{x \in \mathcal{E}}$ of geodesic distances is computed in $O(n \log(n))$ operations, and the geodesic descriptor $p_{x_k} \in \mathbb{R}^d$ is obtained by computing the quantiles (4) from these distances. This process of iteratively adding the furthest point to the set \mathcal{S} is continued until a given number of points n is reached. Examples of this farthest point sampling method on a 2-D shape and a surface are shown in Fig. 6.

Global descriptor as a point cloud. The global descriptor is then defined as a uniform sampling of the local descriptors

$$\varphi(\Omega) = \{p_{x_i}\}_{i \in I} \subset \mathbb{R}^d,$$

and is thus a cloud of n points in \mathbb{R}^d .

This point cloud $\varphi(\Omega)$ should be thought as being drawn from a probability distribution. Each shape has its own distribution, that is invariant under isometric bending of the shape.

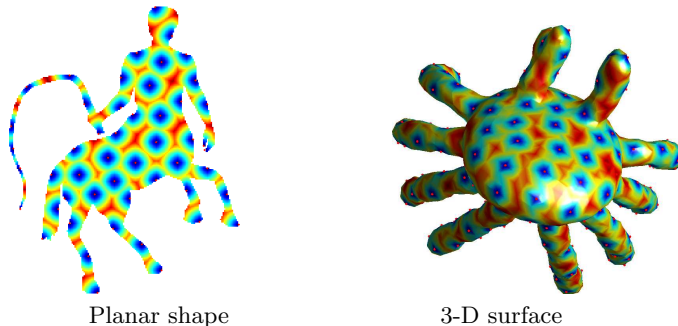


Fig. 6. *Illustration of the farthest point sampling strategy for a 2-D shape and a 3-D surface. For each case $|\mathcal{S}| = n = 100$ points are sampled (red dots). The geodesic distances between each $x \in \Omega$ and these points are plotted as a colormap.*

An alternative representation is to compute the d -dimensional histogram of the distribution. We prefer in this paper to work directly using discretized point cloud, because it offers a more precise matching.

The resulting global descriptor $\varphi(\Omega)$ is invariant to isometric deformations, in the sense that $\varphi(\Omega) = \varphi(R\Omega)$ if R is a deformation of Ω that maintains geodesic distances. More generally, if R does not modify too much the distances, meaning

$$\forall (x, y) \in \Omega^2, \quad d_\Omega(x, y) \approx d_{R\Omega}(Rx, Ry), \quad (5)$$

then $\varphi(\Omega) \approx \varphi(R\Omega)$. This is the case for bending deformations and articulations, see [17]. Observe that geodesic distances $\{d_\Omega(x, y)\}_{x, y \in \Omega}$ have to be normalized (according to maximum distance) to achieve invariance to scaling.

4 Optimal Transport Retrieval

Our shape retrieval method uses a similarity measure that compares the geodesic descriptors $\varphi(\Omega)$ using the Wasserstein metric related to the Monge-Kantorovich optimal transport problem (see [19] for an in-depth study).

4.1 Similarity Measure

In classical settings, shapes are generally represented by histogram-based descriptors and thus compared with L^p distances or the Kullback-Leibler divergence. In our setting, the descriptors are high dimensional discrete distributions, and we propose to use the Wasserstein distance [19] which is well adapted to compare statistical distributions [20], and is known to be more robust than Hausdorff distance [29]. Our similarity measure is thus defined as

$$\Delta(\Omega_0, \Omega_1) = W(\varphi(\Omega_0), \varphi(\Omega_1)),$$

where $W(X, Y)$ is the Wasserstein distance between two point clouds $X, Y \subset \mathbb{R}^d$, defined in the next section. Since our geodesic descriptors satisfy the approximate invariance (5) for bendings and articulations $R \in \mathcal{R}$, our similarity measure satisfies $\Delta(\Omega_0, R(\Omega_0)) \approx 0$.

4.2 Wasserstein Distance

Given two point clouds $X, Y \subset \mathbb{R}^d$ of n elements, the L^2 -Wasserstein distance is defined as

$$W(X, Y)^2 = \min_{\sigma \in \Sigma_n} \sum_{i \in I} \|X_i - Y_{\sigma(i)}\|^2, \quad (6)$$

where Σ_n is the set of all permutations of n elements and $I = \{0, \dots, n-1\}$. Our framework extends to arbitrary strictly convex cost such as L^p -metric for $p > 1$. Computing this distance boils down to estimate the optimal assignment $i \mapsto \sigma^*(i)$ minimizing Formula (6). This can be computed exactly using linear programming or other dedicated algorithms in $O(n^{2.5} \log(n))$ operations [30].

One-dimensional case. It is well known that the Wasserstein assignment problem in 1-D can be easily solved in $O(n \log(n))$ operations by sorting the values [19]. Indeed, σ_X and σ_Y being two permutations such that $\{X_{\sigma_X(i)}\}_i$ and $\{Y_{\sigma_Y(i)}\}_i$ are sorted in increasing order, the optimal assignment is

$$\sigma^* = \sigma_Y \circ \sigma_X^{-1}. \quad (7)$$

4.3 Approximate Wasserstein Distance

For large point clouds in high dimension ($d \geq 2$), computing exactly (6) is too demanding. Following an idea recently introduced in [31], we propose to use the an approximate transport cost $\tilde{W}(X, Y)$ defined as

$$\tilde{W}(X, Y) = \|X - X^{(\infty)}\|, \quad (8)$$

where $X^{(\infty)}$ is computed using an iterative algorithm described in the following paragraph. Starting from $X^{(0)} = X$, this algorithm computes points clouds $\{X^{(k)}\}_k$ that progressively evolves $X^{(k)}$ toward Y , minimizing at each iteration an energy E_Y which is a sum of 1-D Wasserstein distances on the unit sphere S^{d-1} in \mathbb{R}^d :

$$E_Y(U) = \frac{1}{2} \int_{\theta \in S^d} \sum_{i \in I} \langle U_i - Y_{\sigma_\theta^*(i)}, \theta \rangle^2 d\theta, \quad (9)$$

where $\langle \cdot, \cdot \rangle$ is the L^2 -scalar product, and where σ_θ^* is the optimal 1-D assignment according to direction θ of $\{Y_i, \theta\}_i$ with $\{U_i, \theta\}_i$ following Formula (7).

Algorithm. Finding a minimum of energy (9) can be done using a classical gradient descent strategy. For numerical considerations, this energy is estimated at each iteration k from a restricted set of directions, thus resulting in a *stochastic gradient descent* scheme (see e.g. [32]), which relies on three steps:

▷ **Step 1** Define the direction set $\Psi^{(k)} \subset S^{d-1}$, a collection of vectors *randomly and uniformly sampled* on S^{d-1} . The corresponding energy $E_Y^{(k)}$ is therefore

$$E_Y^{(k)}(U) = \frac{1}{2} \sum_{\theta \in \Psi^{(k)}} \sum_{i \in I} \langle U_i - Y_{\sigma_\theta^*(i)}, \theta \rangle^2. \quad (10)$$

- ▷ **Step 2** Compute, for each direction $\theta \in \Psi^{(k)}$, the optimal 1-D assignment σ_θ^* of 1-D distribution $\{\langle Y_i, \theta \rangle\}_{i \in I}$ with $\{\langle X_i^{(k)}, \theta \rangle\}_{i \in I}$ using Formula (7).
- ▷ **Step 3** The set $\{\sigma_\theta^*\}$ being computed, update the point cloud using a Newton step with parameter $\lambda \in]0, 2[$ to minimize the energy $E_Y^{(k)}$, such that $\forall i \in I$

$$\begin{aligned} X_i^{(k+1)} &= X_i^{(k)} - \lambda \left(\nabla^2 E_Y^{(k)}(X^{(k)}) \right)^{-1} \nabla E_Y^{(k)}(X_i^{(k)}), \\ &= (1 - \lambda) \cdot X_i^{(k)} + \lambda H^{-1} \sum_{\theta \in \Psi^{(k)}} \langle Y_{\sigma_\theta^*(i)}, \theta \rangle \theta, \end{aligned} \quad (11)$$

where $H = \sum_{\theta \in \Psi^{(k)}} \theta \theta^T \in \mathbb{R}^{d \times d}$ is the Hessian matrix of $E_Y^{(k)}$ at point $X_i^{(k)}$.

Convergence. Stochastic gradient descent is known to converge if one uses a properly chosen step size $\lambda = \lambda_k$ that decays through the iterations, see [32]. In numerical simulations, we always observed convergence of $X^{(k)}$ to some $X^{(\infty)}$ using the fixed step size $\lambda = 1/|\Psi^{(k)}|$. Furthermore, $X^{(\infty)}$ is actually equal (up to a permutation) to Y , so that the algorithm computes an assignment between the two distributions –which is not necessarily optimal in the sense of Formula (6).

Implementation. In the numerical simulations of Sect. 5, we used a fixed number of $K = 100$ iterations and $|\Psi| = d$ directions, and we noticed that using more iterations does not improve significantly the retrieval results. The complexity of the proposed algorithm is therefore $O(|\Psi|Kn \log(n))$.

This method extends the algorithm proposed by [33] that makes use of an orthogonal set of direction Ψ^k , and a descent step size $\lambda = 1$. Using a smaller step, *e.g.* $\lambda = 1/|\Psi^k|$ is important to ensure the convergence of the method. Using a larger set of directions is useful to obtain a assignment that is closer to the optimal one, see [31] for more details. Observe that other approximation methods has been previously proposed in the literature, *e.g.* making use of metric embedding [21] or wavelet approximations [34].

5 Numerical Examples

Given a database of manifolds $\{\Omega_j\}_{j \in J}$, our algorithm for shape retrieval begins by computing the global signature $\varphi(\Omega_j)$ for each manifold in the database. This is performed in parallel to the farthest point sampling algorithm described in Sect. 3.2. When an input manifold Ω is queried in the database, its global signature $\varphi(\Omega)$ is computed, and the shape in the database are ordered according to the Wasserstein distance approximation $\tilde{W}(\varphi(\Omega), \varphi(\Omega_j))$.

To evaluate the retrieval performance of the proposed descriptor, two classical performances curves are displayed:

- the *average recall* curve shows the average number of correct shapes (or “true-positives”) retrieved per query among the r most similar ones in the dataset. This curve, plotted depending of the rank r , is obtained by querying every shape in the dataset (the query itself is no used to compute the score);

- the *average precision-recall* curve plots the rate of correct shapes retrieved among r as a function of the average recall rate.

In order to show the interest of considering high dimensional geodesic statistics, each performance curves are shown for two different descriptors: the aforementioned multi-dimensional descriptor and also a simple 1-D descriptor corresponding to the distribution of *eccentricity* (maximal geodesic distances).

5.1 2-D Shape Retrieval

In this setting, the domain \mathcal{E} of ending points $\{y_j\}_{j \in J}$ is the boundary $\partial\Omega$ of the manifold. 4-dimensional distributions of $n = 500$ points are used as global descriptors, considering 3 quantiles $Q_x(\alpha_j)$ (minimum, median and maximum) in addition with the mean values of geodesic distances.

We consider first the ‘‘Articulated Shapes’’ dataset of [17] (see Fig. 7(a)), a small dataset being designed to evaluate the robustness of retrieval to bending deformations. Performance curves are shown in Figs. 7(b) and 7(c) for both 4-D and 1-D descriptors. Comparison with state-of-the-art methods [17,22] is also provided in Table 1. One can see that considering several geodesic statistics at the same time enables to catch more sophisticated information on the shape while being more robust to bending deformations. Note that it is not the case when using Euclidean metric with descriptors made of 4 1-D histograms instead of the approximate Wasserstein metric with 4-D discrete distributions.


Table 1. Retrieval results on the articulated shapes dataset [17]. Scores correspond to the number of correct shapes retrieved among 40 depending on their rank.

METHOD		RANK			
DESCRIPTOR	METRIC	1 st	2 nd	3 rd	4 th
Geodesic quantile distribution (4-D)	Approx. Wasserstein	39	34	30	24
Maximal geodesic distribution (1-D)	Wasserstein	27	24	16	18
4 Geodesic quantile histogram (1-D)	Euclidean	21	15	7	11
Inner Distance Shape Context	χ^2 [17]	40	34	35	27
Inner Distance Shape Context	EMD- L_1 [22]	39	39	34	32

In order to evaluate the robustness of our approach for 2-D shape retrieval on a larger database, we consider now the MPEG-7 dataset of 1400 shapes (see Fig. 8(a)). Results are shown in Figs. 8(b) and 8(c). Again, one can see that using multi-dimensional statistics on geodesic distances yields far better results. The state-of-the-art method of [22] yields a bullseye score (average recall rate for rank $r = 40$) of 86.56%, whereas we obtain 59.7%. This can be explained by the strongly non-isometric variations of the objects in the MPEG-7 classes, which makes our representation quite inefficient for this retrieval. An important avenue for future work is to design a large benchmark of planar shapes undergoing bendings and articulation deformations, to explore the performance of our algorithm and related methods.



(a) **Articated shapes dataset of [17].** Pairs of shapes from different classes. The complete dataset is composed of 8 classes of 5 elements.




(b) Recall vs Image Rank.

(c) Average Precision-Recall.

Fig. 7. Retrieval results on the articulated shapes database [17]. Figure 7(a): database overview. Figure 7(b): Average recall rate depending on the shape rank threshold for each query shape in the dataset (the score does not include the query itself). Figure 7(c): Average Precision versus Recall.



(a) **MPEG7 dataset.** Pairs of shapes from different classes. The complete dataset is composed of 8 classes of 5 elements.



(b) Recall vs Image Rank.

(c) Average Precision-Recall.

Fig. 8. Retrieval results on MPEG7 database. Figure 8(a): database overview. Figure 8(b): Average recall rate. Figure 8(c): Average Precision versus Recall.

5.2 3-D Shape Retrieval

For surface, the domain $\mathcal{E} = \Omega$ is the whole manifold. Hence, the first order quantiles $Q_x(0)$ are discarded since they are zero, so that we handle $d = 3 -$

dimensional geodesic statistics of $n = 500$ points, considering from now *max* and *median* in addition with *mean* values of geodesic distances.

To evaluate the robustness of such global descriptor for 3-D shapes, we used the McGill dataset of 3-D articulated objects [35] (see Fig. 9(a) for an overview). Retrieval results are shown in Figs. 9(b) and 9(c). Again, it is clear that the combination of several geodesic distance characteristics achieves better retrieval results than considering only one. A comparison with state-of-the-art approach described in [36] is also provided in Figure 9, where we obtain similar results. Following the same protocol as in [36], retrieval results are given for a subset of the complete McGill Database (14 classes out of 19).

6 Conclusion

The first contribution of this paper is a generic framework to represent manifolds with statistical signatures based on geodesic distances, which are robust to bendings. The second contribution of this paper is an algorithm to compute a similarity measure between multi-dimensional joint distributions, which yields a fast approximation of the Wasserstein metric. This algorithm is applied to perform shape retrieval using our geodesic framework.


Our framework extends naturally to include additional information such as texture. One can indeed use a non-constant Riemannian metric that takes into account this information.

References


1. Veltkamp, R.C., Latecki, L.: Properties and performance of shape similarity measures. In: Proc. of Conference on Data Science and Classification. (2006)
2. Zhang, D.S., Lu, G.J.: Review of shape representation and description techniques. Pattern Recognition **37** (2004) 1–19
3. Bustos, B., Keim, D.A., Saupe, D., Schreck, T., Vranić, D.V.: Feature-based similarity search in 3D object databases. ACM Comput. Surv. **37** (2005) 345–387
4. Tangelder, J.W.H., Veltkamp, R.C.: A survey of content based 3D shape retrieval methods. Multimedia Tools Appl **39** (2008) 441–471
5. Teague, M.: Image analysis via the general theory of moments. Journal of the Optical Society of America **70** (1980) 920–930
6. Teh, C.H., Chin, R.T.: On image analysis by the methods of moments. IEEE Trans. Patt. Anal. and Mach. Intell. **10** (1988) 496–513
7. Liao, S., Pawlak, M.: On image analysis by moments. IEEE Trans. Patt. Anal. and Mach. Intell. **18** (1996) 254–266
8. Zahn, C.T., Roskies, R.Z.: Fourier descriptors for plane closed curves. IEEE Transactions on Computer **21** (1972) 269–281
9. Prokop, R.J., Reeves, A.P.: A survey of moment-based techniques for unoccluded object representation and recognition. CVGIP **54** (1992) 438–460
10. Reuter, M., Wolter, F.E., Peinecke, N.: Laplace-spectra as fingerprints for shape matching. In: Symposium on Solid and Physical Modeling. (2005) 101–106
11. Osada, R., Funkhouser, T., Chazelle, B., Dobkin, D.: Shape distributions. ACM Transactions on Graphics **21** (2002) 807–832




(a) McGill dataset of articulated and non-articulated objects [35] (respectively composed of 9 classes of 202 elements and 10 classes of 255 elements).



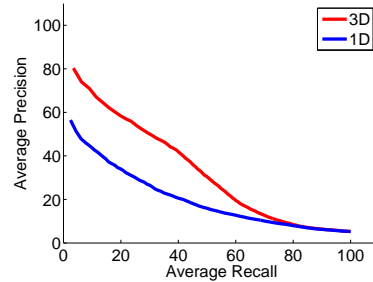
(b) Recall vs Image Rank on articulated dataset.



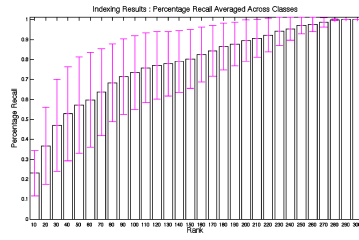
(c) Average Precision-Recall on articulated dataset.



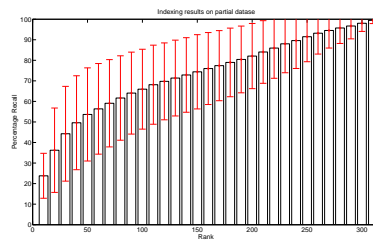
(d) Recall vs Image Rank on complete dataset.



(e) Average Precision-Recall on complete dataset.



(f) Average Recall on partial McGill database (14 classes) with [36] method.



(g) Average Recall with our approach.

Fig. 9. Retrieval results on McGill database [35]. Figure 9(a): overview of the database. Figure 9(b) and 9(d): Average recall rate on articulated and complete dataset. Figure 9(c) and 9(e): Average Precision versus Recall on articulated and complete dataset. Figures 9(f) (from [36]) and 9(g): comparison of our approach with state-of-the-art method described in [36] on partial McGill Dataset (14 classes out of 19). Results are shown as average recall rate curve (plotted in black) along with the average intra-class standard deviation (in red).

12. Ben Hamza, A., Krim, H.: Geodesic matching of triangulated surfaces. *IEEE Trans. Image Proc.* **15** (2006) 2249–2258
13. Ion, A., Peyré, G., Haxhimusa, Y., Peltier, S., Kropatsch, W., Cohen, L.: Shape matching using the geodesic eccentricity transform - a study. In: Proc. Workshop of the Austrian Association for Pattern Recognition, OCG (2007) 97–104
14. Ion, A., Artner, N., Peyré, G., López Mármol, S., Kropatsch, W., Cohen, L.: 3D shape matching by geodesic eccentricity. In: Proc. Workshop on Search in 3D, IEEE (2008)
15. Gorelick, L., Galun, M., Sharon, E., Basri, R., Brandt, A.: Shape representation and classification using the Poisson equation. *IEEE TPAMI* **28** (2006) 1991–2005
16. Belongie, S., Malik, J., Puzicha, J.: Shape matching and object recognition using shape contexts. *IEEE Trans. Patt. Anal. and Mach. Intell.* **24** (2002) 509–522
17. Ling, H., Jacobs, D.W.: Shape classification using the inner-distance. *IEEE Trans. Patt. Anal. and Mach. Intell.* **29** (2007) 286–299
18. Johnson, A., Hebert, M.: Using spin images for efficient object recognition in cluttered 3D scenes. *IEEE Trans. Patt. Anal. and Mach. Intell.* **21** (1999) 433–449
19. Villani, C.: *Topics in Optimal Transportation*. American Mathematical Society (2003)
20. Rubner, Y., Tomasi, C., Guibas, L.J.: The Earth Mover’s Distance as a Metric for Image Retrieval. *International Journal of Computer Vision* **40** (2000) 99–121
21. Grauman, K., Darrell, T.: Fast contour matching using approximate earth mover’s distance. In: Proc. of IEEE CVPR’04. (2004) I: 220–227
22. Ling, H., Okada, K.: An Efficient Earth Mover’s Distance Algorithm for Robust Histogram Comparison. *IEEE Trans. PAMI* **29** (2007) 840–853
23. Sethian, J.A.: A fast marching level set method for monotonically advancing fronts. *Proc. of the National Academy of Sciences* **93** (1996) 1591–1595
24. Tsitsiklis, J.: Efficient Algorithms for Globally Optimal Trajectories. *IEEE Trans. on Automatic Control* **40** (1995) 1528–1538
25. Kimmel, R., Sethian, J.: Computing Geodesic Paths on Manifolds. *Proc. of the National Academy of Sciences* **95** (1998) 8431–8435
26. Bornemann, F., Rasch, C.: Finite-element discretization of static Hamilton-Jacobi equations based on a local variational principle. *Comput. Visual Sci.* **9** (2006)
27. Gonzalez, T.F.: Clustering to minimize the maximum intercluster distance. *Theoretical Computer Science* **38** (1985) 293–306
28. Peyré, G., Cohen, L.D.: Geodesic remeshing using front propagation. *International Journal of Computer Vision* **69** (2006) 145–156
29. Veltkamp, R.C., Hagedoorn, M.: State of the art in shape matching. (2001) 87–119
30. Burkard, R., Dell’Amico, M., Martello, S.: *Assignment Problems*. SIAM (2009)
31. Rabin, J., Peyré, G., Delon, J., Bernot, M.: Wasserstein Barycenter and its Application to Texture Mixing. (<http://hal.archives-ouvertes.fr/hal-00476064/en/>)
32. Bottou, L.: Online algorithms and stochastic approximations. In Saad, D., ed.: *Online Learning and Neural Networks*. Cambridge University Press, Cambridge, UK (1998)
33. Pitié, F., Kokaram, A., Dahyot, R.: Automated colour grading using colour distribution transfer. *Computer Vision and Image Understanding* (2007)
34. Shirdhonkar, S., Jacobs, D.: Approximate Earth Mover’s Distance in linear time. In: Proc. CVPR’08. (2008) 1–8
35. McGill 3-D shapes dataset: (<http://www.cim.mcgill.ca/~shape/benchMark/>)
36. Siddiqi, K., Zhang, J., Macrini, D., Shokoufandeh, A., Bouix, S., Dickinson, S.: Retrieving articulated 3-D models using medial surfaces. *Mach. Vision Appl.* **19** (2008) 261–275